The Epistemic Problem Does Not Refute Consequentialism

TYLER COWEN

George Mason University

Should radical uncertainty about the distant future dissuade us from judging options by referring to their consequences? I argue no. Some short-run benefits are sufficiently high that we should pursue them, even if our long-run estimates possess a very high variance. I discuss the relationship between the epistemic argument and 'fuzzy' rankings and also 'arguments from infinity'. Furthermore, extant versions of the epistemic argument require the assumption that we have no idea about the major consequences of our acts. Even a slight idea about some major consequences will render the epistemic argument less plausible. In most applications of the epistemic argument, long-run uncertainty is not the relevant confounding variable; on close examination the epistemic argument tends to trade on other principles altogether.

INTRODUCTION

Some commentators have suggested that epistemic problems force us to abandon moral consequentialism. By moral consequentialism, I mean the view that we should judge actions primarily or solely in terms of their consequences.[1] But how can we pursue good consequences if we have no serious idea what we are doing?

Shelley Kagan summarizes the critical argument as follows:

> Perhaps the most common objection to consequentialism is this: it is impossible to know the future ... This means that you will never be absolutely certain as to what all the consequences of your act will be ... there may be long term bad effects from your act, side effects that were unforeseen and indeed unforeseeable ... So how can we tell which act will lead to the best results overall – counting *all* the results? This seems to mean that consequentialism will be unusable as a moral guide to action. All the evidence available at the time of acting may have pointed to the conclusion that a given act was the right act to perform – and yet it may still turn out that what you did had horrible results, and so in fact was morally wrong. Indeed, it will never be possible to say for sure that any given act was right or wrong, since any event can continue to have further unseen effects down through history. Yet if it is impossible to tell whether any act is morally right or wrong, how can consequentialism possibly be a correct moral theory?[2]

[1] See Smart and Bernard Williams, *Utilitarianism: For and Against* (Cambridge, 1973); Samuel Scheffler, *The Rejection of Consequentialism* (Oxford, 1982); and Peter Singer, *Practical Ethics* (Cambridge, 1993) for some standard treatments of consequentialism, which of course includes utilitarianism as one variant.

[2] Shelley Kagan, *Normative Ethics* (Boulder, 1998), p. 64. Note that Kagan is summarizing this view, not endorsing it.

© 2006 Cambridge University Press
Utilitas Vol. 18, No. 4, December 2006
doi:10.1017/S0953820806002172 Printed in the United Kingdom

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

We can express the relevant critique of consequentialism in different language. We wonder if we have the correct moral theory in the first place, if we cannot know 90 percent, or perhaps 99.9 percent, of what is to count toward a good outcome.[3]

These epistemic problems run deeper than the usual political disagreements about what is the right thing to do. As a practical matter, we know very little about how our current choices will alter the world to come. Many of our decisions set off extensive long-run chain reactions, which make the distant future difficult to predict.

To view the point in its most extreme form, what if John bends down to pick up a banana peel? If nothing else, this action will likely affect the identities of all his future children, if only by changing the timing of future sex acts by a slight amount, or by reconfiguring the position of John's sperm within his testicles.[4] And a different set of people born in the future will, in many cases, cause the world to take a very different path.

Ray Bradbury raised such an issue in his short story 'A Sound of Thunder'; the general mechanism has since become a common feature of popular culture.[5] If you go back in time, and alter one small event, the entire history of the world can change. One extra sneeze from one caveman, millennia ago, probably would suffice to change the entire course of world history. But time travel is not needed for the basic mechanism to apply to current choices. Today, when you stop at a traffic light, rather than plowing through the yellow, you likely affect the length of other commutes and thus change the timing of millions of future conceptions. Subsequent genetic identities will change as well. Come the next generation, these different identities lead to different marriage patterns and thus an entirely new set of individuals in the future.

In many cases we might think that aggregate outcomes are stable to small perturbations. The long-run random effects will sometimes cancel out or offset each other; Tolstoy went so far as to suggest that the 'great men' of history had no impact at all, citing Napoleon in this connection. Or some physical systems may be stable with regard to small perturbations; for instance your conduct in traffic may be offset by subsequent traffic lights and obstacles.

That being said, we cannot count on all or even most perturbations to cancel out. Note that our traffic behavior need only lead other male drivers to jostle enough to redistribute their sperm, thus affecting the identities of future conceptions. Given that a male body holds millions

[3] I am indebted to David Schmidtz for this formulation of the point.
[4] Derek Parfit, *Reasons and Persons* (Oxford, 1987).
[5] Ray Bradbury, 'A Sound of Thunder', *The Stories of Ray Bradbury* (New York, 1980).

