Economics and Philosophy, **21** (2005) 177–197
doi:10.1017/S0266267105000556

VALUE AND UNACCEPTABLE RISK

Gustaf Arrhenius
Stockholm University

Wlodek Rabinowicz
Lund University

Consider a transitive value ordering of outcomes and lotteries on outcomes, which satisfies substitutivity of equivalents and obeys "continuity for easy cases," i.e., allows compensating risks of small losses by chances of small improvements. Temkin (2001) has argued that such an ordering must also – rather counter-intuitively – allow chances of small improvements to compensate risks of huge losses. In this paper, we show that Temkin's argument is flawed but that a better proof is possible. However, it is more difficult to determine what conclusions should be drawn from this result. Contrary to what Temkin suggests, substitutivity of equivalents is a notoriously controversial principle. But even in the absence of substitutivity, the counter-intuitive conclusion is derivable from a strengthened version of continuity for easy cases. The best move, therefore, might be to question the latter principle, even in its original simple version: as we argue, continuity for easy cases gives rise to a *sorites*.

Larry Temkin has recently argued that the following four principles are mutually inconsistent: (i) sufficiently large chances of small improvements can compensate risks of small losses (continuity for easy cases); (ii) the value ordering of outcomes and lotteries on outcomes is transitive; (iii) a lottery's value remains unchanged if one of its possible outcomes is replaced by an equally good outcome or lottery (substitutivity of equivalents); (iv) no chance of a small improvement, however large, can compensate an arbitrarily small risk of a huge loss (discontinuity for the extreme case).

Temkin takes continuity for easy cases and discontinuity for the extreme case to be highly intuitive assumptions and suggests that the same applies to the principle of substitution of equivalents. The reader is thus left with the impression that it is transitivity that should be given up. In this paper, we show that Temkin's inconsistency proof is fatally flawed but that a better proof can be provided. However, it is more difficult to determine what conclusions should be drawn from this result. We consider this issue in the last section. Contrary to what Temkin suggests, substitution of equivalents is a notoriously controversial principle. But the best move, as we shall see, might be to question Temkin's very point of departure: contrary to appearances, continuity for easy cases probably is the culprit in this story.

1. THE PROBLEM

Temkin (2001) asks the reader to consider a finite series of outcomes, $o_1, \ldots, o_n$, in which each element is *slightly* worse than its immediate predecessor. The last outcome in this gradually descending sequence is supposed to be *much* worse than the one we start with. The sequence begins with a very satisfactory outcome and it ends up with something disastrous.

Temkin takes it that the standard decision-theoretic *continuity* assumption is satisfied by the outcomes in the sequence that are adjacent to each other.

> *Continuity for Easy Cases*: For all outcomes o_{i-1}, o_i, and o_{i+1}, if o_i is slightly worse than o_{i-1}, and o_{i+1} is slightly worse than o_i, there is some probability value p, $0 < p < 1$, such that $o_i \approx (o_{i-1}, p, o_{i+1})$.[1]

Here, "$\approx$" stands for "is equally good as," and "(o_{i-1}, p, o_{i+1})" denotes a binary lottery in which outcome o_{i-1} is assigned a real-valued probability p, while the remaining probability $1 - p$ is assigned to outcome o_{i+1}. Thus, the assumption states that in each triple of outcomes the second of which is slightly worse than the first and slightly better than the third, the intermediate outcome is equally good as some lottery on the outcomes that flank it from above and from below. In this sense, there is continuity in the downward movement in the sequence. If each of the steps is small, none of them involves a radical break. Given an appropriate choice of probabilities, the risk of ending up with o_{i+1} instead of o_i is compensated by the chance of an improvement from o_i to o_{i-1}.

[1] This is the assumption that Temkin actually relies on in his proof (see below). But in the discussion that precedes the proof, he instead argues for a closely related assumption that is in one way stronger and in one way weaker, namely:
In each triple of outcomes the second of which is worse than the first [slightly worse or much worse; this is the strengthening] and the third of which is only slightly worse than the second, there is a lottery on the first and the third outcome that is at least as good as [i.e., equally good as or better than; this is the weakening] the intermediate outcome.

At the same time, we are asked to suppose that the last outcome in the sequence is so much worse than the outcomes the sequence starts with that the continuity assumption does *not* hold for o_1, o_2, and o_n.

> *Discontinuity for the Extreme Case*: There is no probability value p such that $o_2 \approx (o_1, p, o_n)$.

No chance of a slight improvement from o_2 to o_1 is worth the risk, however small, of ending up with o_n instead of o_2. In fact, for that reason, not only does continuity fail for this "extreme case", but there is no lottery with o_1 and o_n as possible prizes that is *at least as good as* o_2. Letting $\preccurlyeq$ stand for "at least as good as," we can express this claim as follows:

> There is no probability value $p < 1$ such that $o_2 \preccurlyeq (o_1, p, o_n)$.

Temkin's example of a sequence $o_1, \ldots, o_n$ involves a gradually decreasing series of alternative incomes. The series starts with the agent's income exceeding $1 million per year and ends with his income being equal to zero. For each outcome o_i in the series, the agent earns in o_{i+1} just one dollar less than in o_i. Temkin takes it that our intuitions about this sequence are as follows: (i) A small risk of losing a dollar can always be compensated, by an appropriately large chance of gaining an extra dollar. (ii) But it is different with large losses. If you are well off, a chance of gaining an extra dollar, however large, is not worth the risk, however small, of losing everything.

That a chance of a small improvement is not worth the risk, however small, of a large loss is, of course, a standard objection to the unrestricted continuity assumption: it is mentioned in practically any textbook in decision theory.[2] One typical answer to this objection is that if the risk of a large loss is *very* small, say one in 10^{1000}, taking it for the sake of a small improvement may well not be unreasonable (cf. Resnik 1987: 103). This may so, but the question is rather whether *refusing* to take such a risk, however small, on the grounds that it is not worth taking, must be unreasonable. If it needn't be, the objection to continuity as a requirement of rationality still stands.

Another typical reaction is that, in our actual behavior, we do in fact regularly accept small risks of catastrophic outcomes, for the sake of small gains. We do it each time we cross a heavily trafficked street to hunt for a bargain, or when we don't bother to put on the safety belt while driving. But, again, the question is rather whether those of us who refuse to take such risks, because they are in their view not worth taking, are being unreasonable.

