

FRED FELDMAN

ACTUAL UTILITY, THE OBJECTION FROM IMPRACTICALITY, AND THE MOVE TO
EXPECTED UTILITY*

ABSTRACT.

Utilitarians are attracted to the idea that an act is morally right iff it leads to the best outcome. But critics have pointed out that in many cases we cannot determine which of our alternatives in fact would lead to the best outcome. So we can't use the classic principle to determine what we should do. It's not "practical"; it's not "action-guiding". Some take this to be a serious objection to utilitarianism, since they think a moral theory ought to be practical and action-guiding. In response, some utilitarians propose to modify utilitarianism by replacing talk of actual utility with talk of expected utility. Others propose to leave the original utilitarian principle in place, but to combine it with a decision procedure involving expected utility. What all these philosophers have in common is this: they move toward expected utility in order to defend utilitarianism against the impracticality objection. My aim in this paper is to cast doubt on this way of replying to the objection. My central claim is that if utilitarians are worried about the impracticality objection, they should not turn to expected utility utilitarianism. That theory does not provide the basis for a cogent reply to the objection.

Utilitarians are attracted to the idea that an act is morally right iff it leads to the best outcome. But critics have pointed out that in many cases we cannot determine which of our alternatives in fact would lead to the best outcome. So we can't use the classic principle to determine what we should do. It's not "practical"; it's not "action-guiding". Some take this to be a serious objection to utilitarianism, since they think a moral theory ought to be practical and action-guiding. Some utilitarians respond by revising the principle by appeal to expected utility. The classic principle is replaced by a new principle that says that an act is morally right iff it maximizes *expected utility*. They then suggest that since the revised principle is practical and action-guiding, the objection has been answered.

This pattern of reasoning (or something quite like it) can be found in the writings of Richard Brandt¹, J. J. C. Smart², Frank Jackson³, Bart Gruzalski⁴, and others.

Instead of responding to the objection by replacing the classic utilitarian principle with a principle requiring the maximization of expected utility, some philosophers respond by distinguishing between a *criterion* for rightness and *decision procedure* for rightness. They retain the classic principle as the criterion for rightness. They acknowledge that in many typical circumstances we don't know what is actually the right thing to do. But they then propose to supplement the classic utilitarian principle for rightness by adding a decision procedure to be used for practical purposes. The decision procedure makes use of the concept of expected utility.⁵ Perhaps it advises us to engage in a pattern of reasoning designed to identify the act that maximizes expected utility. The addition of this decision procedure thus creates a "package"

that is practical and action-guiding. As a result, the objection has been answered in a slightly different way.

Mark Timmons ascribes this second response to the utilitarian in his recent book *Moral Theory: An Introduction*.⁶ J. J. C. Smart makes use of something very like it in *Utilitarianism: For and Against*.⁷ Others have used it as well.

Many philosophers have responded to the objection in these ways, or in ways that are relevantly similar. Some propose to modify utilitarianism by replacing talk of actual utility with talk of expected utility.⁸ Others propose to leave the original utilitarian principle in place, but to add a decision procedure involving expected utility. What all these philosophers have in common is this: they move toward expected utility in order to defend utilitarianism against the impracticality objection.

My aim in this paper is to cast doubt on this way of replying to the objection.⁹ My central claim is that if utilitarians are worried about the impracticality objection, they should not turn to expected utility utilitarianism. That theory (in any of the suggested forms) does not provide the basis for a cogent reply to the objection.

1. SOME BACKGROUND

For purposes of illustration, let us make use of an example due to Frank Jackson.¹⁰ A medical doctor, Dr. Jill, is charged with the treatment of a delightful patient, John. John has some nasty disease. There are three available treatments for John's illness. One of them, which we may identify as B, is sure to cure the disease but will bring about some uncomfortable side effects. Dr. Jill knows this. She also knows, with respect to the other two treatments (which we can call 'A' and 'C') that one of them would cure John completely without side effects and that the other would be fatal. Unfortunately, she does not know which one is the excellent treatment and which one is the fatal treatment. So these treatments are alike in two important respects: each of them is such that Dr. Jill's evidence supports the belief that it might cure the patient completely without side effects; each of them is also such that that evidence supports the belief that it very well might be fatal. So (arbitrarily assigning some values to these various outcomes) the case looks like this:

a1. Prescribe A	maybe +100; maybe -100
a2. Prescribe B	certainly +75
a3. Prescribe C	maybe +100; maybe -100
a4. Prescribe nothing	certainly -50

Case 1

The classic utilitarian principle implies that Dr. Jill morally ought to prescribe the treatment that will in fact cure John without side effects. That is either A or C, but of course Dr. Jill does not know which it is. Someone sympathetic to a more subjectivist approach might say that what Dr. Jill morally ought to do depends upon the evidence currently available to her. Maybe she should choose the second best option in order to avoid the risks associated with the options that might be best but might also be terrible. Those who see it this way will say that the

reasonable thing for Dr. Jill to do in this situation is to prescribe the treatment that is guaranteed to cure the patient and that avoids the risk of killing him. That would be a_2 .¹¹

2. THE MOVE TO EXPECTED UTILITY

Reflection on cases such as this has led some philosophers to formulate a version of consequentialism based on the concept of expected utility. When advocates of this approach speak of *expected utility*, they are not thinking of the amount of utility that the agent (or anyone else, for that matter) expects the action to have. That's just a confusion brought about by the slightly misleading name. Rather, they are thinking of a number that can be defined in the following way.