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

of sperm, it is unlikely that various jostlings will have no affect on the genetic identities of future children. Again, only the smallest of changes are needed to have long-term repercussions.

Tolstoy's 'impersonal forces of history' argument suggests that it does not matter much which individuals are born, but this is implausible. Had Hitler, Lenin, and Mao never existed, the world probably would have taken very different paths. For instance, it is unlikely that Nazism or the Bolshevik Revolution would have succeeded or had the same impact on the world stage. The Second World War, as we know it, would not have occurred, nor would the Holocaust. It is an open question whether the state of Israel would have been created. Of course each of these subsequent changes will in turn alter other historical events and thus reshuffle other gene combinations as well.

For small changes to translate into large final effects, we need only postulate that some individuals, or some leaders, play a significant role on the global stage. Even if *most* individuals do not matter, or *most* small changes wash out, *some* of the small changes today will alter future identities, once we look a generation or two into the future. So the argument requires only that a very small number of personal identities matter for the course of history. If Hitler's great-great-grandfather had bent down to pick one more daisy, many of the effects might have washed out; nonetheless Europe today would be a very different place.

Not only can small changes in initial conditions have very large effects in the long run, but we do not have much idea how those effects will play out. If we try to imagine European history without Hitler, we are playing a very poorly informed guessing game. We have only a general sense that many things would be very different. Following the literature, I refer to this as the epistemic problem.[6]

The difficulty of calculating consequences becomes a *much* more important issue when we maintain a deep concern for the distant

[6] In 'Consequentialism and Cluelessness', *Philosophy and Public Affairs* 29 (2000), James Lenman states the epistemic critique of utilitarianism and cites some precursors. For additional perspectives on the epistemic critique, see Alastair Norcross, 'Consequentialism and the Unforeseeable Future', *Analysis* 50 (1990); Robert L. Frazier, 'Act-Utilitarianism and Decision Procedures', *Utilitas* (1994); Frances Howard-Snyder, 'The Rejection of Objective Consequentialism', *Utilitas* (1997), and Fred Feldman, 'Actual Utility: The Objection from Impracticality, and the Move to Expected Utility', unpublished manuscript, University of Massachusetts. In *The Fatal Conceit: The Errors of Socialism* (Chicago, 1991), Friedrich A. Hayek can be thought of as offering a version of the epistemic critique as well. Shelley Kagan, *Normative Ethics*, p. 64, calls the epistemic argument 'the most common objection to consequentialism'. William Whewell provided one early statement of the epistemic critique: see the discussion in John Stuart Mill, 'Whewell on Moral Philosophy', *Essays on Ethics, Religion and Society*, ed. J. M. Robson (Toronto, 1969 [1852]). In *What If? The World's Foremost Military Historians Imagine What Might Have Been* (Berkeley, 2000), Robert Cowley, Stephen E. Ambrose, and David McCullough consider a series of counterfactuals in the context of military history.

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

future. At sufficiently high rates of discount, most future consequences cease to matter much within a thirty-year horizon. But with sufficiently low rates of discount, welfare-relevant consequences continue for a very long time. I have defended low discounting in other writings, but for our purposes here I will take it as given that the future matters enough to worry about.[7]

ONE CASE WHERE CONSEQUENCES CLEARLY SHOULD MATTER

Let us start with a simple example, namely a suicide bomber who seeks to detonate a nuclear device in midtown Manhattan. Obviously we would seek to stop the bomber, or at least try to reduce the probability of a detonation. We can think of this example as standing in more generally for choices, decisions, and policies that affect the long-term prospects of our civilization.

If we stop the bomber, we know that in the short run we will save millions of lives, avoid a massive tragedy, and protect the long-term strength, prosperity, and freedom of the United States. Reasonable moral people, regardless of the details of their metaethical stances, should not argue against stopping the bomber.

No matter how hard we try to stop the bomber, we are not, a priori, committed to a very definite view of how the long run will play out. After all, stopping the bomber will reshuffle future genetic identities, and may bring about the birth of a future Hitler. We can of course imagine possible scenarios where such destruction works out for the better *ex post*. Perhaps, for instance, the explosion leads to a subsequent disarmament or anti-proliferation advances. But we would not breathe a sigh of relief on hearing the news of the destruction for the first time. Stopping the bomber brings a significant net welfare improvement in the short run, while we face radical generic uncertainty about the future in any case.