[2] See, for example, Luce and Raiffa (1957: 27) or Resnik (1987: 103). Temkin himself refers in this connection to Peter Vallentyne's discussion of John Broome's *Weighing Goods* in Vallentyne (1993).

https://doi.org/10.1017/S0266267105000556 Published online by Cambridge University Press

To be sure, if we *repeatedly* refuse taking microscopically small risks of this kind for the sake of small gains, then, in the long run, we may forgo *very large* gains (one dollar + one dollar + one dollar + ...) that we could have obtained at the expense of accumulated microscopically small risks of very large losses.[3] Note that these accumulated tiny risks might still add up to a very small total risk. This is a challenging objection, but a reply to it can be that what is reasonable as an isolated decision may not be reasonable as a general policy. A decision in a single case need not be irrational just because it instantiates a general policy that it would be irrational for us to adopt.

It might seem, therefore, that it is not unreasonable to refuse to take any risk of a very large loss, however small that risk might be, for the sake of a small improvement. But then it is not unreasonable to reject continuity for an extreme case like this. Temkin argues, however, that continuity for the extreme case *follows* from continuity for easy cases, provided we assume two additional principles: the Principle of Substitution of Equivalents (PSE) and the Transitivity of At-Least-as Good-As (Tr $\preccurlyeq$). As is well known, both these principles, along with the unrestricted continuity, are fundamental to expected utility theory (in the von Neumann–Morgenstern version). PSE states that for any two outcomes or lotteries, x and y, and for any lottery l with x as a possible prize, if x and y are equally good, then replacing x with y in l results in a lottery that is equally as good as l. Tr $\preccurlyeq$ states that for any outcomes or lotteries x, y and z, if y is at least as good as x and z is at least as good as y, then z is at least as good as x.

PSE: For all outcomes or lotteries x, y, z and all probability values p,

$$\text{if } x \approx y, \text{ then } (x, p, z) \approx (y, p, z) \text{ and } (z, p, x) \approx (z, p, y).$$[4]

Tr $\preccurlyeq$: For all outcomes or lotteries x, y, z, if $x \preccurlyeq y$ and $y \preccurlyeq z$, then $x \preccurlyeq z$.

Without saying as much, Temkin leaves the reader with an impression that something must be wrong with the transitivity principle.[5] For apart from transitivity we only need PSE to derive the counter-intuitive continuity for the extreme case from the intuitive assumption that continuity holds for easy cases. But, "PSE is difficult to deny. Indeed,

[3] We owe this reminder to Magnus Jiborn.

[4] PSE is entailed by the so-called Strong Independence Axiom, which states that, for all outcomes or lotteries x, y, z and all probability values p, $0 < p < 1$,

$$x \preccurlyeq y \text{ if and only if } (x, p, z) \preccurlyeq (y, p, z), \text{ if and only if } (z, p, x) \preccurlyeq (z, p, y).$$

[5] This impression is strengthened when one takes into consideration that Temkin has been criticizing transitivity assumptions for many years now, using different arguments. See his (1987), (1996), (1999) and (2000).

perhaps even more than the axiom of transitivity, PSE may appear to be a basic principle of logic and mathematics" (Temkin 2001: 96).

Contrary to what Temkin suggests, PSE is *not* that difficult to deny. We shall say more about this later. For now, however, it is sufficient to point out that Temkin misleadingly makes PSE look like a logical axiom – more precisely, like the axiom of substitutivity of co-extensive expressions, *salva veritate* – when he informally presents it as follows: "Roughly, this principle [= PSE] requires that if $x = y$, then x and y are interchangeable in the formulas in which they occur. Note, outside of modal or intensional contexts, PSE holds for all x and y" (*ibid*.: 96).

Needless to say, this is *not* what PSE requires. PSE is not about substitutions of equivalent expressions in formulas. It is about substitutions of equivalent prizes in lotteries. What it does require is that if a prize x is equally good as a prize y, then interchanging x for y in any lottery results in a lottery that is equally good.

2. TEMKIN'S PROOF

What is more worrying, however, is that Temkin's proof is fatally flawed. As we have seen, he wants to show that PSE and the transitivity principle allow us to derive continuity for the extreme case from continuity for easy cases. Now, in his proof, he assumes (i) the existence of a utility function V that assigns real numbers to outcomes and lotteries and thereby represents their betterness ordering: The higher the number, the better the outcome/lottery. Furthermore, he assumes (ii) that the utility of a lottery equals its expected utility. These two assumptions show up when he takes continuity for easy cases to entail that for each triple o_{i-1}, o_i and o_{i+1} of consecutive outcomes in a gradually descending series there is some probability p such that $V(o_i) = pV(o_{i-1}) + (1-p)V(o_{i+1})$.[6] Thus, if o_2 and o_3 are equally as good as (o_1, p, o_3) and (o_2, q, o_4), respectively, Temkin takes this to mean (cf. *ibid*.: 99) that for some probabilities p and q,

(i) $V(o_2) = pV(o_1) + (1-p)V(o_3)$,

and

(ii) $V(o_3) = qV(o_2) + (1-q)V(o_4)$.

He then uses PSE to replace $V(o_3)$ in (i) by $qV(o_2) + (1-q)V(o_4)$.[7] The proof continues by such substitutions. Now, obviously, making the auxiliary

[6] By continuity for easy cases, there is some probability p such that $o_i \approx (o_{i-1}, p, o_{i+1})$. Consequently, if the existence of a utility function V is assumed, $V(o_i) = V(o_{i-1}, p, o_{i+1})$. Furthermore, if the utility of a lottery is its expected utility, $V(o_{i-1}, p, o_{i+1}) = pV(o_{i-1}) + (1-p)V(o_{i+1})$. Therefore, $V(o_i) = pV(o_{i-1}) + (1-p)V(o_{i+1})$.

[7] Note that there is no need to rely on PSE for this step. Replacing in (i) the left-hand side of equation (ii) by its right-hand side is just a matter of algebra.

https://doi.org/10.1017/S0266267105000556 Published online by Cambridge University Press

assumptions such as (i) and (ii) in the course of the proof of continuity for the extreme case is misconceived. For these assumptions *themselves* entail what is to be proved. More precisely, if (i) the utilities of outcomes and lotteries are representable by real numbers, and if (ii) the utility of a lottery coincides with its expected utility, then continuity trivially holds for all cases, extreme or not. This immediately follows from the mathematical fact that any real number lying between two real numbers is representable as their linear combination. In other words:

For any reals r, r' and r'', if $r' > r > r''$, then for some $p, 0 < p < 1$, $r = pr' + (1 - p)r''$.[8]

If, therefore, as required by (i), there exists a function V from outcomes and lotteries to real numbers that represents the betterness ordering, and if $V(o') > V(o) > V(o'')$, then there *must* exist some $p, 0 < p < 1$, such that

$$V(o) = pV(o') + (1 - p)V(o'').$$[9]

Consequently, if – as required by (ii) – the utility of a lottery (o', p, o''), $V(o', p, o'')$, equals its expected utility, $pV(o') + (1 - p)V(o'')$, it follows that $V(o) = V(o', p, o'')$. Therefore, o and (o', p, o'') are equally good. Thus, continuity trivially follows.

