Burch-Brown, J. M. (2014). Clues for consequentialists. *Utilitas*, 26(1), 105-119. <https://doi.org/10.1017/S0953820813000289>

Publisher's PDF, also known as Version of record

Link to published version (if available):
[10.1017/S0953820813000289](https://doi.org/10.1017/S0953820813000289)

[Link to publication record in Explore Bristol Research](#)
PDF-document

Copyright © Cambridge University Press 2014

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Utilitas

<http://journals.cambridge.org/UTI>

Additional services for **Utilitas**:

Email alerts: [Click here](#)

Subscriptions: [Click here](#)

Commercial reprints: [Click here](#)

Terms of use : [Click here](#)

Clues for Consequentialists

JOANNA M. BURCH-BROWN

Utilitas / FirstView Article / January 2014, pp 1 - 15

DOI: 10.1017/S0953820813000289, Published online: 02 January 2014

Link to this article: http://journals.cambridge.org/abstract_S0953820813000289

How to cite this article:

JOANNA M. BURCH-BROWN Clues for Consequentialists . Utilitas, Available on CJO 2014 doi:10.1017/S0953820813000289

Request Permissions : [Click here](#)

Clues for Consequentialists

JOANNA M. BURCH-BROWN

University of Bristol

In an influential paper, James Lenman argues that consequentialism can provide no basis for ethical guidance, because we are irredeemably ignorant of most of the consequences of our actions. If our ignorance of distant consequences is great, he says, we can have little reason to recommend one action over another on consequentialist grounds. In this article, I show that for reasons to do with statistical theory, the cluelessness objection is too pessimistic. We have good reason to believe that certain patterns of action will tend to have better consequences, and we have good reason to recommend acting in accordance with strategies based on those advantageous patterns. I close by saying something about the strategies that this argument should lead us to favour.

In his 2000 paper entitled ‘Consequentialism and Cluelessness’, James Lenman argues that consequentialism can provide no basis for ethical guidance.¹ Our actions have profoundly far-reaching consequences, only a small portion of which are within view when we make our decisions. There is every reason to think that overall consequences will largely be determined by unforeseeable ones, and there is no reason to believe one thing or another about how these unforeseeable consequences will turn out. In our epistemically impoverished states, maximizing subjective expected utility is hardly more promising as a strategy for bringing about better consequences overall than, say, choosing acts randomly or choosing the acts that look worst.²

For reasons related to statistical theory, Lenman’s account is too pessimistic. We have good reason to think that murder tends to have terrible consequences. We also have grounds for making many other consequentialist judgements. Of course, we will often make mistakes, including believing that we have a more accurate view of

¹ J. Lenman, ‘Consequentialism and Cluelessness’, *Philosophy and Public Affairs* 29 (2000), pp. 342–70.

² See also: G. E. Moore, *Principia Ethica* (London, 1903), ch. 5, secs. 93–100; S. Kagan, *Normative Ethics* (Boulder, 1998), p. 64; D. Parfit, *Reasons and Persons* (Oxford, 1987); R. Frazier, ‘Act-Utilitarianism and Decision Procedures’, *Utilitas* 6 (1997), pp. 241–8; E. Wiland, ‘Monkeys, Typewriters, and Objective Consequentialism’, *Ratio* 18 (2005), pp. 352–60; F. Howard-Snyder, ‘The Rejection of Objective Consequentialism’, *Utilitas* 9 (1997), pp. 241–8; E. Mason, ‘Consequentialism and the Principle of Indifference’, *Utilitas* 16 (2004), pp. 316–21; D. Miller, ‘Actual-Consequence Act-Utilitarianism and the Best Possible Humans’, *Ratio* 16 (2003), pp. 49–62; G. Lang, ‘Consequentialism, Cluelessness, and Indifference’, *The Journal of Value Inquiry* 42 (2008), pp. 477–85. An account that might be interesting to compare with my own is C. Hare, ‘Obligation and Regret When There is No Fact of the Matter About What Would Have Happened if You Had Not Done What You Did’, *Noûs* 45 (2011), pp. 190–206.

the consequences than we do. Nevertheless, I shall argue that we have more reason for optimism than Lenman supposes.