Furthermore, if we can stop the bomber, our long-run welfare estimates will likely show some improvement as well. The bomb going off could lead to subsequent attacks on other major cities, the emboldening

[7] See, for instance, Tyler Cowen and Derek Parfit, 'Against the Social Discount Rate', *Philosophy, Politics, and Society*, 6th series, ed. P. Laslett and J. Fishkin (New Haven, 1992), and Tyler Cowen, 'Policy Implications of Zero Discounting', *Social Philosophy and Policy* (2003). Note that this strong concern for the future refers to human well-being or utility and to the idea of good consequences more generally. Economists (e.g. John Broome, 'Discounting and Welfare', *Philosophy and Public Affairs* 23 (1994)) typically discount future *dollar* benefits positively, because those benefits can be reinvested at a positive rate of return. But the economic approach provides no particular reason to discount future utility in positive terms. A deep concern for the distant future, as outlined here, is consistent with standard practices of applied economics.

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

of terrorists, or perhaps broader panics. There would be a new and very real doorway toward the general collapse of the world. While the more distant future is remixed radically, we should not rationally believe that some new positive option has been created to counterbalance the current destruction and its radically negative potential implications. To put it simply, it is difficult to see the violent destruction of Manhattan as *on net* – in *ex ante* terms – favoring either the short-term or long-term prospects of the world.

Even if the long-run expected value is impossible to estimate, we need only some probability that the relevant time horizon is indeed short (perhaps a destructive asteroid will strike the earth). This will tip the consequentialist balance against a nuclear attack on Manhattan. Now it is not a legitimate response simply to assume away the epistemic problem by considering only the short time horizon. But if the future is truly radically uncertain, as the epistemic argument suggests, we cannot rule out *some chance* of a short time horizon. And if everything else were truly incalculable and impossible to estimate, we should be led to assign decisive weight to this short time horizon scenario. We again should stop the bomber.[8]

If the Manhattan example does not convince you, consider the value of stopping a terrorist attack that would decimate the entire United States. Or consider an attack that would devastate all of Western civilization, or the entire world. At some point we can find a set of consequences so significant that we would be spurred to action, again in open recognition of broader long-run uncertainties.

Surely at *some* point the upfront change must be large enough to provide a persuasive reason for or against it. What if a cosmological disaster destroyed 99.9999 percent of all intelligent life across the universe? Yes, it is possible that subsequent cosmological events could lead to an even greater blossoming of wonders, but at some point of comparison this point is simply fatuous. Most of the life in the universe is being destroyed and more likely than not this is a horrible catastrophe even in the much longer run. So we can argue 'how large' an upfront event is needed to sway us toward an evaluative judgment, but a sufficiently large upfront event should do the trick.

We should not confuse the epistemic argument with the very different argument from infinity. It might be argued, for instance, that the universe contains an infinite amount of expected value or welfare.[9]

[8] I am not suggesting that a concern with consequences is the only reason for stopping the bomber, only that it is sufficient reason. Obviously our duties will push in the same direction; the point is whether consequences provide a good enough reason on their own.

[9] Nick Bostrom, 'Infinite Ethics', unpublished manuscript (2004), available at http://www.nickbostrom.com/.

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

Any apparently positive or negative change therefore has zero import relative to the larger infinity. Whether or not this argument is powerful, it is not the problem at hand. The infinity argument does not require uncertainty at all, as it would hold even if we had perfect information about all future consequences. In contrast, the epistemic argument cites a high variance for potential future outcomes, or an inability to forecast the future, as reasons not to assign moral importance to current consequences. We are not comparing an upfront benefit to infinity, we are comparing it to some hard-to-estimate magnitude in the future.

LENMAN'S COUNTEREXAMPLES AND THEIR IMPLICATIONS

James Lenman, a leading purveyor of the epistemic critique, presents several cases to make us doubt the importance of consequences as an account of moral rightness.[10] But at best Lenman's arguments might downgrade the importance of some kinds of consequences; they do not discriminate against consequentialism more generally.

One of Lenman's central intuitions involves a D-Day example.[11] Imagine that we must decide which French beach to invade to fight Hitler. Our decision might have monumental consequences for the future of civilization. (Imagine that the war is sufficiently close that it would matter. I draw the example from Lenman, noting that we could improve the details if necessary to sidestep fact-based objections.) Imagine further that we have two potential candidate beaches for the invasion, but no particular military reason to favor one beach or the other. One beach is probably much better than the other, but when making the choice we do not know which one will be better. We do know, however, that if we land at beach A, we will cause inconvenience to a dog. As Lenman constructs the example, the dog will experience a broken leg, but the inconvenience could be less painful to make the contrast with the human consequences all the more glaring. If we land at beach B, no dog is injured.