Suppose some possible action, a , is one of your alternatives. Suppose that there are several possible outcomes, O_1, O_2, \dots, O_n , each of which might happen if a were performed. Suppose that for each outcome, O_i , there is an amount of value, V , associated with O_i . Suppose also that for each of these outcomes, there is a probability, $\text{Pr}(O_i/a)$, that is the conditional probability of outcome O_i on action a – in other words, this is the likelihood that outcome O_i will occur, if a is performed. The expected utility of a is the sum, for all these possible outcomes, of $V(O_i) \times \text{Pr}(O_i/a)$.¹²

As I indicated at the outset, there are several different ways in which one could make use of the concept of expected utility. One of these ways is discussed by Mark Timmons in his *Moral Theory: An Introduction*. After defining the concept of expected utility, Timmons formulates two principles (p. 124). They are:

ACU: An action A is *right* if and only if A has as high an actual utility as any alternative action that the agent could perform instead.

PCU: An action A is *right* if and only if A has as high an expected utility as any alternative action that the agent could perform instead.

Timmons explains the point of the distinction in this way (p. 124):

Since expected utility is something we can apparently calculate (at least in many cases), PCU, unlike ACU, represents a useful version of the basic utilitarian principle that satisfies the applicability standard.

He goes on to “propose PCU as a proper utilitarian decision procedure” (p. 125). He describes a case involving a charity event that can be held either indoors or outdoors. He specifies some possible outcomes (rain, no rain) and some values (+2,000, +400, -80) and some probabilities (.2, .8). He sketches the calculations and shows that PCU implies in this example that you should hold the charity event indoors. He seems to endorse this outcome when he says that this is “the most reasonable course of action given your limited knowledge” (p. 126).

I think Timmons is speaking loosely when he says that PCU is a “decision procedure”. PCU is not a procedure of any sort; it is merely a statement of alleged necessary and sufficient conditions for the rightness of actions. (And, given that acts deemed right according to ACU are distinct from acts deemed right according to PCU, we know at the start that they can't both be true.) Surely Timmons' actual point is rather that PCU points the way toward, or suggests, a

decision procedure. A decision procedure would have to be a *procedure* – that is, a series of steps one can take. And it would have to yield a *decision* – that is, a determination of the action that should be performed. So what is the decision procedure associated with PCU?

Perhaps the decision procedure associated with PCU is something like this:

1. List your alternatives.
2. For each alternative, write down its expected utility.
3. Determine which ones maximize expected utility.
4. Perform one of those.

A moment's reflection should make it obvious that Timmons could not have been thinking of this decision procedure. There are several obvious problems. For one thing, Timmons suggests that the use of the decision procedure would involve calculation; yet this decision procedure does not mention anything about calculation. It simply instructs the user to “write down” the expected utilities of his alternatives. A second problem is even more pressing: most of us, most of the time, do not know the expected utilities of our alternatives. So the procedure seems incomplete. Surely the intended procedure must give us a recipe for determining expected utilities before we write them down.

Perhaps this is the intended decision procedure:

1. List your alternatives.
2. List all the possible outcomes of the first alternative.
3. For each outcome of the first alternative, write down its value.
4. For each outcome of the first alternative, write down its probability given the alternative on your evidence.
5. For each possible outcome of the first alternative, multiply the value of the outcome by its probability on the first alternative.
6. Sum these products. This sum is the expected utility of the first alternative.
7. Repeat steps 2-6 for each of the other alternatives.
8. Identify the acts that maximize expected utility
9. Perform one of those acts.

Call this decision procedure ‘PCU(DP)’.

Timmons seems to be saying (on behalf of the utilitarian) that ACU is the criterion of right action. So what you actually should do is whatever would be best. However, if you don't know which of your alternatives maximizes actual utility, you should follow a decision procedure associated with PCU. I have suggested that the intended decision procedure must be something like PCU(DP). Timmons recognizes that following this procedure will not ensure that you end up doing the right thing, but he says that the utilitarian should advocate it anyway, since if you follow it you will at least avoid being a “proper object of blame” (p. 126).

This approach may seem to provide an attractive solution to the medical problem described in Case 1. The doctor knows that the patient will be cured either by treatment A or by treatment C; she does not know which it is. She also knows that the one that does not cure him will kill him. She also knows that treatment B will produce a satisfactory but not spectacular cure. She knows the values and probabilities of all these outcomes. Under these circumstances, it may seem that the right course of action for the doctor is to prescribe treatment B. Given some

natural assumptions, the current proposal supports this intuition. Here, for purposes of illustration, are some suitable assumptions: suppose that act a1 (the choice of treatment A) has two possible outcomes – O1 (complete cure with no adverse side effects) and O2 (death of the patient). Suppose the probability of each is .5. Suppose the value of O1 is +100 and the value of O2 is -100. Then the expected utility of a1 = 0. Similarly for a3 (the choice of treatment C). But the expected utility of a2 (the choice of treatment B) is +75, since it has just one possible outcome (cure with side effects) and that outcome has a value of +75. Therefore, if Dr. Jill were to follow PCU(DP), she would quickly reach the conclusion that a2 maximizes expected utility, and then she would perform a2. This is the result desired by the subjectivist consequentialist.

3. PROBLEMS WITH THIS PROPOSED USE OF EXPECTED UTILITY

In spite of its apparent success in Case 1, the move to expected utility is indefensible. Recall the alleged problem with which we started: some of us are initially attracted to the idea that right acts are ones that maximize actual utility; but ordinary people in ordinary circumstances often don't know which of their alternatives will in fact maximize utility; thus we need a form of utilitarianism that can be applied in real-life circumstances – something that is more “practical”. This was the context in which we moved toward forms of consequentialism involving expected utility. The current suggestion is that instead of groping in the dark for actual utility, we should instead just make use of a decision procedure that enables us to determine what maximizes expected utility. We will be morally blameless if we follow the procedure and perform the selected act.