I begin by introducing the cluelessness objection, along with the limitations of one popular response. I then present some basic statistical theory. This theory suggests that if we repeat similar actions many times under similar circumstances, the overall consequences will tend to be determined by whichever effects are *systematically* associated with the sort of actions we have taken, and the innumerable other influences will tend in the long run to wash out. Thus our ability to give good consequentialist guidance does not depend upon knowing all of the diverse causal factors that will influence the outcome of isolated actions. Rather, it depends upon knowing the main systematic effects of the sorts of actions that we are considering, and then acting on the basis of those systematic effects.

When it comes to actions for which we cannot reasonably hope to evaluate the systematic effects, we have little reason for optimism. Thus my argument by no means justifies an indiscriminate assumption that unforeseeable consequences will balance out, nor does it suggest that consequentialists have a strong basis for guidance in all circumstances. It favours a more limited conclusion: we have a good basis for providing consequentialist guidance just in so far as we are in a position to evaluate the systematic effects of the sorts of actions we are considering.

CLUELESSNESS

Lenman worries that, for many actions, the portion of consequences we foresee is minute compared to the portion we do not. To make his case, Lenman focuses on acts that have highly unpredictable influences upon which future agents come to exist. Drawing on various examples concerning the reproductive choices of Hitler's ancient ancestors, Lenman argues that many acts influence which people come to exist, and that these acts are implicated counterfactually in the vast arrays of good and bad things that those future agents get up to. The agents who come to exist if I choose *A* may get up to very different things from the agents who come to exist if I choose *B*. Surely the difference in good produced by one line of descendants and another will be of greater magnitude than the foreseeable expected advantage of my act.

Sometimes our influence on which agents come to exist will be direct, as when we make decisions about whether and when to have children, or when we make a concerted effort to save lives by contributing to UNICEF or Oxfam. At other times, the influence will be indirect. Say that Ida decides not to go to a New Year's party. Had she gone, she would have introduced her friend Max to a fellow attendee, who would

have recommended a particular bed and breakfast for a holiday in the hills. At that bed and breakfast, Max would have fallen in love with the gardener next door, and they would have gone on to raise a family together. Thus Ida's decision not to go to the New Year's party has far-reaching repercussions that are quite impossible for her to foresee; many of our acts will have such indirect repercussions, and we often cannot tell in advance which acts these will be.

Lenman's most prominent example is of two people living many centuries ago. Bandit Richard is on a murderous rampage, but he decides in a moment of conscience to spare the life of a young woman, Angie. Having been spared by Richard, Angie lives and goes on to have a child. A hundred generations later, Adolf Hitler turns out to be one of her descendants. Had Richard carried on with his rampage, the Holocaust might never have happened. In objective act-consequentialist terms, therefore, Richard's choice to spare Angie has marks against it – a disquieting suggestion. Yet, Lenman says, the consequentialist is in no position to say one thing or another about Richard's action, for who knows what would have happened had Richard gone on to kill Angie? Perhaps, he says, killing Angie would have led to the birth of Malcolm the Truly Appalling. The unforeseeable consequences of either murder or not murder are so vast that the consequentialist has almost nothing to go on. With no grounds on which to reason about the balance of these hidden effects, Lenman argues, we should resist the temptation to assume that they will behave in any well-mannered way, such as tending on the whole to balance out.

Although his headline example involves extreme events, Lenman's worry is not about bad luck in rare cases. It is about our pervasive ignorance at every scale of action, from everyday common encounters to large-scale strategic projects. Ending smallpox has overwhelmingly positive foreseeable effects. But the unforeseeable effects of such large-scale projects are also enormous, involving the survival or death of entire populations of future agents. Given how much is hidden from us concerning the overall actions of entire populations of possible people, the consequences we foresee will remain dwarfed by those we do not.³ On Lenman's account, if we ought to prevent smallpox or prohibit murder, it is because of non-consequentialist moral requirements – not because we can reasonably claim to have evaluated the overall effects.

³ Tyler Cowen suggests that cluelessness only undermines small-scale actions, and not large-scale ones. But it seems reasonable to worry that larger-scale actions will tend to have comparably greater hidden consequences, in which case it is not clear that Cowen's argument overcomes the cluelessness objection. T. Cowen, 'The Epistemic Problem does not refute Consequentialism', *Utilitas* 18 (2006), pp. 383–99.