In most plausible moral theories we attach some weight to the suffering of animals. It is bad if the dog on beach A must suffer. Yet Lenman suggests that the fate of the dog provides at best an extremely weak reason to favor landing at beach B rather than beach A. The fate of the dog is tiny relative to what stands at stake in the comparison. According to Lenman, the real comparison, if we knew it, would swamp the factor of the dog. Our mere ignorance of the true relevant factors

[10] James Lenman, 'Consequentialism and Cluelessness', *Philosophy and Public Affairs* 29 (2000).
[11] Ibid.

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

should not allow us to elevate one small factor in importance, when that factor would otherwise become negligible in the final calculations. In Lenman's view, our uncertainty about beach B vs. beach A does not cause those two options to cancel, leaving the fate of the dog to tip the scales one way or the other. Instead, he wonders whether we can evaluate actions in terms of their consequences altogether.[12]

The hard-line response, of course, may dismiss Lenman's intuition rather than responding to it. We can imagine the consequentialist yelling out: 'Save the dog from a broken leg, damn the uncertainty, the variance of the invasion decision is high in any case.' The epistemic critique of course rejects this response; I believe it is drawing on some version of what I call the Principle of Roughness:

The Principle of Roughness

> Some of our choice options will differ in complex ways. We might nonetheless, *ex ante*, make a reasoned judgment that they are roughly equal in value, and that we should be roughly indifferent across the two options. After making a small improvement to one of these choices, we still might be roughly indifferent to which option is better.

Versions of this principle are found in the literature on incommensurability, incomparability, and related concepts.[13] For instance, we often rely on The Principle of Roughness in aesthetics. Assume that we are trying to judge whether Beethoven or Mozart is the better composer. We might judge the two composers as being roughly equal, or judge that neither composer can be elevated over the other. Assume then that we discover one new work by Beethoven, a lovely two-minute bagatelle for piano. We are not obliged to assert that Beethoven is now the better composer. Our original judgment was sufficiently 'rough' that it can survive this new discovery. In contrast, a very exact comparison, such as that of weight or measure, could be upset by a small change at the appropriate margin.

We also find the Principle of Roughness in some of our judgments of goodness. We might, for instance, be choosing between a new health care program, and a new poverty reduction program, each with significant and complex benefits for different groups of people. We might judge that the two policies are roughly equal in value. We might then

[12] In 'Consequentialism and Cluelessness', Lenman appears to favor 'ethical theories for which the focus is on the character of agents and the qualities of their wills, for theories that are broadly Kantian or Aristotelian in spirit'. Note that Lenman also offers a 'personal integrity' objection to the consequentialist pursuit of uncertain benefits, which I do not consider here.

[13] Ruth Chang (ed.), *Incommensurability, Incomparability, and Practical Reason* (Cambridge, 1997); Ruth Chang, 'The Possibility of Parity', *Ethics* 112 (2002).

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

discover that one of the policies was slightly better than we previously had thought. Yet we would not then have to declare that the revised policy suddenly is clearly morally better than the other policy.

I am not seeking to argue for the Principle of Roughness; rather I am content to note that many people hold this intuition. And this Principle will give some weight to the epistemic critique. The Principle of Roughness, to the extent it applies, leads us to discriminate against relatively small benefits and losses. For this reason the Principle of Roughness may be operating in the epistemic critique. The future amounts at stake seem so large that (relatively) small changes in upfront benefits and costs do not matter much and do not change the fundamental nature of the comparison.

In most applications of the Principle of Roughness (e.g. Mozart vs. Beethoven), small changes (i.e. discovery of an extra sonata) are swamped by high *absolute totals* (of achievement) in the first place. In the D-Day example, the small change is swamped by the high *variance* in our estimates of consequences. We can think of the epistemic critique as extending the Roughness Principle to uncertainty and perhaps to higher moments of the probability distribution.

The Principle of Roughness, however, does not refute consequentialism. Instead the Principle reflects the importance of pursuing large benefits. The rougher the initial comparison, the larger the upfront benefits required to make the new comparison a clear one. To refer back to the music example, we probably would suddenly judge Beethoven as superior to Mozart if we discovered fifty new Beethoven symphonies, each as wonderful as his Fifth or Ninth.

Large upfront benefits are less likely to be overwhelmed by the roughness of our comparisons, and thus it is compelling to pursue large benefits. To borrow a metaphor, anything we try to do today is 'floating in a sea of long-run outcome variance'. Only big, important goals will, in reflective equilibrium, stand above the ever-present froth and allow the comparison to be anything more than a rough one. When small goals are at stake, our moral intuitions become confused, and as a result perhaps we should downgrade the importance of those goals, at least when the level of background uncertainty is high.