But the implausibility of this idea should be obvious. The calculation of regular utility is problematic. To perform the calculation, one needs to know his alternatives, their total consequences until the end of time, and the values of those consequences. A daunting epistemic task! But the determination of expected utility is even more problematic. To perform the relevant calculations, one needs to know his alternatives, and for each alternative, one also needs to know every one of its possible outcomes, the actual value of each of those outcomes, and the probability of each outcome, given the act, on one's current evidence. One also needs to know some mathematical facts: the product, for each outcome, of its value and its probability; and the sum of those products. If the epistemic task in the actual utility case was daunting, this task is double-daunting.¹³ Clearly Timmons was wide of the mark when he said that PCU “satisfies the applicability standard”.

Even if we focus on the very simplest sort of case, the epistemic problems will be overwhelming. Let's look more carefully at the example involving Dr. Jill. Suppose that she's in a real-life situation roughly like the artificial situation sketched in Case 1. Let's consider what would happen if she were to attempt to use the decision procedure I called PCU(DP). According to this decision procedure, Dr. Jill must first list all of her alternatives. We have been imagining that her alternatives are:

- a1: Prescribe A
- a2: Prescribe B
- a3: Prescribe C
- a4: Prescribe no treatment.

But this is unrealistic. In fact, any real-world counterpart of Dr. Jill would have far more alternatives. How about aspirin? or acetaminophen? or ibuprofen? or any other available drug? I checked a website named “drugs.com” and I found that there are at least 24,000 drugs available for prescription in the US. Dr. Jill could prescribe any of those. She could prescribe large doses or small doses or multiple doses. She could also prescribe any combination of two of them; or any combination of three; etc. So even if we stick just with her alternatives involving the prescription of drugs licensed for use in the US, we can see that a real-life Dr. Jill in fact would have well over a million alternatives. If we add over the counter medications, holistic treatments, different possible prayers she could recite, etc. we see that there are in fact many millions of alternatives available to her. The idea that she should actually write something down for each alternative is simply preposterous. If she can’t list her alternatives, she cannot perform the first step of the decision procedure.¹⁴

Suppose somehow Dr. Jill manages to identify a few hundred of the most salient alternatives. Suppose we understand step 1 of PCU(DP) in such a way as to make it permissible for her to list not all her alternatives, but only the most salient ones. We move then to step 2. Now she must list, for the first alternative, all of its possible outcomes.

Earlier we assumed that there were just two possible outcomes for a_1 – O_1 (complete cure with no adverse side effects) and O_2 (death). But this too is completely unrealistic. I take a possible outcome of an action to be a total way the world might be if the action were performed. This of course is relativized to the body of evidence available to the agent around the time of the action. So we must ask ourselves this question: in light of the evidence available to Dr. Jill at the time of her decision, how many different outcomes of a_1 are epistemically possible for her – how many such outcomes are such that for all she knows they might take place if she were to prescribe treatment A? Clearly, the answer is that there are at least thousands of possible outcomes. Perhaps there are infinitely many such possible outcomes. Imagine some way the whole world could be after the performance of a_1 . Imagine describing one of these and then asking Dr. Jill this question: ‘Is it possible, given what you now know, that the world could turn out like *that* if you were to perform a_1 ?’. Clearly, there are thousands of cases in which her answer would have to be ‘yes; that outcome is possible given what I know’. If she is following PCU(DP), she will have to write down every one of these possible outcomes. Clearly, this cannot be done.¹⁵

Suppose again that this problem can be overcome. Perhaps Dr. Jill will just write down the most salient outcomes of her first alternative – how she will know which ones are salient is another question. But suppose she does it.

Next Dr. Jill must write down two numbers for each possible outcome of a_1 . The first number is the value of that outcome. There is, of course, considerable controversy about the evaluation of outcomes. We don’t know which axiological theory gives us the truth about the evaluation of outcomes. For purposes of illustration, let’s suppose some form of hedonism is true. Suppose that for each possible outcome, there is some number that represents the value of that outcome according to the true form of hedonism. It is absurd to suppose that Dr. Jill could know, or even guess, what these numbers are. If we understand value in this way, no ordinary human being would be able to follow the instruction that tells her to write down, for each possible outcome, its value. Dr. Jill cannot perform step 3 of PCU(DP).

Some may think that we make Dr. Jill’s efforts easier if we adopt a very “subjective” account of value. Perhaps we require no more than this: for each possible outcome of a_1 , Dr. Jill is to write down the number that represents how good *she now thinks* that outcome would be.

But even this is unrealistic. Imagine that Dr. Jill focuses her attention on one of the salient possible outcomes of a_1 . Call it 'O_{23,456}'. We ask her, 'what is the number, n , such that you now think that the value of O_{23,456} = n ?' Obviously, Dr. Jill will have to respond by saying that there is no such number. I cannot imagine why anyone would think that a real-life doctor would have a complete set of beliefs about the values of all possible outcomes (or even all the salient possible outcomes). It seems clear that if we define value in this highly subjective way, there simply are no such values in the vast majority of cases. People do not have the relevant beliefs.

A slightly less subjective conception of value might be proposed.¹⁶ We might say that there is a certain body of evidence available to an agent just prior to the performance of an action. For each possible outcome, O_i , of that action, there is some amount of value such that the agent's evidence supports the view that O_i has that value. If you have an extraordinary body of excellent evidence, this value might be equal to O_i 's actual value. If your evidence is scanty or misleading, this value might diverge dramatically from the actual value of the outcome. Following Zimmerman, we can call this the "probable value" of O_i (for the agent, at the time).