SUBJECTIVE CONSEQUENTIALISM AND CLUELESSNESS

In light of epistemic considerations, many consequentialists favour a move to subjective consequentialism. Whereas objective consequentialists suggest that the rightness of an act depends on its actual consequences, subjective consequentialists focus on reasonable beliefs. They hold that right actions are those which a reasonable and conscientious agent would identify as having the best subjectively expectable consequences.⁴

In calculating subjective expected utility, the agent discounts the value of each possible outcome in proportion with how confident she is that it will take place. When it comes to outcomes of which she is entirely ignorant, the rational thing is to equivocate – to treat all such outcomes as equally likely, and thus without influence in her decision. Indeed it is sometimes argued that we are obliged, as a requirement of rationality, to treat irredeemably hidden consequences as equally likely, in much the same way that we are obliged to ignore the possibility that the sun will not rise tomorrow.⁵

Subjective consequentialism has an appealing practicability since it associates rightness with a standard that is epistemically accessible to agents, rather than being determined by external facts to which an agent may or may not have access. Moreover, focusing on an agent's reasonable beliefs allows evaluation of acts to be closely related to evaluation of blameworthiness and responsibility. This is attractive, since it means that a subjective consequentialist standard might play a unifying role in relation to many of the subjects in which we are interested when we come to studying ethics.

Nevertheless, subjective consequentialism is made less attractive by epistemic worries. If we are largely ignorant of what will tend to be for the better, it is hard to see what reason there could be for us to follow subjective consequentialist recommendations. A theory holding that we should *try* to bring about better consequences will be more attractive if we have reason to believe that our efforts will tend to pay off in the long run.

Of course, we may be tempted to think that in the absence of evidence one way or another, it *will* generally be the case that our subjective

⁴ Such views have been discussed at length and by many, but see for instance F. Howard-Snyder, 'It's the Thought that Counts', *Utilitas* 17 (2005), pp. 265–81. For an in-depth study of perspectives on subjective rightness, see H. M. Smith, 'Subjective Rightness', *Social Philosophy and Policy Foundation* 27 (2010), pp. 64–110.

⁵ See e.g. Mason, 'Consequentialism and the Principle of Indifference'; W. Shaw, *Moore on Right and Wrong* (Dordrecht 1995), pp. 114–20; D. Dorsey, 'Consequentialism, Metaphysical Realism and the Argument from Cluelessness', *Philosophical Quarterly* 62 (2012), pp. 48–70. See also W. Sinnott-Armstrong, 'Consequentialism', *Stanford Encyclopedia of Philosophy* (2011).

expectations will reflect the way things are with the world, and thus that maximizing subjective expected utility is indeed a good strategy for bringing about better consequences.

Alastair Norcross argues for this optimistic conclusion in his paper 'Consequentialism and the Unforeseeable Future'.⁶ In his primary example, a capricious ruler requires you to decide whether Smith will live or die, and a chance mechanism links your choice to the saving or killing of 7 million other people. A coin is flipped out of view. If you choose to save Smith and the coin lands heads, then 7 million other people will also be saved, while if it lands tails, 7 million others will be killed. Conversely, if you choose to let Smith die and the coin lands heads, then 7 million people will be killed alongside him, while if it lands tails, then the 7 million others will be saved. Even though chance is involved, Norcross concludes that you ought to save Smith, and to expect that making decisions of that sort over and over will probably yield a benefit overall.

Norcross proposes that our real-world decisions are analogous to this example. Although the outcomes beyond our control may be vast and depend on chance, if we act with the expected utilities, then the net effect in the long run should be beneficial.

The problem with Norcross's example is that his chance mechanism exploits physical symmetries. Knowing the set-up, we are led to assume that the physical probabilities are evenly balanced for and against us, because the outcome is being decided by the flip of an implicitly symmetrical, fairly balanced coin. Thus the example builds physical symmetry into the picture just where Lenman thinks we can make no symmetry assumptions.

That means that in Norcross's example our subjective probabilities unambiguously match up with the physical ones, thus motivating his correct inference that maximizing subjective expected utility in this case will reliably bring about better consequences. But in decision situations in which we do not have access to similar evidence about objective probabilities, our subjective probabilities will be evenly distributed out of ignorance, not out of positive evidence about how the world is; and so it seems that we have little basis for believing that our choice of acts will yield better consequences than choosing at random.