To refer back to Denman's D-Day example, consider the magnitude of the upfront benefit in Lenman's example, namely preventing a dog's broken leg. To be sure we may find real merit in animal welfare arguments. But do we really know how to weigh the welfare of dogs against the welfare of humans, or against other species? We start off being uncertain about the value of the dog, and then we find that the epistemic critique intensifies this uncertainty. We wonder something like: 'given the possible variance of outcomes, could the dog's leg be so important as to sway the issue?'

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

To put those questions into stark relief, let us consider another moral dilemma, again involving an invasion decision. Instead of preventing the broken leg of a dog, assume that one beach invasion will kill one hundred innocent civilians (if need be, make it one thousand lives). The other beach will involve no comparable danger. By assumption, these civilian lives have nothing to do with the final outcome of the war, but obviously one hundred innocent lives carry considerable moral value.

Obviously we should choose the landing that has no chance of killing the civilians, given that we have no other reason for favoring one beach over the other. Consequentialism gives us a straightforward answer, even though our actions will set off unforeseeable long-run effects. Once the upfront benefit is sufficiently high, the epistemic critique has less force, though of course we remain uncertain as to whether we will choose the correct beach for defeating Hitler.

VARYING OUR DEGREE OF KNOWLEDGE

So far we have discussed whether the epistemic critique is robust to large rather than small upfront benefits. A second question is whether the epistemic critique is robust to having a very small sense that one future course is better than another. In the D-Day example, we have absolutely no information about how to attain the primary goal of importance, namely defeating the Nazis. I will vary this assumption to see which aspects of the problem are driving our intuitions.

We can characterize our knowledge of the future in at least two ways:

(1) We can always attach a Bayesian probability, whether explicit or implicit, to various outcomes.[14] Even if we are very uncertain, in principle there exist betting odds that we would or would not be willing to take on a given choice. So we might say that we have 'no idea' as to whether our Middle Eastern policy will work, but in reality we would take some bets at some odds on the success of the policy and refuse other odds. These counterfactual bets help us pin down implicit probability estimates. That being said, the variance of the probability distribution for future events still might be extremely high.

(2) Under 'Knightian uncertainty' we simply cannot assign probabilities to some events.[15] Those events are unique; some of them might be 'unimaginable', and in any case not subject to a probability calculus.

[14] Bryan Caplan, 'The Austrian Search for Realistic Foundations', *Southern Economic Journal* 65 (1990).

[15] Gerald P. O'Driscoll, Jr. and Mario Rizzo, *The Economics of Time and Ignorance* (Routledge, 1996).

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

Under the Bayesian view, the D-Day example postulates that each D-Day invasion beach has exactly the same odds of succeeding. But of course this equality is one very special relationship, and it will not generally hold. In most cases, no matter how great our long-run uncertainty, we will have some reason for assigning *differing* probabilities to the value of each beach for a successful invasion. And in those cases we should choose the more appropriate beach.

Consider this scenario where we have a slight (rational) sense of which beach is better for the invasion. Assume we know that if a windstorm comes that day, beach A is better for the military campaign against Hitler. It so happens that the chance of a windstorm is very small in France at that time of the year, but still the chance of the windstorm is not zero. Otherwise, if no windstorm comes, we have no idea which beach is better for the invasion, although one beach will turn out to be much better than the other beach, *ex post*. (If we wish, we could stipulate also that beach A also avoids the dog's broken leg, although we no longer need this benefit to reach a conclusion.) Given that all other matters are held equal, we should invade the beach that will turn out to be better in the windstorm.

Lenman's example assumes that we know literally nothing about the major consequences of our acts; we know only the minor consequence concerning the dog. In contrast, the windstorm example assumes that we know a small amount about the major consequences of our acts, albeit not very much. Once we know a small amount about major consequences, however, the case for counting consequences appears more robust. And in most real world cases, no matter how great our uncertainty, we do know at least a small amount about major consequences, if only in stochastic terms. So the epistemic critique does not much weaken consequentialism when we have *some* information about *some* consequences of major importance.[16]

Now the epistemic critique may be relying on 'Knightian uncertainty' rather than Bayesian estimates. But even in these cases we still have degrees of uncertainty. I may have 'no idea' about my forthcoming birthday surprise, but this uncertainty is not comparable to my 'no idea' about intelligent life on other planets. Background social context will give us some expectations, even if we cannot assign definite numbers to probability forecasts. Knightian uncertainty still appears to admit of differences of degree. We are back to the likelihood that we often will have some idea, however slight, or however non-quantifiable, as to which beach is better for the invasion.