3. A BETTER PROOF

Since continuity is one of the axioms of expected utility theory, assuming that very theory *as an auxiliary assumption* in a proof of continuity is worse than question-begging. Should we then simply consign Temkin's idea to the dustbin? That would be rash. There is an important kernel of truth in his paper that is worth saving. As it turns out, PSE and the transitivity principle do allow us to derive from continuity for easy cases something very close to continuity for the extreme case. More precisely, the following can be proved:

Consider any finite descending outcome sequence $o_1, \ldots, o_n$, that is, any sequence in which every element $o_i (i < n)$ is better than its immediate successor o_{i+1}. Note that, for the proof to follow, the sequence need not be gradually descending: Its elements need not be only slightly better than their immediate successors. The sequence will be said to obey *continuity for adjacent elements* if for every three consecutive elements o_{i-1}, o_i, o_{i+1}, the middle element in that triple, o_i, is equally good as some lottery

[8] To see this, let $p = (r - r'')/(r' - r'')$. For example, for $r' = 10$, $r = 7$ and $r'' = 2$, $p = (7\text{-}2)/(10\text{-}2) = 5/8$.

[9] Rather amusingly, Temkin suggests that "most of us will accept" the existence of a probability p for which this equation holds, in those cases in which both o' and o'' are close to o (ibid: 99). That's certainly true! But the Gricean implication of this statement (that "some of us might not") is more problematic.

(o_{i-1}, p, o_{i+1}) on the first and the third element, with p higher than 0 and lower than 1.

The relation $\approx$ will be said to be *transitive* if and only if for all outcomes or lotteries x, y, z, if $x \approx y$ and $y \approx z$, then $x \approx z$. Note that the transitivity of $\approx$, which we need for the proof, is entailed by Tr$\preccurlyeq$, if we define $x \approx y$ in the standard way, as $x \preccurlyeq y$ & $y \preccurlyeq x$. But the converse entailment does not hold. Thus, the full power of Tr $\preccurlyeq$ is not needed for the result that follows.

Observation 1: Consider any descending outcome sequence $o_1, \ldots, o_n$, with $n \geq 3$. Continuity for the adjacent elements in that sequence, together with PSE and the transitivity of $\approx$, entail that its second outcome is equally good as some compound lottery on the outcomes in the sequence that involves a risk of ending up with o_n.

Since the assumption of continuity for adjacent elements would in the present context be question-begging for some descending sequences, one might find Observation 1 not very interesting. After all, the triple o_1, o_2, o_n in which the second element is only slightly worse than the first one and the third element is radically worse than the other two, is also an example of a descending sequence. The claim that o_2 in such a sequence is equally good as a lottery on the adjacent elements (i.e., on o_1 and o_n) is what is being questioned in the first place. However, since Observation 1 is perfectly general, it also holds for all *gradually* descending sequences that slowly take us all the way to a disaster. For such sequences, continuity for adjacent elements seems to be a more reasonable requirement.

Corollary: Suppose an outcome sequence $o_1, \ldots, o_n$ is gradually descending, that is, every element in the sequence is only *slightly* worse than the immediately preceding one. Suppose also that the initial elements of the sequence are quite satisfactory while the terminal elements are disastrous. Continuity for the adjacent elements in such a gradually descending sequence, that is, Continuity for Easy Cases, together with PSE and the transitivity of $\approx$, entail that the second-best outcome in the sequence, o_2, is equally as good as some compound lottery that involves a risk of the disastrous outcome o_n, and that at best can end up with o_1, which is only slightly better than o_2.[10]

Note that we do not quite prove continuity for this extreme case. Continuity would require that the second-best outcome is equally good as

[10] If the assumption of continuity for adjacent elements would be question-begging for some descending sequences, why do we bother with proving Observation 1 in the first place, instead of directly concentrating on its Corollary? (We are indebted to Peter Vallentyne for pressing this point.) Well, we proceed in this roundabout manner because we want to make clear that the proof we provide below has a very simple structure. In particular, nothing in that proof depends on the somewhat vague notions of a "gradual descent" or "slight worsening" that are central to Corollary.

a *simple* lottery on the best and the worst outcome. We haven't proved that this is the case. But we prove something that is weaker but just as counter-intuitive. For our resistance to continuity for the extreme case rests on the intuition that o_2 is so satisfactory and o_n so disastrous that no chance of a slight improvement from o_2 to o_1 is worth the risk, however small, of ending up with o_n instead.

Proof of Observation 1. For simplicity, let us suppose that a descending outcome sequence has just five elements: $a \succ b \succ c \succ d \succ e$. A generalization of this proof for sequences of arbitrary length is provided in Appendix 1.

We take continuity to hold for the adjacent elements, that is, for any three consecutive outcomes in the sequence. We want to show that b is equally good as some compound lottery that can be constructed out of the elements of this sequence and that involves a risk of ending up with e.

By continuity for adjacent elements, there are some p, p' and p'', $0 < p$, p', $p' < 1$, such that

(i) $b \approx (a, p, c)$,
(ii) $c \approx (b, p', d)$,

and

(iii) $d \approx (c, p'', e)$.

By PSE, (iii) entails that

(iv) $(b, p', d) \approx (b, p', (c, p'', e))$.

From (ii) and (iv), by the transitivity of $\approx$, it follows that

(v) $c \approx (b, p', (c, p'', e))$.

By PSE again, (v) entails that

(vi) $(a, p, c) \approx (a, p, (b, p', (c, p'', e)))$.

From (i) and (vi), by the transitivity of $\approx$, we get

(vii) $b \approx (a, p, (b, p', (c, p'', e)))$.