To see why this is problematic, consider what happens if we drop the assumption that the coin is evenly balanced. Once we do this, then the subjective probabilities and objective probabilities come apart.

⁶ A. Norcross, 'Consequentialism and the Unforeseeable Future', *Analysis* 50 (1990), pp. 253–6.

If the coin is biased just slightly, at .49 heads and .51 tails, then saving Smith will have an objective expectation of *killing* 139,999 people, whereas letting him die will have an objective expectation of *saving* 139,999 people. If we choose to save lots of Smiths under these circumstances, we will probably wind up letting hundreds of thousands of people die. Given that the error systematically works against us in this case, we would have a much *better* chance of saving lives by choosing at random, rather than choosing on the basis of foreseeable consequences.

As this example suggests, it is important not to take ignorance of unforeseeable consequences as evidence that they will tend to balance out. As Michael Strevens puts it, 'the nature of the world is independent of our epistemic deficits. The fact we do not know anything about *A* does not constrain the way things are with *A*.'⁷ If we do not have reason to think that the coin is fairly balanced, and we do not have evidence about the direction in which it is biased, then we will not be justified in expecting our recommendations to yield better outcomes – we will just be justified in admitting ignorance.

However, there is a second lesson to be found in this example, which is that we often assign equal probabilities not out of mere ignorance, but as an expression of reasonable background beliefs. It is not unreasonable to think that the coin is likely to be fair or indeed biased towards heads – after all, most real coins are. Since we know something about what coins are generally like, it makes sense in Norcross's example to think that we have a reasonable basis on which to make a decision. Of course, we also know something else about this case, which is that it is a set-up by a capricious dictator. Capricious dictators are likely to take pleasure in other people's discomfort. That gives us less reason to think that the coin is equally balanced, leaving the example rather less benign. But the general point holds, regardless of which of these considerations we find more compelling. In deliberating about what to do, we often draw on background beliefs which give us some reasonable basis for assigning probabilities.

The question, then, is whether the consequentialist can ever justify holding optimistic background beliefs about the far-reaching consequences of our actions (such as the belief that chancy and unforeseen consequences will tend to wash out) given the vast arrays of inscrutable causal influences that will shape those consequences. In the rest of this article I will argue that we do have a reasonable basis for some such beliefs.

⁷ M. Strevens, 'Inferring Probabilities from Symmetries', *Nous* 32 (1998), pp. 231–46.

STATISTICAL THEORY, PREDICTION AND PATTERNS

Statistical theory gives us greater reason for optimism about predicting consequences than Lenman supposes. The key lies in switching focus from individual events to broader patterns, and probabilities based on them. Predicting patterns is often less demanding than predicting isolated events. The reason is that the diverse causal effects which make it difficult to predict individual events often interact to yield surprisingly simple overall patterns.

Consider, for example, one of the most important and foundational results of statistical theory, called the Central Limit Theorem (CLT). The CLT tells us that if something we are measuring is the additive result of a vast number of small and independent influences, then it will tend to follow what is called a normal (Gaussian) distribution, or, in common parlance, a bell-curve. A normal distribution is shaped roughly as shown in Figure 1.

Figure 1. (Colour online) A histogram following a normal distribution.

That the distribution is normal tells us about how outcomes will be dispersed around the mean, e.g. what percentage of the outcomes will be at any given distance from the average. In a normal distribution, extreme events are possible but unlikely, and equally likely in either negative or positive directions. As it happens, this distribution is ubiquitous: a great many phenomena can be thought of as resulting from the interaction of innumerable, roughly independent effects.⁸

⁸ Good introductions to the CLT include: M. W. Crofton, 'On the Proof of the Law of Errors', *Philosophical Transactions of the Royal Society London* 160 (1870), pp. 175–87; S. Page, *Diversity and Complexity* (Princeton, 2011); H. Tijms, *Understanding Probability*:

Various other central limit theorems alter one or more of the conditions of the CLT, in order to define under what circumstances a parameter will be characterized by another sort of distribution. For instance, some important variations introduce interdependence between the contributing factors, so that if one factor is high this affects the likelihood of another factor being high as well. Introducing dependence between variables often results in much greater spread, or variance, and can have other effects as well.