[16] On the argument that we often have some idea about consequences, see Mill, 'Whewell on Moral Philosophy'.

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

The epistemic critique relies heavily on a *complete* lack of information about initial circumstances. This is not a plausible general assumption, although it may sometimes be true. At the same time, the epistemic critique appears to be using a *more plausible assumption*, namely that of a high variance for the probability distribution of our estimates concerning the future. But simply increasing the level of variance or uncertainty does not add much force to the epistemic argument. The original force came from the assumption of no information about major events of consequence.

To see this more clearly, consider another case of a high upfront benefit. Assume that the United States has been hit with a bioterror attack and one million children have contracted smallpox. We also have two new experimental remedies, both of which offer some chance of curing smallpox and restoring the children to perfect health.

If we know for sure which remedy works, obviously we should apply that remedy. But imagine now that we are uncertain as to which remedy works. The uncertainty is so extreme that each remedy may cure somewhere between 300,000 and 600,000 children. Nonetheless we have a slight idea that one remedy is better than the other. That is, one remedy is slightly more likely to cure more children, with no other apparent offsetting negative effects or considerations. Despite the greater uncertainty, we still have the intuition that we should try to save as many children as possible. We should apply the remedy that is more likely to cure more children. We do not say: 'We are now so uncertain about what will happen. We should pursue some goal other than trying to cure as many children as possible.'

Nor would we cite greater uncertainty about longer-run events as an argument against curing the children. We have a definite good in the present (more cured children), balanced against a radical remixing of the future on both sides of the equation. The definite upfront good still stands firm. Alternatively, let us assume that our broader future suddenly became less predictable (perhaps genetic engineering is invented, which creates new and difficult-to-forecast possibilities). That still would not diminish the force of our reason for saving more children. The variance of forecast becomes larger on both sides of the equation – whether we save the children or not – and the value of the upfront lives remains. A higher variance of forecast might increase the required size of the upfront benefit (to overcome the Principle of Roughness), but it would not refute the relevance of consequences more generally.

We could increase the uncertainty more, but consequentialism still will not appear counterintuitive. The remedies, rather than curing somewhere in the range of 300,000–600,000 children, might cure in the broader range of zero to all one million of the children. By all

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

classical statistical standards, this new cure scenario involves more uncertainty than the previous case, such as by having a higher variance of possible outcomes. Yet this higher uncertainty lends little support for the view that curing the children becomes less important. We still have an imperative to apply the remedy that appears best, and is expected to cure the greater number of children.

This example may appear excessively simple, but it points our attention to the non-generality of the epistemic critique. The critique appears strongest only when we have absolutely no idea about the future; this is a special rather than a general case. Simply boosting the degree of background generic uncertainty should not stop us from pursuing large upfront benefits of obvious importance.

SOME PHILOSOPHIC CONCLUSIONS

The epistemic critique increases the plausibility of what I call 'big event consequentialism'. In this view, we should pursue good consequences, but with special attention to consequences that are very important and very good, or correspondingly, very bad. This includes stopping the use of nuclear weapons, saving children from smallpox, making progress against global poverty, and maintaining or spreading liberal democracy. Big events, as I define them, typically are of significant practical importance, involve obvious moral issues, and their value is not controversial to benevolent onlookers.

In contrast, consider 'small events'. Preventing a broken leg for a single dog, however meritorious an act, is a small event as I define the concept. Making American families wealthier by another $20 also would count as a small event. We should not count small events for nothing, but epistemic issues may well lower their importance in reflective equilibrium. Of course we do not need a strict dividing line between big and small events, but rather we can think in terms of a continuum.

In some cases a large number of small benefits will sum up to a big benefit, or equal the big benefit in importance. It then can be argued that we should treat the large benefits and the small benefits on a par. If we lift a different person out of poverty one billion times, this is no less valuable than lifting one billion people out of poverty all at once.

Here two points are relevant. First, sometimes we are facing a single choice in isolation from other choices, rather than examining a rule or general principle of behavior. In this case it does not matter whether or not the small benefits would, if combined in larger numbers, sum up to a greater benefit. The small benefits will not be combined in greater numbers, and we should still upgrade the relative importance of larger benefits in our decision calculus.

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

Second, not all small benefits sum into equivalence with larger benefits. Sometimes one value has a lexical relationship to (all or some) other values. For instance arguably a large number of canine broken legs, even a very large number, do not sum in value to make a civilization. It does not matter how many dogs and how many broken legs enter the comparison. In other words, civilization may be a lexical value with respect to canine broken legs. And when lexical elements are present, the mere cumulation of numbers of broken legs does not trump the more significant value.