This is what we have been after. According to (vii), the second-best outcome is equally good as a compound lottery which is constructed out of the outcomes in the sequence, and which involves a non-zero probability for *e*, the worst outcome in the sequence. The probability in question equals $(1 - p)(1 - p')(1 - p'')$, which is greater than zero given that each of p, p' and p'' is lower than one.

4. DISCUSSION

What can we say about these results? If we do not want to be driven to accepting risks of disastrous outcomes for the sake of small improvements, Corollary to Observation 1 implies that we must either (i) deny that Temkin sequences are possible, that is, deny that disastrous outcomes can be reached from satisfactory outcomes by finite gradually descending sequences, or (ii) reject PSE, or (iii) reject Continuity for Easy Cases, or – finally – (iv) reject the transitivity of $\approx$. Since we take this last alternative to be a desperate measure, we won't discuss it in this paper. Instead, we shall consider the first three options.

4.1 Are Temkin's sequences possible?

Denying that a gradually descending sequence can lead from satisfactory outcomes to a disaster is a possibility that Temkin never considers in his paper. Still, it is an option that has been around for many years now. One standard proposal is to postulate *infinite* value distances between satisfactory and disastrous outcomes.[11] That way, we can never reach the latter from the former by a finite sequence of small steps, if a small step *by definition* makes the outcome worse only by a finite value amount. In a descending sequence in which the starting-point and the end-point infinitely differ in value, there must be at least one step, from o_i to o_{i+1}, for some i such that $1 \leq i < n$, in which value radically decreases – by an infinite amount. In that step, then, continuity for easy cases will not be applicable: outcomes o_i and o_{i+1} are too far apart. The expected value of a lottery (o_{i-1}, p, o_{i+1}) is infinitely lower than the value of o_i, for any $p < 1$.

Another standard proposal, which in some ways is closely related to the previous one, avoids postulating infinite values but instead introduces multidimensionality into betterness comparisons and imposes some ordering of lexical priority on the dimensions. In other words, the solution is to forbid any trade-offs between losses on the more important dimensions and improvements on the less important ones.[12] In this way, breaks in continuity can easily be accounted for. Lexical priority works somewhat like infinite differences in value: If one outcome is superior to another on the lexically prior dimension, then the value of the former outcome belongs, so to speak, to a different order of magnitude than the

[11] For the references to relevant literature on 'non-Archimedean' decision theory that allows violations of continuity by introducing infinite utilities, see note 6 in Hájek (2001).

[12] The classical reference is Hausner (1954).

value of the latter outcome. If we now assume that a slight worsening of an outcome by definition must take place on the lexically less important dimension, it follows that really bad outcomes – the ones that we would never want to risk for the sake of small improvements – can *never* be reached by a sequence of slight worsenings from a satisfactory point of departure. Consequently, Temkin's argument does not get off the ground.

This observation should not, however, be assigned too much importance. To begin with, the proposal above has some rather unacceptable consequences. The lexical modeling implies that no trade-offs *at all* are allowed between gains and losses on the different dimensions. If o' has a lower value on the first dimension than an outcome o, by however small amount, then it is worse than o, independently of how much higher value it might have on the second dimension. This seems counter-intuitive.

One could also put this problem in another way: why assume that *slight* worsenings can only take place on the less important dimension? Can't an outcome be slightly worse than another if it is worse on the lexically prior dimension, but only by a very small amount? A remedy might be to suppose that, on the lexically prior dimension, each outcome is assigned only one of, say, two extreme values, 1 or 0. (Say, 1 for life and 0 for death; or – to take another example – 1 for heaven or 0 for hell in the afterlife. Instead of just two values, we could, of course, have several, for example, 3, 2, 1 and 0. What is important is that the value losses with respect to the lexically prior dimension proceed in discrete jumps rather than in a continuous manner.) Under these circumstances, outcome differences with respect to the lexically prior dimension will always be significant. But such a maneuver radically limits the applicability of the modeling. Most relevant dimensions of comparison we can think of are much more continuous than this.

What is more, when we move to *lottery* comparisons, the counter-intuitive implications will arise even for this special 1-or-0 case. Even if the value of an *outcome* on the lexically prior dimension is either 1 or 0, the expected value of a lottery on outcomes can on that dimension be anywhere between these two extremes. If lotteries are compared in terms of their expected values on different dimensions, with the more important dimension being ascribed lexical priority, the following will hold: If a lottery l' has a lower expected value on the first dimension than a lottery l, by however small amount, then it is worse than l, independently of how much higher expected value it might have on the second dimension.

Even if one ignores this difficulty, one might question the basic idea of a multi-dimensional approach to discontinuity: It is by no means obvious that breaks in continuity always are due to multidimensionality in betterness comparisons. Even if outcomes are characterized along just one dimension, as in Temkin's example with a descending sequence of yearly incomes, it may be the case that we are prepared to accept continuity for adjacent outcomes in the sequence but not for the outcomes that are far apart. Thus, breaks in continuity seem to be possible even as one descends along just one dimension.

Or so it might seem. The appearances may be misleading, however. In real life, a factor such as the income size stands for several dimensions that are causally connected to one's earnings. Food, housing, sex, entertainment, travel, etc. are all different dimensions on which one's life may vary depending on income. On the other hand, the relations between such life dimensions do not seem to be lexical, even though some of these dimensions are more important than others. The same considerations apply if we think of infinitistic explanations of continuity breaks instead of the lexical ones. It is not obvious that Temkin's sequence of gradually decreasing yearly incomes must involve some step in which the value of an income *infinitely* decreases, as would be required by the infinitistic account we have sketched earlier. Thus, Temkin may be right in his suggestion that disastrous outcomes, which we don't want to risk for small improvements' sake, sometimes can be reached from satisfactory outcomes by finite gradually descending sequences.

4.2 Is PSE unassailable?

Another option, then, that we should better consider is to question PSE. Contrary to Temkin's suggestion, it is clear that PSE is a highly controversial principle. Here's how it might be criticized: to assume PSE is to accept that we can always replace a part of a whole, such as a prize in a lottery, by another part that is equally good, without changing the overall value of the whole in question. Obviously, this is a claim that post-Moorean moral philosophers should be wary of accepting. If – as G. E. Moore has taught us – the value of a whole need not be the sum of the values of its parts, then replacing parts of that whole with other parts that are equally good may not keep the value of the whole unchanged. In other words, the *contributive* value of a part should not be confused with the value that this part has on its own, independently from the context.