It's not clear to me that every outcome (or even every salient outcome) in fact has a probable value. Perhaps the agent's evidence simply does not determine any particular value for some outcomes. But even if every (salient) outcome has a probable value, there is no reason to think that ordinary agents, in ordinary epistemic situations, will be able to determine what this value is. Let us try to imagine the difficulty of the task. Suppose an omniscient onlooker is able to provide Dr. Jill with a complete description of one of the possible outcomes of one of her alternatives. Suppose (in vague outline) it goes like this: 'shortly after you prescribe aspirin, 325 mg. twice a day for two weeks, Jack goes to the pharmacy; he purchases the aspirin and takes the first pill. Later in the day, he loses the bottle and decides that Tylenol would be pretty much equivalent. He takes a Tylenol tablet at dinner time. Later that night he has an upset stomach, but by morning he feels much better. He feels so much better that he is motivated to express some affection for his wife. He and his wife then engage in sexual intercourse. They have a child who later grows up to be as bad as Hitler. Millions of innocent Episcopalians are sent to concentration camps. One of the Episcopalians is Wilber, who would have become even worse than Hitler if he had not been sent to the concentration camp, etc. etc.' Call this outcome O_{23,456}. Now we ask Dr. Jill to specify the number, n , such that her current evidence supports the conclusion that the value of O_{23,456} = n . Even if there is such a number, there is no reason to suppose that Dr. Jill knows what it is.¹⁷ Suggesting that she should just "write it down" seems seriously unrealistic.

As before, let us simply bypass this problem by supposing that there is some way for Dr. Jill to come to know the values of all the salient outcomes of all the salient alternatives available to her. Her job is still not done. Next she must write down, for each outcome, a number representing the probability of that outcome. More exactly, if a is one of her alternatives, and O is one of the possible outcomes of a , then there is some number, n , such that $\text{Pr}(O/a) = n$ on Dr. Jill's current evidence.¹⁸ This is just an objective fact about the relationship between that body of evidence and that outcome. Anyone who contemplates this dispassionately will recognize that no real-life agent in any interesting real-life situation of choice would be able to determine what these numbers are. Though Dr. Jill (in some sense) "has" the evidence, she is hardly in a position to say, for every outcome, O , of every alternative, a , what $\text{Pr}(O/a)$ is.

Perhaps someone will suggest that we adopt a more "subjective" account of probability. Perhaps the idea is that Dr. Jill is to write down the number, n , that represents how likely she

currently takes O to be given a . In this case, the probability of O given a on Dr. Jill's evidence is not an objective fact determined by intrinsic features of O and a and Dr. Jill's evidence. It is just Dr. Jill's purely subjective estimate of how likely it would be for O to ensue if a were performed.

As before, however, it seems unlikely that there are any such numbers. If Dr. Jill has never in fact contemplated all the details of the outcome, O, and she has not carefully surveyed the totality of the evidence available to her, she will have to admit that there is no number, n , such that she currently believes that the probability of O given $a = n$. If she were asked for such a number for each outcome of a , she would have to reply in most cases, 'Gee. I never thought of that. I honestly don't have any particular view about how likely that would be.' So, even if we take a highly subjective view of probabilities, no real life Dr. Jill would be able to complete step 4 of PCU(DP).

But let us again assume that this problem could be overcome. Let us assume that Dr. Jill somehow has arrived at views concerning the values and probabilities of all of the possible outcomes of her first alternative. All that remains, then, is a little elementary arithmetic. In order to determine the expected utility of a , she just needs to do some multiplication and some addition. Let us suppose, then, that she has settled on this list of values and probabilities for the possible outcomes of her first alternative:

Outcome	Value	Probability given a
O1	96.7	.108
O2	94.3	.27
O3	87.6	.003
O4	42.8	.0476
O5	-26.93	.032
O6	-27.4	.059
O7	7.69	.108
O8	3.48	.27
O9	862.7	.003
O10	24.3	.0476
O11	-62.39	.032
O12	-4.806	.059

..... etc. for about a thousand more possible outcomes of a

To find the expected utility of her first alternative, Dr. Jill must multiply probability times value for each possible outcome and then sum the resulting products. Given that there are perhaps more than a thousand value/probability pairs, it seems to me obvious that no real life Dr. Jill would be able to determine the expected utility of any alternative. Surely, if she were to set out to try in any serious way, her patient would be dead and buried before she could complete the task.

But even if she could determine the expected utility of one of her alternatives in this way, Dr. Jill's efforts would just be beginning. She would have to repeat this process for every available alternative. If there are a million such alternatives, she would have completed only a millionth of the project.

My conclusion is that no real life person in any interesting situation of moral choice could make use of PCU(DP) to determine the act that maximizes expected utility. Nor could she do it even if the values and probabilities were purely "subjective" – just the values and

probabilities she took the outcomes to have.¹⁹ In light of this difficulty, it seems to me clear that if impracticality is the problem, expected utility is not the solution.