Numerous value relationships have been cited as lexical. A large number of slight headaches, no matter how numerous, may not sum up in value to equal a smaller number of intensely painful deaths or personal tortures.[17] A very large number of 'muzak and potato' lives do not sum to overtake the value of a sophisticated civilization.[18] Rawls put forward liberty and the difference principle as his lexical values for all political comparisons.[19] For our purposes, we do not require a very strict notion of lexicality for these designations to matter. A big value need not be lexical against a (multiplied) smaller value at all possible margins. Instead the big value need only be lexical across the comparisons that arise under relevant policy comparisons. Furthermore a big value need not be lexical in absolute terms against all other smaller values.

We therefore receive further guidance as to which big events are upgraded in the most robust fashion. The big values that receive the most robust upgrading would be those values with some lexical importance, relative to possible comparisons against other smaller values.[20]

To sum up these points, critics of consequentialism would like to establish something like the following: 'We find it hard to predict consequences. Therefore consequences do not matter very much, relative to other factors, such as deontology or virtue ethics. We should abandon consequentialist morality.' But so far epistemic considerations have yet

[17] Alastair Norcross, 'Comparing Harms: Headaches and Human Lives', *Philosophy and Public Affairs* 26 (1997).

[18] Derek Parfit, 'Overpopulation and the Quality of Life', *Applied Ethics*, ed. Peter Singer (Oxford: Oxford University Press, 1986), pp. 145–64; Parfit, *Reasons and Persons*.

[19] John Rawls, *A Theory of Justice* (Cambridge, Mass.: Belknap Press, 1987).

[20] It might be argued that a lexical or partially lexical value does not need further upgrading. Such a conclusion is not justified. Value A may be lexical over value B in a direct comparison. That is, no amount of B can add up in such a way as to surpass a given quantity of A. This notion of lexicality fits the examples below. But at least two questions remain. First, would we give up *any* amount of value B to have any *additional* amount of value A? Second, would we give up any amount of B to raise the *probability* of keeping more of A? Lexicality, like many other ethical concepts, can have more than one dimension. Our intuitions across one of those dimensions do not automatically translate into the other dimensions.

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

to produce a strong argument for this view. The arguments support a different conclusion, namely downgrading the importance of minor consequences, and upgrading the importance of major consequences. The most robust major consequences are those which carry values with some lexical properties, and cannot be replicated by a mere accumulation of many small benefits.

SOME POLICY IMPLICATIONS

Given big event consequentialism, our policy choices should be directed toward achieving big events, especially if they have partially lexical properties. Many policies do pursue goals of this kind, such as poverty reduction, world peace, and preventing environmental collapse.

Furthermore, we should persist in striving for such goals. As we struggle for great achievements, along the way we will encounter many incidental costs and obstacles. We are not certain how much weight we should attach to these costs, just as we were not sure how to value the dog's broken leg in the D-Day example, given the high generic uncertainty attached to our choices. But we now learn that these incidental costs may be less important than we had thought.

A 'Think Big' politics, of course, ought to be properly sophisticated and take proper account of human rights. It does not imply, for instance, that Mao's Great Leap Forward, or the Cultural Revolution, was a good idea. The goal must be desirable, and the means must be well-suited toward achieving the relevant goal. Not all big ideas are good ideas. Instead, we should upgrade the relative importance of pursuits that are both big and good.

In pursuing big goals, it is appropriate to adjust for the likelihood of making mistakes. It might be argued, for instance, that we easily become captive to false ideologies that promise amazing benefits, such as millenarianism. Perhaps we become too easily swept up in the glories of revolution or utopias, just as many Western intellectuals mistakenly pledged loyalty to Mao. So we might, for indirect reasons, find the pursuit of smaller goals the best way to achieve the ultimate large goals of importance. Such a possibility will rest on empirical considerations which are beyond the scope of this article. It remains the case, however, that the epistemic critique has not defeated the importance of good consequences. Furthermore we would have a new recipe for identifying which were the small benefits of importance, namely those most likely to lead to subsequent large benefits, given the relevant human imperfections in making decisions along the way.

The epistemic critique still implies some very strong cautions. Although our political views should involve a commitment to achieving very important ends, our attachment to particular means should be

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

tentative, highly uncertain, and radically contingent. In other words, we should be skeptical of political ideologies that involve very definite views on means–ends relationships. After all, why should we be so sure that our favored ideology will in fact bring good consequences? Our knowledge that big goals are especially important does not imply that we have a very good idea how to achieve such goals.