This quick and easy rejection of PSE is, however, much too quick. Treating lotteries as wholes and the possible prizes as their parts is misleading to some extent. That replacements of parts by their equivalents may change the value of the whole package is a simple "complementarity effect," as economists would put it. As far as the instrumental or extrinsic value is concerned, complementarity is a familiar fact of life: in the context of a given package or bundle of goods, the instrumental/extrinsic value of a part may well be different from its value in another package. Therefore, complementarity effects are to be expected: Parts that are equally good in one context can make unequal value contributions in a different context. For *intrinsic* value, the presence of complementarity effects is less obvious,

https://doi.org/10.1017/S0266267105000556 Published online by Cambridge University Press

but if the intrinsic value of a package depends not just on the intrinsic values of its parts but also on their interrelations within the package, then we should expect complementarity effects even in this case. This is what Moore has taught us: Replacing a part with another part with the same intrinsic value may influence the interrelations between the parts and in this way change the intrinsic value of the package as a whole. However, as has often been pointed out, packages or bundles of goods in one important respect differ from lotteries: The former are "conjunctive" wholes, while the latter are essentially "disjunctive" in nature. In a conjunctive whole, its different parts *coexist* and thus can interrelate, while in a lottery, different prizes are *alternatives* to each other. If one of them is realized, the other one is not. Consequently, they cannot interrelate. Therefore, identifying the contributive value of a possible prize with the value that prize has on its own (discounted by probability) is not *obviously* unreasonable.[13] If we still want to reject PSE, we need an independent motivation, which specifically targets disjunctive wholes. That is, we need an argument that specifically questions replacements of prizes in lotteries.

Such a motivation has been provided by Maurice Allais.[14] Following his lead, it is easy to construct examples in which PSE would lead to counter-intuitive results. Thus, consider a compound lottery (($3000, .9, $0), .5, $0). If we in that lottery replace its risky prize, the sub-lottery ($3000, .9, $0), by its monetary equivalent, say, $1000, this replacement may significantly decrease the value of the lottery of the whole. An agent who sets value on money and an extra value on safety could well consider the *safe* $1000 to be an appropriate equivalent of the somewhat risky ($3000, .9, $0). Nevertheless, such an agent may still go for the chance of a larger gain at the expense of a slight increase in risk when both of his choice alternatives are risky, as they are in the choice between (($3000, .9, $0), .5, $0) and ($1000, .5, $0).

How can examples like this be squared with the intuition that Moore's principle of organic wholes is applicable only to wholes with coexistent parts? Here is a plausible explanation: while alternative lottery outcomes do not coexist in a lottery, their *possibilities* do. In one sense, then, a disjunctive whole such as a lottery may be seen as a *conjunctive whole composed of possibilities*. This allows for the interrelations between such possibilities to play a role in the lottery evaluation.

However, we are not yet in the clear. Even though Allais-type examples manage to show that PSE is a deeply problematic principle, they need not be directly relevant to the issues discussed in this paper. It is essential to the cases like the ones Allais has considered that there is a large distance in

[13] For this observation, see v. Neumann and Morgenstern (1953: 18) and Samuelson (1952).

[14] See Allais (1953). For an able defense of Allais' objection, see McClennen (1988). For a thoughtful criticism of Allais-type counter-examples, see Broome (1991: ch. 5).

value between the best and the worst outcome in a sub-lottery that is being replaced by a safe equivalent. (Thus, in our example, there is a large value distance between $3000 and $0.) It is that feature of these cases that accounts for the violation of PSE. Without the large distance in value between lottery outcomes, some of the intuitions behind Allais-type examples would not come into play. On the other hand, in the case of Observation 1, things are different. There is a proof of that observation, slightly more complicated than the one we have sketched above, in which we only substitute elements in a descending outcome sequence with lotteries between *adjacent* elements.[15] If the sequence is gradually descending, the value differences between such lottery outcomes are small, which means that the Allais-type objections are not directly applicable to cases of this kind. For this alternative proof of Observation 1, see Appendix 2. The proof relies on a *generalized* form of PSE that allows substitutions of equivalents in lotteries at an arbitrary position, however deeply imbedded that position may be in a given lottery. Thus, for example, if an outcome o is equally good as a lottery (o', p, o''), then Generalized PSE allows interchanging o with (o', p, o'') in, say, a compound 2-stage lottery in which one of the possible prizes is itself a lottery that has o as one of its prizes. In Appendix 2, we show that Generalized PSE is derivable from the ordinary PSE.[16]

This does not mean, of course, that it is unproblematic to use (generalized) PSE for substitutions of outcomes with lotteries that exhibit small value distances. It might well turn out that such substitutions also are questionable, but that their problematic nature is less easily discernible. Nevertheless, if we do not want to accept the conclusion that a chance of a small improvement is worth taking a risk of a disastrous outcome, *giving up PSE is not enough*, as we now are going to show. For we seem to be driven to that worrisome conclusion even if we don't assume PSE.[17] More precisely, the conclusion in question can be derived using just two

[15] By contrast, in the proof provided above, an outcome that is close to the beginning of the sequence is being replaced by a compound lottery whose final outcomes can greatly vary in value, if the value distance between the outcomes in the beginning and the end of the sequence is large. This makes that proof vulnerable to Allais-type objections.

[16] In the alternative proof of Observation 1, we only make use of a weak version of the generalized PSE, namely, a version which is restricted to substitutions of outcomes with lotteries on adjacent elements. In view of the equivalence between PSE and generalized PSE, it might therefore seem that we thereby manage to derive the same result as in the original proof but now from a weaker premise. However, this isn't exactly correct. While we do prove in appendix 2 that PSE entails the generalized PSE, it is probably not the case that this entailment would still obtain if both principles were restricted to substitutions with lotteries on adjacent elements. At least, the proof we use in the appendix cannot be adjusted to establish the latter claim.

[17] We are indebted to Sven Danielsson for alerting us to this possibility.

principles: a somewhat strengthened version of Continuity for Easy Cases plus a very natural assumption about the concept of "slightly worse."

Extended Continuity for Easy Cases: For all outcomes or lotteries x, y and z, if y is slightly worse than x and z is slightly worse than y, then there is some probability value p, $0 < p < 1$, such that $y \approx (x, p, z)$.

The difference between this version of Continuity for Easy Cases and the original version of that principle is that x, y, and z now are allowed to vary not just over outcomes but also over lotteries: in particular, what we are going to need for our derivation below is the following assumption. If an outcome o is slightly worse than another outcome o' and if a lottery l is slightly worse than o, then there is some $p < 1$, such that $o \approx (o', p, l)$.