4. WATERING DOWN THE PROPOSAL

It is interesting to note that many of the philosophers who seem to advocate the move to expected utility already recognize that their proposal really does not “satisfy the applicability standard” (in Timmons’ words). Timmons himself remarks parenthetically that we will be able to perform the necessary calculations only in “many cases”. Smart (1973, p. 39) explicitly says that in many cases it is not possible to assign numerical probabilities to particular events. He goes on to point out that in ordinary affairs when making prudential decisions we often make use of rough and approximate estimates of likelihood. He says (p. 40):

‘If we are able to take account of probabilities in our ordinary prudential decisions it seems idle to say that in the field of ethics, the field of our universal and humane attitudes, we cannot do the same thing, but must rely on some dogmatic morality, in short on some set of rules or rigid criteria. Maybe sometimes we just will be unable to say whether we prefer for humanity an improbable great advantage or a probable small advantage, and *in these cases perhaps we shall have to toss a penny* to decide what to do. Maybe we have not any precise methods for deciding what to do, but then our imprecise methods must just serve their turn.’²⁰

So it appears that at least some of the advocates of expected utility utilitarianism recognize that under ordinary circumstances we haven’t got the information or mathematical skills required to determine the expected utilities of our alternatives. Instead, they say that it is good enough for us to proceed in an “intuitive way”. We can just imagine the likely consequences and values; we can use “imprecise methods”; we can “toss a penny”. In short, recognizing that PCU in fact fails the test of applicability, these philosophers seem to say that it doesn’t matter; we can just guess, or follow our hunch, or choose randomly.

It almost seems that these philosophers must have forgotten why they are talking about expected utility in the first place. Recall that the problem arose because some critics of actual utility utilitarianism thought that there was an epistemic shortfall – we don’t know enough to identify the act that actually has the highest utility. Expected utility was introduced precisely to circumvent this alleged problem.²¹ Yet now, when they notice that it is just as hard – or perhaps harder – to determine expected utilities, the advocates of expected utility press on. Now they suggest that it’s OK to make guesses or to rely on “intuitive and imprecise methods”. But if that’s the best they can offer on behalf of expected utility, the advocate of actual utility surely can offer the same thing. Why can’t he also say that it’s OK to make guesses or to rely on intuitive and imprecise methods when trying to identify the act that maximizes actual utility? Why not say that ACU is the criterion of actually right actions, and that when you don’t know which of your actions maximizes actual utility, you may use imprecise methods, or guess, or just flip a coin?²²

5. A NEW ROLE FOR EXPECTED UTILITY

Some of the philosophers who have discussed the move toward expected utility apparently view the topic from a different perspective. Instead of advocating PCU as a criterion of objectively right action, or even as a hint of a decision procedure, they propose an even more limited role for expected utility. The role in question is suggested by Timmons' remark about the mitigation of our moral evaluation of individuals. He says that a person who performs an act that maximizes expected utility might do the objectively wrong thing, but nevertheless he 'may not be a proper object of blame' (p. 126).²³ This suggests a new role for expected utility.

Let us say that a person is "blameless" in performing an action provided that no one who understood his circumstances and the quality and nature of the evidence available to him could reasonably blame him, or reproach him, for having performed the action. Recall Case 1. Dr. Jill had reason to believe that either treatment A or treatment C would be best, but she also had reason to believe that trying either one would be very risky. If she were to prescribe treatment B instead of treatment A, and if this were to be in fact less successful than A would have been, she would have performed the objectively wrong act according to ACU. But we might say that in virtue of the risks involved in prescribing A, she is blameless in prescribing B. No one could justifiably blame her for her failing to prescribe the treatment that would have been best.

There are many ways in which we might try to link the concept of blamelessness with expected utility. Here is a principle that captures this idea:

BEU: An agent, S, is blameless in performing an action A if and only if A has as high an expected utility as any alternative action that S could perform instead.

Talk of expected utility in BEU should be understood as talk of the expected utility of the action *on the evidence available to the agent of the action around the time of the action*.

It seems to me that this proposed application of the concept of expected utility is no more plausible than the ones that have already been considered. It goes wrong in both directions, and for several different reasons.

i. I think it is important to keep in mind that the expected utility of a given action on a given body of evidence is an objective matter of fact. A person could happen by sheer luck to select the act that maximizes this value. Here's a case that illustrates this: suppose some doctor has a patient whose illness requires either treatment A or treatment B. Suppose the doctor is not sure which will be best. He gets out his manuals and looks up the possible outcomes of each treatment. He locates the information about their values and probabilities. He starts to do his calculations but (since he has been drinking too much lately) he keeps making mistakes. In the end, out of sheer frustration, he simply reaches into his drawer of free samples and hands over the first bottle of pills he locates. "Take this", he says. "Call me in the morning."

Suppose that simply as a matter of luck he happened to land on the treatment that maximizes expected utility on the evidence that the doctor had observed in his manuals. In that case, BEU implies that the doctor is blameless in prescribing as he did. This seems wrong.

ii. A second case works in the opposite direction. Suppose some doctor is a careful, responsible, and methodical calculator of expected utilities. Suppose he has a patient whose illness calls for either treatment A or treatment B. Each has various possible outcomes with their associated values and probabilities. Suppose the doctor has manuals containing this information.

The doctor carefully looks up the information and carefully performs his calculations. The calculations involve lots of numbers, so he checks them twice. Unfortunately, a stray burst of solar radiation affects his pocket calculator in such a way as to make it appear that a treatment A maximizes expected utility on his evidence. In fact, treatment B maximizes expected utility on that evidence. He then prescribes treatment A – the one that he thinks will maximize expected utility – and it turns out that the treatment successfully cures the patient with no adverse side effects. Surely the doctor is blameless in prescribing A. He did everything in a responsible manner and achieved the best results. However, BEU implies that he is not blameless in performing this act, because the act he performed in fact did not maximize expected utility on the evidence he had around the time of the action.

A natural revision of BEU suggests itself:

BEU': An agent, S, is blameless in performing an action A if and only if S believes that A has as high an expected utility as any alternative action that S could perform instead.