To be sure, we should opt for the policies that are most likely to bring about big and important ends. But our particular policy recommendations, though we believe them to be the best available, will stand only a very slight chance of being correct. They ought to stand the *highest* chance of being correct, of all available views, but this chance will not be very high in absolute terms. We should think of the details of our political views as analogous to betting on a slightly crooked roulette wheel, designed to produce the number two more than a proportionate amount of the time. We should bet on the slightly favored outcome, namely the number two, and by doing so we can improve our prospects. But most of the time we are likely to predict the wrong number; we will be betting on two and other numbers will come up.

Our political stances and policy recommendations should therefore be accordingly tolerant. Imagine a world where your chance of being right is 2 percent, and your chance of being wrong is 98 percent. Each particular opposing view, however, is right with a chance of only 1.5 percent, slightly less than the chance of your view. Furthermore, if any one of these other people is right, and you are wrong, your view will have grave negative consequences, such as bringing about the premature end of civilization. That is, your view has grave negative consequences with probability 0.98. OK, your view is the best one available. But in this scenario, how intellectually arrogant should you be about the details of your beliefs? Above and beyond your basic loyalty to achieving good ends, how firm should your ideological views be?

The resulting tolerance and meta-rationality may make contractarian or contractualist accounts of politics easier to justify. Focusing on big events, and downplaying our ability to predict means–ends relationships, ought to move our political beliefs closer together. That is, it ought to be easier to reach 'agreement in principle', even if we continue to disagree on particular facts or how to interpret them.

The epistemic problem also suggests a new argument for constitutional democracy and the rule of law, as opposed to dictatorship. Democracy may make epistemic problems marginally less severe, and thus make it easier to choose better policies.[21]

[21] The arguments of Hayek, *The Fatal Conceit*, concerning the difficulty of central planning, are well known. That being said, many dictatorships rely on the forces of the

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

Most specifically, a greater reliance on rules, rule of law, and predictable public institutions – all central features of constitutional democracy – alleviate epistemic problems to some degree. Those changes would introduce greater predictability into our environments. Rules limit the ability of one madman to do harm, and in that sense they make identity reshufflings less important. It is no accident that Hitler, a powerful dictator, is such a common example in the literature on the epistemic problem. So the epistemic problem provides an additional justification for the rule of law. To the extent that politics is delinked from individual identities, the chance is lower that one person will matter too much or that our entire future will hinge upon a very small number of personal identities.[22]

This argument also suggests that the presence of a rule-bound democracy should alter the nature of a practiced pluralism. To the extent that predictability is a smaller practical problem, the epistemic critique militates less against consequentialism. More specifically, we would have less of a reason to dismiss the importance of small events and small benefits. The rule of law implies that the potential benefits of small events are less likely to be overwhelmed by the long-run uncertainty of our more general forecasts. Under a dictatorship, in contrast, the case for 'big event consequentialism' is stronger.

We can imagine other cross-sectional results of a comparable nature. In some communities, for instance, long-run effects may be more bounded than in others. Consider the world prior to widespread international travel. Changes in the genetic composition of Greenland might not have much altered events in the rest of the world. Consequentialism for Greenland therefore should have placed greater stress on small events than, say, consequentialism for the territories of central Europe.

I do wish to stress, however, that the rule of law does not eliminate the epistemic problem. Rules diminish but do not eliminate long-run epistemic problems. Nor do rules eliminate the danger that individual leaders may matter too much. Given complex and non-linear chains of cause and effect, the long-run consequences of any rule will be very difficult to calculate. For instance, it remains difficult to predict exactly which sets of rules will prove most stable over time. Would the United States be more stable if it kept its current constitution, or should it move to proportional representation? Should we keep or abolish the

market more than they use central planning. In any case the argument in the text can supplement Hayek's analysis when relevant.

[22] In one of the earliest responses to the epistemic problem, John Stuart Mill, 'Whewell on Moral Philosophy', noted that it is easier to predict the effects of rules than to predict the effects of detailed individual policies.

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press

Electoral College? Which choices are most likely to sustain the rule of law over time? It is unlikely that we will ever have very firm estimates to answer these questions in every regard. We will have some slight idea of what is better or worse, but great uncertainties remain. So a greater reliance on the rule of law, no matter how useful or wise, weakens epistemic problems but does not eliminate them.[23]

tcowen@gmu.edu

[23] The author wishes to thank Bryan Caplan, Robin Hanson, Daniel Jacobson, Kevin McCabe, Nathan Nobis, Tom Round, David Schmidtz, and Alex Tabarrok for useful comments and discussions.

https://doi.org/10.1017/S0953820806002172 Published online by Cambridge University Press