The assumption about "slightly worse" that will be used in the derivation says that being slightly worse than a given object is a feature that is inheritable by equivalents.

Slightly Worse Equivalents: For all outcomes or lotteries x, y, z, if $x \approx y$ and x is slightly worse than z, then y is slightly worse than z.

This assumption is partly related to Tr $\preccurlyeq$: One of the implications of the latter principle is that being *worse* than a given outcome or lottery is inheritable by equivalents. Since whatever is slightly worse must *ipso facto* be worse, Tr $\preccurlyeq$ entails that if $x \approx y$ and x is slightly worse than z, then y must be worse than z. But whether it has to be slightly worse is another matter. Thus, the relationship between the two principles is not straightforward: Tr $\preccurlyeq$ neither entails Slightly Worse Equivalents nor is entailed by it.

Observation 2: Consider any outcome sequence $o_1, \ldots, o_n$, with $n \geq 3$, that is gradually descending, that is, a sequence in which every successive element is only slightly worse than its immediate predecessor. Together with Slightly Worse Equivalents, Extended Continuity for Easy Cases entails that the second-best outcome o_2 is equally good as some compound lottery on the outcomes in the sequence that involves a risk of the worst outcome o_n and that at best can end up with o_1, which is only slightly better than o_2.

For the proof of this observation, see Appendix 3.

What lessons should one draw from this result? We might, of course, put into question Slightly Worse Equivalents, but this seems to be a very unattractive option. The first lesson is that PSE is not essential for generating the problematic implications. The second lesson, we shall now argue, is that we should reject Continuity for Easy Cases. We should give up that principle not only in its extended version but also in its original simple form, in which it is meant to apply just to the basic outcomes. Despite the appearances, this premise of Temkin's argument is inherently problematic.

4.3 What about continuity for easy cases?

We want to question, then, whether Continuity for Easy Cases holds. Think of all triples of consecutive outcomes in Temkin's gradually descending sequence – a sequence which starts with a satisfactory outcome and ends in a disaster. Must there exist, for any such a triple o_{i-1}, o_i and o_{i+1}, a lottery on the first and the third outcome that is equally good as the second outcome in the triple?[18]

Since the outcome sequence is gradually descending, o_i is slightly worse than o_{i-1} and o_{i+1} is slightly worse than o_i. Now, suppose the following is the case:

Borderline: While o_{i-1} and o_i are *satisfactory* outcomes, o_{i+1} is not.

In other words, as one moves from o_i to o_{i+1}, the borderline is crossed between what is satisfactory and what is not. Note that this borderline *must* be crossed somewhere in Temkin's sequence. Indeed, as one moves down the sequence, one may well pass several borderlines: one between outcomes that are excellent and those that are less than excellent but still satisfactory, another between satisfactory outcomes and those that are not, and yet another between outcomes that are unsatisfactory but still not disastrous and the truly disastrous outcomes. To simplify the argument, however, let us focus on just one borderline – the one between satisfactory outcomes and those that are unsatisfactory.

There are problems with the potential vagueness of such a borderline, but let us ignore this issue just for a moment. Now, if risks of disasters are not worth chances of small improvements, the same might also be said about risks of unsatisfactory outcomes: Why risk getting an unsatisfactory outcome for a small gain? One might object, however, that there is a disanalogy between borderline cases and risks-of-disaster cases since the latter involve large differences among the outcomes whereas borderline cases only involve small differences. Given that there are only small differences among the outcomes under consideration, one might claim

[18] Would Temkin need to argue that Continuity for Easy Cases holds for *all* Temkin sequences, i.e. for all sequences that gradually descend from highly satisfactory outcomes to disasters? Or is it enough if he can establish that it holds for at least some sequences of this kind? This issue has been raised by an anonymous referee. Well, if Temkin were to argue that Continuity for Easy Cases, together with the other assumptions, suffices to establish unrestricted continuity, then he would need to defend that principle for all sequences of this kind (indeed, for all descending sequences). But if he only wants to show that we *sometimes* in this way are being led to what we consider to be a counter-intuitive conclusion about chances of small improvements being worth taking risks of disasters, then his task is much easier: he only needs to establish that Continuity for Easy Cases holds for at least one sequence of this kind. In what follows, we choose this second interpretation and thus try to provide considerations that put Continuity for Easy Cases in doubt with regard to all sequences like this.

that it is irrational to attach much significance to the fact that one of the outcomes is not satisfactory.[19]

Although we agree that there is in this respect a disanalogy between the two kinds of cases, we think that even in borderline cases there is something intuitive about unwillingness to take certain risks. This intuition can be accounted for in terms of the categorical evaluations of outcomes (such as the outcome being good or bad, satisfactory or unsatisfactory), rather than in terms of the comparative differences between them. If, as things stand, I have a guarantee of a satisfactory outcome, why would it be irrational of me to refuse to risk an unsatisfactory outcome just for the sake of a chance of a small improvement? For example, it doesn't seem irrational to attach importance to leading at least a satisfactory life, so that one would decline any chance of a small improvement that carries with it a risk of leading an unsatisfactory life.

Therefore, if borderline holds, it may well be the case that no chance of a small improvement from o_i to o_{i-1} can compensate the risk, however small, of ending up with an unsatisfactory outcome o_{i+1} instead of the satisfactory o_i. But then, since the difference between o_i and o_{i+1} is small, we have a violation of Continuity for Easy Cases in this particular case.[20] Thus, if the borderline between satisfactory and unsatisfactory outcomes is precise, then Continuity for Easy Cases must be rejected. And without that principle the worrisome conclusion for the extreme case no longer follows.

What happens, though, if it is indeterminate where in the sequence the borderline should be drawn? Well, if the borderline is vague, this vagueness affects some of the applications of Continuity for Easy Cases: for some triples of consecutive outcomes in the sequence, those in the vague area, it will be indeterminate whether continuity holds for these triples or not.[21] More precisely, for every triple of consecutive outcomes for which it is indeterminate whether its first two elements are satisfactory while the third element is not, it will be indeterminate whether Continuity for Easy cases holds for the triple in question.

[19] We owe this point to Peter Vallentyne.

[20] Indeed, if there are several significant borderlines as one moves down the sequence, then whenever any one of them is crossed, Continuity for Easy Cases is put into question.