It's reasonable to suppose that BEU' generates correct results in the latest case. However, it generates apparently incorrect results in many other cases. For example: suppose a young and naive person has been brought up in a fairly strict religious tradition. He has been taught always to turn to the Rabbi for moral guidance. He's now troubled about a minor matter of religious observance. He asks the Rabbi for guidance. They pray together on this. The Rabbi says that he thinks the young man should resolve his dilemma in a certain way. The young man accepts and acts on the Rabbi's advice.

In this case, BEU' implies that the young man is subject to blame. That follows from the fact that he has not gathered any information about possible outcomes, probabilities, or values. He knows nothing about expected utility. We may even assume that he does not have the concept of expected utility. So he is not justified in thinking that the suggested course of action maximizes expected utility. But in spite of this it seems to me that he has behaved in such a way as to evade legitimate blame.

BEU' also goes wrong in the opposite direction. Suppose someone knows himself to be a very poor calculator of expected utilities. Suppose on some occasion of important moral choice he carelessly considers only a few dozen alternatives and, in each case, only a fraction of their possible outcomes. Suppose anyway he comes to believe that alternative a345 maximizes expected utility. Several wise advisors warn against a345, saying that it is too risky. He then performs a345 anyway and the results are horrible. Suppose by accident it turns out that he was right: given the evidence available to him just prior to acting, in fact no alternative to a345 had higher expected utility. BEU' implies that he is blameless in performing this act. That judgment seems wrong to me.²⁴

6. WHAT TO DO WHEN YOU DON'T KNOW WHAT TO DO

So this leaves us with a question: suppose a person does not know what he morally should do, but wants to keep himself above justified moral reproach. What should he do? Of course, the consequentialist will say that what he *should do* is whatever will be best. I should state the question more carefully: when a person does not know what he morally should do, but wants to

remain morally blameless, what course of action can be recommended for him?

In some circumstances it might be a good idea to seek expert advice. This might happen, for example, in a case in which a particularly wise and experienced expert was at hand and ready to give helpful advice. Obviously, it would not be a good idea in any case in which no experts were readily available and it was important to make a decision quickly. But where such an expert is handy, and where one does not have or understand the evidence oneself, or one knows oneself to be a bad calculator of expected utilities, one can sometimes avoid legitimate blame by allowing oneself to be guided by recognized experts.

In some circumstances you might just as well flip a coin. This might happen, for example, if the possible outcomes were of indeterminate but probably nearly equal and small value, and one had only the scantiest evidence about their likelihoods, and no experts were in the vicinity. In such a case it might be especially important to pick *some* course of action, and it might not be too important to pick the right one. In such a case one might remain blameless if one just picked more or less at random. This could be the case even if the selected alternative did not maximize expected utility and it was not the case that the agent believed that it would maximize expected utility. In such a case considerations of utility might be irrelevant. Better in such a case to pick at random than not to pick at all.

In still other circumstances it might make sense to abide by the requirements of conventional morality, or to try to figure out what Jesus would do, or to follow the advice of Jiminy Cricket and let your conscience be your guide.

Finally, it must be admitted that in some unusual circumstances it might be a good idea to spend some time trying to calculate expected utilities. This might happen, for example, in a medical case in which you know that there are just a few alternatives available, and in which it would be possible to look up the success rates and the frequencies and severities of the adverse side effects of all the available treatments. With that sort of evidence in hand (and assuming that no other factors were morally very important, and that time was not of the essence) it might make sense to evaluate possible outcomes, weight them for likelihood, do the arithmetic and choose the alternative that maximizes expected utility. This would be completely pointless in cases in which you cannot determine the possible outcomes or their values or their probabilities. But if the expected utilities can fairly readily be calculated, and no other method is available, and there are no experts on the scene, and you don't have to administer the drug any time soon, who could legitimately blame you for taking a few days with your pocket calculator?²⁵

I think it's important to recognize that none of the answers listed above could provide the basis for a universally suitable decision procedure for avoiding moral blame. Depending upon the details of the situation, different courses of action would be appropriate. But my aim here is not to offer or to defend any particular view about courses of action that will be guaranteed to keep one blameless. My aim, rather, is to indicate the implausibility of the notion that in every case the only way (or the best way) to ensure that one will be blameless is to perform the act that maximizes expected utility (or the act that you believe to be the act that maximizes expected utility). I grant, of course, that there might be cases in which considerations of expected utility are relevant to moral obligation or to keeping oneself above legitimate moral reproach. My point is that such cases are rare and unusual. It's a mistake to think that such considerations are always relevant.

APPENDIX: ODDIE AND MENZIES

In ‘An Objectivist’s Guide to Subjective Value’ Graham Oddie and Peter Menzies seem to be presenting a slightly different reaction to the epistemic objection to ACU. As they use the term, the “actual outcome value” of an action is ‘the value of the outcome the [action] would have produced had that [action] been chosen’ (Oddie and Menzies 1992, pp. 515-6). Accordingly, they use the term ‘actual outcome consequentialism’ to indicate the view that the agent ought to maximize the actual outcome value of acts. Clearly, this is a view relevantly like ACU. Oddie and Menzies quickly dismiss this view, saying that it generates incorrect results in “chancy moral set-ups”. They describe medical cases like the ones described here. Suppose a doctor happens to prescribe the medicine that has the best actual outcome though he did not know it would have that outcome and the probability of that good outcome was small. It was more likely that it would prove fatal. Then, according to Oddie and Menzies, ‘... the doctor subjected the patient to an unnecessary risk of fatality and, thereby acted wrongly’ (p. 516).