The CLT and cognate results do not in themselves solve the problem, but they do point towards an important lesson. The *mere* fact that consequences are shaped by vast arrays of causal influences does not mean that their distributions will be entirely erratic and unpredictable. On the contrary, it is the very diversity and variety of influences that turns out to be helpful: what allows a simple macro-level pattern to emerge is often precisely the fact that the outcomes are determined by the combined influence of a vast and diverse array of causal effects.⁹ We should not assume that we are clueless about an overall *pattern* of events just because the outcome of each individual event will be shaped by innumerable causal influences, each of which is quite unpredictable on its own.

This brings us to a second foundational theorem from statistics, known as the Law of Large Numbers. The Law of Large Numbers explains more fundamentally why we can gain power in prediction by focusing our attention on broader patterns, and then assigning probabilities to individual events on the basis of these broader patterns.

If we try to predict the outcomes of an isolated few events, chance influences mean that the resulting patterns may be quite wild and unpredictable. But the Law of Large Numbers tells us that when it comes to the outcomes of very large numbers of similar events, the chance influences will tend to wash out, leaving the overall pattern of outcomes following the underlying probability distribution, whatever that may be.¹⁰

Chance Rules in Everyday Life (Cambridge, 2007); W. J. Adams, *The Life and Times of the Central Limit Theorem* (London, 2009); L. LeCam, 'The Central Limit Theorem around 1935', *Statistical Science* 1 (1986), pp. 78–91; S. L. Zabell, *Symmetry and its Discontents* (Cambridge, 2005).

⁹ M. Strevels, *Bigger than Chaos: Understanding Complexity through Probability* (London, 2003). See also S. Page, *Diversity and Complexity* (Princeton, 2011).

¹⁰ An important assumption must be made if we are to apply the Law of Large Numbers: we must assume that we are dealing with the sorts of events for which it is appropriate to reason as if they are characterized by an underlying probability distribution (i.e. where it yields predictive power to do so). Just how to determine which events are reasonably treated as having an underlying probability distribution is a matter of some dispute. One proposal is that simple behaviour will emerge for complex systems when the micro-properties affecting the outcome are sufficiently numerous and

It is because of the Law of Large Numbers that we can predict accurately the relative frequencies of outcomes that would result from rolling dice a great many times. It is also because of this law that we can improve our chances of success on individual rolls by ascribing to each roll the probabilities characterizing the overall ‘urn’. If we know the macro-probabilities associated with taking similar acts many times under similar circumstances, then we will tend to enjoy successful prediction by regularly assigning those macro-probabilities to individual events.¹¹

The same holds true when we switch to thinking about more demanding patterns. It is notoriously difficult to predict weather more than a short while in advance. Weather systems are chaotic, which means their behaviour is extremely sensitive to initial starting conditions. Minute differences in initial conditions can lead to dramatic differences in the resulting weather. But the *climate* does not share the weather’s level of unpredictability. The climate can be understood as the statistical properties of the weather – it is the ‘urn’ from which are drawn specific weather events. We know a great deal about the composition of the urn, and can make strong predictions about weather events on that basis, even if we are not good at predicting individual draws. Moreover, we can improve our prediction for individual weather events by making inferences based on the relative frequencies of the urn as a whole.

As with dice and weather, so with actions – to an extent. It is easier to predict overall consequences of patterns of action than it is to predict precise consequences of individual acts. The eventual consequences of each act will be influenced by a vast array of separate effects. If we try to predict the outcome of a given act, we are likely to fail, because we are unlikely to surmise accurately the state of each of the factors influencing the outcomes. We are more likely to succeed if we concentrate on broad patterns of effects following from similar actions (actual or hypothetical), and then use these as the basis from which to assign values to individual acts.

We are at a disadvantage with actions, though, compared with weather and dice, for a number of reasons. The first is that whereas we have opportunities to observe symmetries of dice and patterns of

independent, or with interdependencies that for one reason or another tend to wash out. We see below that this happens in the case of climate and dice – indeed, a washing-out of micro-level influences must happen anytime outcomes in the long run follow a clear probability distribution. Specifying precisely under what conditions this washing-out will take place is, however, difficult. Thanks to Michael Strevens for discussion on these points. Interested readers might see Strevens, *Bigger than Chaos* and also Page, *Diversity and Complexity*.