[21] One might add that, on the so-called supervaluationist approach to vagueness, it *is* determinate that Continuity for Easy Cases is violated at *some* point, even though the identity of that point is indeterminate. On that approach, a statement is vague (= indeterminate as far as its truth value is concerned) if it is true on some precisifications and false on others. Now, for different precisifications of the borderline between satisfactory and unsatisfactory outcomes, the point at which Continuity for Easy Cases fails will be different. But there will always be some such point, on every precisification. For a classical exposition of supervaluationism, see Fine (1975). A detailed recent defense of this account is provided in Keefe (2000). It should, however, be emphasized that the argument in our paper does not presuppose any particular approach to vagueness.

If it is vagueness that accounts for the apparent plausibility of Continuity for Easy Cases, then the derivation of the counter-intuitive conclusion for the extreme case turns out to rely on a principle that is just as problematic as the principle that lies behind the classical *sorites*. When we remove from a heap of sand one grain of sand after another, there will never be a point at which we cross a determinate borderline between a heap and a non-heap. This accounts for the apparent plausibility of the principle: "Removing just one grain cannot make a difference: it can never make a non-heap out of a heap." This principle *seems* plausible because none of its specific applications is determinately false. However, if we start with a heap composed of n grains of sand and then repeatedly apply the principle in question, we are driven to the absurd conclusion that one grain of sand still makes a heap. Likewise, Continuity for Easy Cases *seems* plausible because none of its applications is determinately false. However, repeated applications of that principle drive us to the counter-intuitive conclusion that something like continuity would hold even in the extreme case.

To sum up, since Temkin's result rests on Continuity for Easy Cases, and since that principle should be rejected, there is, *pace* Temkin, no need to reject Transitivity.

5. CONCLUSION

Temkin sequences of outcomes, which gradually lead all the way down to disaster, do seem to be possible and we have also shown that Temkin to some extent was right in his claims about such sequences: in the presence of PSE and given a weak transitivity assumption, Continuity for Easy Cases leads to something like continuity for the extreme case. We seem to be driven to the conclusion that a chance of a small improvement is worth running a risk of disaster. To be sure, PSE, contrary to Temkin's suggestion, is a highly questionable principle, but, as we have shown, it is not really needed for this derivation. We can give up PSE and still derive the counter-intuitive conclusion, if we only extend Continuity of Easy Cases somewhat and let it apply to lotteries as well as to outcomes. On the other hand, Continuity for Easy Cases, even in its original simple form, can be put in doubt. Its apparent plausibility trades on the vagueness of the borderline between satisfactory and unsatisfactory outcomes in a Temkin sequence. The derivation of the counter-intuitive conclusion about the extreme case appears therefore to be just another instance of *sorites*.

APPENDIX 1: PROOF OF OBSERVATION 1

Observation 1: Consider any descending outcome sequence $o_1, \ldots, o_n$, with $n \geq 3$. Continuity for the adjacent elements in that sequence, together with PSE and the transitivity of $\approx$, entail that its second outcome is equally good

https://doi.org/10.1017/S0266267105000556 Published online by Cambridge University Press

as some compound lottery on the outcomes in the sequence that involves a risk of ending up with o_n.

We prove this observation for descending sequences of any length $n \geq 3$, by mathematical induction on the sequence length:

Base step: For $n = 3$, Observation 1 immediately follows from continuity for adjacent elements.

Induction hypothesis: Observation 1 holds for all descending sequences of length n.

We want to establish that it then also holds for all descending sequences of length $n + 1$. Let $o_1, o_2, o_3, \ldots, o_{n+1}$ be such a sequence. We need to show that o_2 is equally good as some compound lottery that involves a risk of ending up with o_{n+1}.

Continuity for adjacent elements implies that

(i) $o_2 \approx (o_1, p, o_3)$, for some p such that $0 < p < 1$.

Since the subsequence $o_2, o_3, \ldots, o_{n+1}$ is of length n, our induction hypothesis implies that

(ii) $o_3 \approx l$, for some compound lottery l on the elements of $o_2, o_3, \ldots, o_{n+1}$ that involves a risk of ending up with o_{n+1}.

Given PSE, (ii) implies that

(iii) $(o_1, p, o_3) \approx (o_1, p, l)$,

and given the transitivity of $\approx$, (i) and (iii) imply that

(iv) $o_2 \approx (o_1, p, l)$.

Since $p < 1$ and l involves some risk of ending up with the worst outcome o_{n+1}, the same applies to (o_1, p, l). Q. E. D.

APPENDIX 2: AN ALTERNATIVE PROOF OF OBSERVATION 1

There is an *alternative proof* of that observation, which, while more complicated, may still be preferable, for reasons suggested in section 4.2 above. It has the following important feature: Whenever some outcome in the sequence in the course of the proof gets replaced by an equally good lottery, the possible outcomes of that lottery are *adjacent* to the outcome that is being replaced.

In this alternative proof, we make use of Generalized PSE which allows substitutions of equivalents at an arbitrary position in a compound lottery, however deeply imbedded that position may be. As is shown below, Generalized PSE is derivable from the ordinary PSE.

Like the proof in Appendix 1, the alternative proof also proceeds by induction on the sequence length. The base step, for $n = 3$, is the same as in the previous proof and the induction hypothesis is the same, namely, that Observation 1 holds for all descending sequences of length n. To prove Observation 1 for descending sequences of length $n + 1$, we consider any such sequence $o_1, o_2, \ldots, o_n, o_{n+1}$. Since its initial subsequence $o_1, o_2, \ldots, o_n$ is of length n, the induction hypothesis implies that o_2 is equally good as some compound lottery l on $o_1, o_2, \ldots, o_n$ that involves some risk of ending up with o_n. (Note that we now apply the induction hypothesis to the initial subsequence rather than to the terminal one, as we did in the previous proof.) Continuity for the adjacent elements implies that $o_n \approx (o_{n-1}, p, o_{n+1})$, for some p such that $0 < p < 1$. We now use Generalized PSE to replace o_n in l by (o_{n-1}, p, o_{n+1}). The resulting lottery l′, which is equally good as l, involves some risk of ending up with the worst outcome o_{n+1}. And by the transitivity of $\approx$, it follows that o_2, which is equally good as l, must be equally good as l'. Q. E. D.

It remains to prove that PSE entails Generalized PSE. To formulate the latter principle in a precise way, we need some definitions.

The *depth* of a lottery is defined recursively:

(i) a lottery on basic outcomes is of depth 1;
(ii) if a lottery l has at least one prize of depth k and if none of its prizes is of a greater depth, then l is of depth $k + 1$.