In my earlier sketch of expected utility consequentialism, I said that the expected utility of an act is the sum of the products, for all possible outcomes of that act, of the value of the outcome times the probability of the outcome on the act. I did not say much about probability. Oddie and Menzies describe two ways of understanding the concept of probability in this context. For each, we get a way of understanding expected utility.

One of these is characterized as “subjectivist”. Suppose that an agent has different degrees of belief, or “credences” concerning possible outcomes on possible actions. Let ‘Cr(O_j/A)’ represent the agent’s credence concerning outcome O_j given action A. If we allow these credences to play the role of probability, then we get a subjectivist conception of expected value:

$$SV(A) = \sum_j Cr(O_j/A) Value(O_j \& A)$$

Subjectivists will say that this is the sort of value that is of significance for moral theory, and that the “regulative ideal” is always to perform the option with maximal subjective value (Oddie and Menzies 1992, pp. 515).

The second way of construing value is “objectivist”. Suppose that chance is an objective feature of the world. Certain outcomes have certain objective probabilities on certain actions whether we know about them or even have beliefs about them. Let ‘Ch(O_j/A)’ indicate the objective chance of outcome O_j on action A. Then “objective value” may be defined this way:

$$OV(A) = \sum_j Ch(O_j/A) Value(O_j \& A)^{26}$$

Oddie and Menzies say that they will defend a form of objectivism that ‘claims that maximizing objective value is the moral agent’s regulative ideal’ (p. 515).

Oddie and Menzies recognize that ordinary agents generally do not know the objective chances of various outcomes. Some might say that this casts doubt on the relevance of the concept of objective value. But Oddie and Menzies claim that rational agents are required to abide by a constraint: they should aim to bring their credences into line with true objective chances, and thus to bring the values they assign to options into line with the true objective values (p. 518). Thus, the notion of objective value has primacy in moral theory. We really ought to do what maximizes objective value.

But in light of the fact that we generally don't know the objective probabilities of the outcomes, Oddie and Menzies attempt to formulate a "selection rule" that can be used by the ordinary moral agent. The selection rule does not tell the agent simply to do the act that maximizes objective value: after all, she does not know what act that is. In order to formulate this selection rule, Oddie and Menzies introduce the concept of "K-value". Let the K_i 's be a bunch of hypotheses about the probability of an outcome, B, on an action, A. For each of these K_i 's, you have some credence level – some extent to which you think that it properly represents the objective chance of B given A.²⁷ Let Cr be your credence function. Then the "credence mean estimate" of the chance of B given A is:

$$\sum_i Cr(K_i) Chi(B/A)$$

The K-value of an option is just the sum of the products, for all outcomes of that option, of value times the credence-mean estimate of the probability of the outcome given the action. They define it as follows:

$$KV(A) = \sum_j [\sum_i Cr(K_i) Chi(O_j/A)] Value(A \& O_j)$$

Oddie and Menzies then state their selection rule: 'we propose that the selection rule that the objectivist should endorse is: select the act with maximal K-value' (p. 520). They make a number of remarks concerning this rule. The one that is most relevant for present concerns is this (p. 521):

First, the rule counts as a practical guide to action since it employs only materials accessible to the mind of an ordinary agent: unlike objective value, the K-value of an option can be determined by an agent who lacks complete information about the true configuration of chances in the world.

In addition, they suggest that this concept of K-value may be relevant to attributions of blameworthiness. They say it is a 'plausible conjecture that an agent is judged worthy of blame for a wrongdoing if, or to the extent that, the K-value of the wrongful act was less than that of some other act available to her ... ' (p. 521).

If I have understood them correctly, Oddie and Menzies mean to be making several related claims. One is that a person's real moral obligation is to perform the act that maximizes what they call 'objective value'. But since ordinary people often don't know which of their alternatives actually maximizes this value, they offer a selection rule: do the act that maximizes K-value. They claim that this selection rule provides a practical guide to action since it "employs only materials accessible to the mind of an ordinary agent". Finally, they claim that it's plausible to suppose that agents are blameworthy for wrong actions insofar as those actions fail to maximize K-value.

If they are indeed saying that an ordinary agent should use this selection rule to determine what she should do, and that failure to perform the act that satisfies it will be blameworthy, and that this will somehow get around the epistemic problem suggested earlier in connection with ACU (or Oddie and Menzies's "actual outcome consequentialism"), then I think their view is deeply implausible. While in some sense the "materials" are accessible to the mind of an ordinary agent, no ordinary agent could gather those "materials" together in the required way and do the required calculations. It's about as plausible as saying that ordinary untalented people

should carve statues as beautiful as Michelangelo's "David" since after all they can easily locate a chisel, a hammer, and block of stone. The materials may be accessible, but the skills required to employ them are in short supply. I suspect that in fact there has never been, and could never be, an actual human being capable of calculating the K-value of any act. Keep in mind that it would require infinitely many calculations (one for each possible probability level for each outcome for each alternative).

My conclusion is that the concept of K-value might have some interest for some project in moral philosophy, but it has no relevance to the problem under consideration here: the formulation of a practically useful moral principle that ordinary agents can use to determine what they ought ("subjectively") to do; or at least what they can do in an effort to remain above moral reproach.