¹¹ See Strevens, *Bigger than Chaos*.

weather, we do not have the opportunity to observe instances of the total consequences of actions. We only have examples of the *discernible portion* of the effects of past actions – that is, the effects for which the causal relationships are obvious enough for us to pick up on. There may be many systematic effects that we are not in a good position to see. They might follow much later, or the causal structures linking our acts to these effects might be such that the links between our acts and their consequences are, as Kant puts it, impenetrably obscured from view.¹² For this reason, William Shaw overstates his criticism when he says that G. E. Moore's highly cautionary and Lenman-like account of the problem of practical ethical guidance displays a sort of Humean scepticism about induction.¹³ For that to be the case, we would need past instances of total consequences from which to infer forward, which we do not have.¹⁴

Consider, for example, that it has taken extensive research to uncover many of the likely environmental effects of increasing CO_2 emissions. Many systematic effects have been obscured from view until relatively recently, and our understanding of the particular, local-level ecological changes likely to result from anthropogenic climate change remains weak in many respects. Another example concerns the fact that we are often ignorant of the macro effects of our everyday micro-level economic choices. Patterns of consumption often have far-reaching economic, social and environmental impacts that are quite invisible to the consumer. Such examples suggest that our oft-repeated actions may have obscure, systematic and non-negligible effects that we cannot observe due to the complexity of causal processes.

Another complication is that although human beings may take or reject similar actions a large number of times, they will rarely be acting in just the same circumstances. Diversity in social context means that superficially similar acts may be systematically different in their effects: a lie told in one circumstance may be quite different from a lie told in another. Similar acts that are taken in different contexts or towards different ends may need to be appraised in different ways. Sometimes, it will be our actions themselves that cause context and culture to change. Telling the truth at one point in time may alter

¹² I. Kant, *Critique of Practical Reason* (1788). ch. 1.8, theorem IV, remark II.

¹³ W. Shaw, *Moore on Right and Wrong* (Dordrecht, 1995).

¹⁴ Lenman interprets Moore as being optimistic about action guidance. I do not think that is right. The passage that Lenman cites comes in the midst of a highly cautionary account, with Moore vividly portraying the scale of our ignorance. Moore argues that the assumptions that must hold if we are to give action guidance – namely that unforeseeable consequences will not overturn foreseeable advantages – are large ones which have never been shown to be true, and that without them we shall have no hope for practical ethics. See chapter V of *Principia Ethica*, and in particular sections 93–100.

the social context in which people later make their choices around truth-telling. These considerations raise questions about which act descriptions are most appropriate – a non-trivial point to which I will return shortly.

Finally, one of the most serious challenges to guidance for consequentialists is that although we may have an accurate impression of many systematic effects of our actions, we may at the same time have a profoundly inaccurate impression of the broader context in which our acts are being taken, and thus of the larger significance of the consequences we do foresee. Oedipus was not undermined fundamentally by a failure of prediction.¹⁵ He simply acted in ignorance of the identities of his mother and father. This missing information radically affected the significance of his actions. The fact that we often have a limited grasp of the broader contexts of our actions is a serious challenge indeed, and one that my argument here does nothing to address.¹⁶

These are significant challenges, and they are certainly great enough to mean that consequentialists should be circumspect in many of their claims. However, it would be too pessimistic to say that these challenges leave us clueless. In the final section of this article, I will make some brief suggestions about the practical implications of my argument for consequentialists.

IMPLICATIONS FOR CONSEQUENTIALISTS

Whether we will be successful in tending to bring about better consequences will depend on the extent to which we have identified the effects that systematically tend to result from adopting one sort of action or strategy rather than another.¹⁷ We can say that an event *B* is a systematic effect of *A* if *A* has a tendency to bring about or promote *B*, in the positive case, or a tendency to inhibit *B* in the negative case. What will undermine our prospects for bringing about better consequences is if the evaluands that look better to us in subjectively expectable

¹⁵ J. Bennett makes this observation as well. J. Bennett, *The Act Itself* (Oxford, 1998). p. 53.

¹⁶ Thanks to Bob Adams for reminding me of the seriousness of this worry.