An outcome or a lottery x will be said to be *contained* in a lottery l iff

(i) l has x as one of its prizes, that is, if for some outcome or lottery y, l is of the form (x, p, y) or (y, p, x),
or
(ii) l has as one of its prizes some l' that in turn has x as one of its prizes,
or
(iii) l has as one of its prizes some l' that has as one of its prizes some l'' that has x as one of its prizes,
or, etc.

In case (i), we say that l contains x at *position* 1 or 2, depending on whether x is the first or the second prize in l. Thus, if $l = (x, p, y)$, then l contains x at position 1 and y at position 2. In case (ii), l contains x at position 11, 12, 21 or 22, depending on whether l' is the first or the second

https://doi.org/10.1017/S0266267105000556 Published online by Cambridge University Press

prize in l and on whether x is the first or the second prize in l'. For example, if $l = ((x, p, y), z)$, then l contains x at position 11 and y at position 12. We define other positions in a lottery in a similar way. If a lottery l contains x at a position α, then $l(y/x)_\alpha$ will be a lottery in which x in l at position α is replaced by y. Note that a lottery may contain x at several different positions. The replacement operation must therefore specify the position at which replacement is to be made.

We now have all we need to formulate the generalized version of PSE:

Generalized PSE: For any outcomes or lotteries x and y and for any lottery l, if l contains x at a position α and $x \approx y$, then $l \approx l(y/x)_\alpha$.

Using this principle we can make interchanges in a lottery at any position, however deeply imbedded that position may be in the lottery in question.

Lemma: PSE entails Generalized PSE.

Proof by strong induction on lottery depth:

Base step: If a lottery l is of depth *1*, the Generalized PSE immediately follows from PSE.
Induction hypothesis: The Generalized PSE holds for the lotteries of all depths up to k.
Now, consider any lottery l of depth $k + 1$. If l contains x at a position α, and $x \approx y$, then x at α either is a prize in l, in which case PSE immediately entails that $l \approx l(y/x)_\alpha$, or l has as its i-th prize ($i = 1$ or 2) some lottery l' that contains x at some position β such that $\alpha = i\beta$. Since l' is at most of depth k, the induction hypothesis implies that $l' \approx l'(y/x)_\beta$. Consequently, by PSE, $l \approx l(l'(y/x)_\beta/l')_{\mathrm{i}}$. Since $l(l'(y/x)_\beta/l')_i = l(y/x)_{i\beta} = l(y/x)_\alpha$, it follows that $l \approx l(y/x)_\alpha$. QED

APPENDIX 3: PROOF OF OBSERVATION 2

Observation 2: Consider any gradually descending outcome sequence $o_1, \ldots, o_n$, with $n \geq 3$. Together with Slightly Worse Equivalents, Extended Continuity for Easy Cases entails that the second-best outcome o_2 is equally good as some compound lottery on the outcomes in the sequence that involves a risk of the worst outcome o_n and that at best can end up with o_1, which is only slightly better than o_2.

Proof by induction on the length of the sequence:
Base step: If $n = 3$, Observation 2 immediately follows by Continuity for Easy Cases.
Induction hypothesis: Observation 2 holds for all sequences of length n.
Now, consider any sequence of length $n + 1$: $o_1, \ldots, o_{n+1}$. By the induction hypothesis, o_3 is equally good as some compound lottery l on the sequence $o_2, \ldots, o_{n+1}$ such that l involves a risk of ending up with o_n and at best can

yield o_2. Since o_3 is slightly worse than o_2, the same must apply to l, by Slightly Worse Equivalents. Consequently, by Extended Continuity for Easy Cases, there is some p, $0 < p < 1$, such that $o_2 \approx (o_1, p, l)$. Since p < 1, and l involves a risk of ending up with o_{n+1}, the same must apply to the lottery (o_1, p, l). Q. E. D.

REFERENCES

Allais, M. 1953. Le comportement de l'homme rationnel devant le risque: critiques de postulates et axiomes de l'école américaine, *Econometrica* 21:503–46

Broome, J. 1991. *Weighing goods*. Cambridge University Press

Fine, K., 1975. Vagueness, truth and logic. *Synthese* 54:235–59. Reprinted in *Vagueness: a reader*, ed. R. Keefe and P. Smith. MIT Press, 1996:119–50

Hájek, A. Pascal's Wager. *The Stanford encyclopedia of philosophy (Spring 2004 edition)*, Edward N. Zalta (ed.), URL + <http://plato.stanford.edu/archives/spr2004/entries/pascal-wager/>

Hausner, M. 1954. Multidimensional utilities. In *Decision processes*, ed. R. M. Thrall, C. H. Coombs, and R. L. Davis, 167–80. John Wiley

Keefe, R. 2000. *Theories of vagueness*. Cambridge University Press

Luce, R. D, and H. Raiffa, 1957. *Games and decisions: introduction and critical survey.* John Wiley

McClennen, E. F. 1988. Sure-thing doubts. In *Decision, probability, and utility*, ed. P. Gärdenfors and N.-E. Sahlin, 166–82. Cambridge University Press

v. Neumann, J. and O. Morgenstern. 1953. *Theory of games and economic behavior*, 2nd edn. John Wiley

Resnik, M. D. 1987. *Choices: an introduction to decision theory*, University of Minnesota Press

Samuelson, P. A. 1952. Probability, utility and the independence axiom, *Econometrica* 20:670–8

Temkin, L. 1987. Intransitivity and the mere addition paradox. *Philosophy and Public Affairs* 16:138–87

Temkin, L. 1996. A continuum argument for intransitivity. *Philosophy and Public Affairs* 25:175–210

Temkin, L. 1999. Intransitivity and the person-affecting principle: a response. *Philosophy and Phenomenological Research* 59:777–84

Temkin, L. 2000. An abortion argument and the threat of intransitivity. In *Well-being and morality: essays in honour of James Griffin*, ed. R. Crisp and B. Hooker, 336–56. Oxford University Press

Temkin, L. 2001.Worries about continuity, expected utility theory, and practical reasoning. In *exploring practical philosophy: from action to values*, ed. D. Egonsson, J. Josefsson, B. Petersson, and T. Rønnow-Rasmusen, 95–108. Ashgate Publishers

Vallentyne, P. 1993. The connection between prudential and moral goodness. *Journal of Social Philosophy* 24:105–28

https://doi.org/10.1017/S0266267105000556 Published online by Cambridge University Press