¹⁷ I am using 'strategy' as an umbrella term, much in the way that Lenman uses 'policy' (p. 343). One sort of strategy (or policy) would be directly to attempt to maximize expected utility; another would be to adopt a given rule or practical principle; and another would be to develop a particular aspect of one's character. A strategy need not be conscious or intentional (it could simply denote how one goes about things) and it need not aim at producing good consequences. This is the familiar point that being affectionate out of love for one's friend does not aim at the greater good, but a consequentialist can recommend it as a good way to go about things.

consequences have sufficient hidden systematic effects working against us.

We have little chance of identifying the systematic effects of individual, maximally specified acts. If we assume determinism for a moment, then the systematic effects of a fully specified act would just be its actual consequences. (If we could specify fully the starting conditions and the causal laws, then the objective probabilities would be 1 for actual consequences, and 0 for all other consequences.) Such calculations are beyond us. One is hopelessly in the dark when it comes to predicting the systematic effects of this or that fully specified instance of not-murder. The presumption, in the case of individual instantiations, should be of substantial ignorance.

When we switch to reasoning about broader patterns, however, our capacities for prediction improve. In selecting our actions, therefore, we should reason using descriptions which are general enough to allow us to make reliable evaluations, but narrow enough that we do not exclude useful information. Refine the description too much, and one will lose the advantage one gains from thinking about statistical patterns instead of individual, maximally specified acts. Leave the description too broad, and one will leave out information that would improve prediction.

Consider this example. Jim is trying to decide whether to lend money to his friend, Johnny.¹⁸ Lending money to friends is often a helpful thing to do; but Johnny has told Jim that he wants to buy a gun and get back at Jack. Thus Jim has good reason to think that this particular instance of lending money to friends is likely to go poorly – it will be far to the left on the ‘lending money to friends’ distribution. There is no reason to think that Jim should lend money to Johnny, since he knows enough to refine the description and foresee that the act will probably turn out badly – instead, he should use the information he has. The epistemically best practical principles will be ones based on descriptions optimally tailored to circumstances – which is to say, ones that are specified just up to the point where including further information reduces rather than improves the reliability of our inferences.

It should be emphasized that for many sorts of acts the optimal level of description will remain rough, rather than fine. In addition to the epistemic reasons already discussed, there are reasons related to what have been called ‘expectation effects’ – the positive effects that result when members of society know that they can trust one another to act in certain ways and not in others, and the negative effects that result from the absence of such trust. Societies do better

¹⁸ Thanks to Tim Lewens for this example.

at supporting the flourishing of their members when individuals can trust exploitation to be prevented, reasonable agreements to be kept, and conflicts to be handled without violence. These effects would be undermined if individuals thought that others would be tempted to violate the relevant norms, or would be able to get away with doing so. Thus many important practical principles that consequentialists should favour should take the form of broadly applied social rules.¹⁹

What about actions that are fairly unique, not to be repeated many times? For these we must be more creative, but the consequentialist should still often have a good deal to say. Here is an example. In choosing policy responses to climate change, we face particular difficulties of prediction. Any response sufficient to reduce emissions substantially will have large-scale systematic effects which we are in no position to anticipate at the time of deliberation, for they are bound to involve the alteration of many interdependent, non-linear, complex adaptive systems. What policies should we favour?

In these cases, the best strategies might be ones recommending particular methods of risk management, informed by our best understanding of relevant causal systems. They might include principles like favouring diversified portfolios of responses where possible, so that individual systematic errors are less likely to dominate overall effects.²⁰ They might also recommend using cyclical processes of risk assessment; favouring approaches allowing for later adaptation in response to new information; requiring certain kinds of transparency and public accountability, and so on. Or it might be that some of these processes are too slow-moving to bring about effective changes, in which case there may be arguments to be made for less incremental approaches.²¹ These are debates to be had within the context of the particular risks being managed. The central point, for present purposes, is that for actions that are non-repeatable in important respects, we may often still be able to make reasoned judgements about more advantageous and less advantageous strategies.

¹⁹ On expectation effects, see J. Harsanyi, 'Morality and the Theory of Rational Behaviour', *Social Research* 44 (1977), pp. 623–56; J. Riley, 'Defending Rule Utilitarianism', *Morality, Rules, and Consequences*, ed. B. Hooker, E. Mason and D. E. Miller (Edinburgh, 2000), pp. 40–70; B. Hooker, *Ideal Code, Real World* (Oxford, 2000).

²⁰ Such considerations might add support to the 'stabilization wedges' approach to climate change mitigation, advocated in S. Pacala and Robert Socolow, 'Stabilization Wedges: Solving the Climate Problem for the Next 50 Years with Current Technologies', *Science* 305 (2004), pp. 968–72. Thanks to Robert Socolow for this suggestion.

²¹ Michael Oppenheimer, Robert Keohane and other audience members drew my attention to this point at the 'Communicating Uncertainty in Science' seminar (Princeton University, December 2012).

In formulating practical guidance, a consequentialist should eliminate as useless any strategies we cannot know how to implement (e.g. 'take the act with the best actual consequences') or that are in other ways impracticable (unrealistic or outside our capacities); eliminate strategies about whose overall consequences we cannot make reasonably reliable evaluations; and recommend following whichever remaining strategies have the best expectable consequences.

To judge whether we stand a chance of picking out the main systematic effects associated with following a particular strategy, we should reason about what sorts of causal mechanisms might lead to hidden systematic effects, and about how likely it is that such causal mechanisms exist. For many sorts of actions, the hidden systematic effects seem likely to be great, and our prospects for prediction will be poor. But for others they seem unlikely to be great, comparatively speaking, or if great they seem likely to be positively correlated with effects we foresee. For instance, it is conceivable that murders tend systematically to have positive results hundreds of years hence, but it is difficult to imagine what mechanism could cause this pattern. On the basis of the evidence that murder tends to have terrible consequences, and the difficulty of imagining countervailing, obscured systematic benefits, consequentialists are justified in claiming that murder is wrong. Even if they are unsure about outcomes in individual cases, they should be confident about the practical principle of prohibiting murder in all normal sorts of cases. It is with that principle that we should act.

Finally, the kinds of examples on which Lenman focuses involve effects that we really should *not* believe to be systematically associated with the sorts of acts being considered. In fact, the worry appears to be that we face overwhelming unpredictability precisely *because* the distant effects are not systematically associated with the acts in question, at least under the incompletely specified descriptions we use when we deliberate about them. The distant effects on identities of future agents arise randomly, as often with introducing friends as with failing to introduce them; as often with murder as with not-murder. But when we switch to reasoning about broader patterns of action, the fact that these hidden effects are not systematically associated with the acts we are considering turns out to work with rather than against us.

CONCLUSION

Lenman argues that the consequentialist is in no position to provide practical guidance; the profound unpredictability of consequences means that consequentialists have only the faintest of grounds for

recommending one action over another. I have argued that, for reasons related to statistical theory, Lenman's account is too pessimistic. The choice to save a life might occasionally lead, by a perverse chain of events, to the birth of a future tyrant, but this does not undermine our basis for giving consequentialist action guidance. Births of distant tyrants are probably not systematically associated with saving lives, any more than they are with failing to save lives. Consequentialists should recommend acting in accordance with the best strategies out of those for which we are able to evaluate the systematic effects. If similar decisions are repeated a large number of times, and if we are accurate in our appraisal of systematic effects, then acting on this basis should tend to lead to better consequences overall.²²

j.burch-brown@bristol.ac.uk

²² Many people have offered helpful comments on this article. I am particularly indebted to Brad Hooker, who has been exceptionally generous with his support, and to my doctoral supervisors, Tim Lewens and Hallvard Lillehammer. I would also like to thank Bob Adams, Chris Bertram, Luc Bovens, John Broome, Kam Chadha, Roger Crisp, John Cusbert, Andrew Ellis, Marc Fleurbaey, Stephen John, Jimmy Lenman, Brian McElwee, James Morauta, Martin Peterson, Onora O'Neill, Simon Rippon, Michael Strevens and Peter Vallentyne; and workshop audiences in Cambridge, London, Oxford, Princeton, and Bristol. For funding and institutional support, I thank the Cambridge Political Economy Society Trust; the Rausing Trust; the History and Philosophy of Science Department at Cambridge; the G. L. S. Shackle studentship at St. Edmund's College, Cambridge; and King's College, Cambridge. Finally, most of all I would like to thank Frank Burch-Brown, Carol Burch-Brown and Ann Kilkelly, and my partner William Baker.