REFERENCES

- Bales, R. Eugene. (1982): 'Act Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure?', *American Philosophical Quarterly* 8: 257-265.
- Brandt, Richard. (1959): *Ethical Theory: The Problems of Normative and Critical Ethics*, Englewood Cliffs: Prentice-Hall.
- Carlson, Erik. (unpublished manuscript): 'Deliberation, Foreknowledge, and Morality as a Guide to Action'.
- Feldman, Fred. (1986): *Doing the Best We Can*, Dordrecht: Reidel.
- Frazier, Robert L. (1994): 'Act Utilitarianism and Decision Procedures', *Utilitas* 6: 43-53.
- Gibbard, Allan. (1990): *Wise Choices, Apt Feelings*, Cambridge, MA: Harvard University Press.
- Gruzalski, Bart. (1981): 'Foreseeable Consequence Utilitarianism', *Australasian Journal of Philosophy* 59: 163-176.
- Howard-Snyder, Frances. (1997): 'The Rejection of Objective Consequentialism' *Utilitas* 9: 241-248.
- Jackson, Frank. (1991): 'Decision-theoretic Consequentialism and the Nearest and Dearest Objection', *Ethics* 101: 461-482.
- Kagan, Shelly. (1998): *Normative Ethics*, Boulder, CO: Westview Press.
- Lenman, James. (2000): 'Consequentialism and Cluelessness', *Philosophy and Public Affairs* 29: 342-370.
- Mason, Elinor. (2003): 'Consequentialism and the "ought implies can" principle', *American Philosophical Quarterly* 40: 319-331.
- Mason, Elinor. (2004): 'Consequentialism and the Principle of Indifference', *Utilitas* 16: 316-321.
- Mill, J.S. (1863): *Utilitarianism*.
- Moore, G. E. (1912): *Ethics*, London: Oxford University Press.
- Oddie, Graham and Menzies, Peter. (1992): 'An Objectivist's Guide to Subjective Value', *Ethics* 102: 512-533.
- Smart, J. J. C. (1973): 'An Outline of a System of Utilitarian Ethics', in Smart, J. J. C. and Williams, B., *Utilitarianism: For and Against*, Cambridge: Cambridge University Press.
- Strasser, Mark. (1989): 'Actual versus Probable Utilitarianism', *Southern Journal of Philosophy* 27: 585-597.
- Tännsjö, Torbjörn. (1998): *Hedonistic Utilitarianism*, Edinburgh: Edinburgh University Press.

Timmons, Mark. (2002): *Moral Theory: An Introduction*, New York and Oxford: Rowman & Littlefield.

Zimmerman, Michael. (forthcoming): 'The Relevance of Risk to Wrongdoing'.

Zimmerman, Michael. (forthcoming): 'Is Moral Obligation Objective or Subjective?'.

NOTES

* Originally presented at the 2004 Bellingham Summer Philosophy Conference, Bellingham Washington, August 2, 2004, with comments by Gustaf Arrhenius and Elizabeth Harman. I am grateful to Arrhenius and Harman for their challenging criticism, as well as to Chris Heathwood, Michael Zimmerman, Owen McLeod, Elinor Mason, Eric Moore and other participants at the Bellingham Conference for comments and suggestions. In October, 2004 I discovered Mark Strasser's 'Actual Versus Probable Utilitarianism'. I see that in that paper Strasser anticipates a number of the points I make here.

¹ Brandt (1959, pp. 381-5) distinguishes between the actual utility form of utilitarianism and a form based upon "expectable utility". He acknowledges some advantages to the actual utility form, but comes down in favor of the expectable utility form in part because of an alleged "pragmatic" gain: 'our discussion will better illumine the reasoning we all have to perform in life.' (p. 382) It should be pointed out that Brandt's concept of expectable utility (defined on pp. 382-3) is a bit obscure. It is not clear that he intended it to coincide with what we now call 'expected utility'.

² When Smart (1973) presents his first statement of utilitarianism, he puts it this way: '... the only reason for performing an action A rather than an alternative action B is that doing A will make mankind (or, perhaps, all sentient beings) happier than will doing B' (p. 30). Evidently, he is thinking of a version of the theory that appeals to the actual values of consequences of actions. He immediately mentions that '... in fact we can have only probable belief about the effects of our actions, and so our reason should be more precisely stated as that doing A will produce more *probable benefit* than will doing B' (p. 30, emphasis added). A few pages later he reminds the reader that 'Unfortunately we cannot say with certainty what would be the various total situations which could result from our actions. ... All we can do is to assign various probabilities to the various possible effects of an action' (p. 38). He gives an example involving the act of giving ice cream to Smith. The choice could lead to pleasure of eating ice cream or the pain of toothache; to sympathetic enjoyment or envy on the part of a bystander, Jones. Each outcome is assigned a probability (4/5, 1/5, 2/5, 3/5, etc.). Smart then sketches the outcomes of an alternative (taking a nap; going for a walk). Smart points out that this is a simplified example. He mentions how difficult it would be to assign actual probabilities to the outcomes of actions that really arise in our moral life (p. 39). He says that the development of a good measure of objective probability is an important bit of unfinished business for the utilitarian. I mention this passage because it seems to illustrate the first pattern of reasoning – the pattern that leads to a version of the utilitarian principle that enjoins the maximization of expected utility. See note 7 below for more on Smart.

³ Jackson (1991) describes a medical case in which a doctor has the choice of various drugs, but does not know which will be best. He says that the doctor should prescribe the drug that has the highest expected utility – we would be horrified if she did anything else, even if it happened to have good actual consequences. He concludes, 'Hence, the fact that a course of action would have the best results is not in itself a guide to action, for a guide to action must in some appropriate sense be present to the agent's mind' (pp. 466-7).

⁴ Gruzalski (1981) describes an example involving a driver who chooses to pass in a dangerous spot. The driver does not know that he will make it safely, but simply as a matter of good luck, he does. Gruzalski claims that the action can be described as "lucky" since everything turned out OK, but that it was not the right thing for the driver to do, because he had no way to foresee that it would have good consequences. Gruzalski goes on to defend a form of expected utility utilitarianism.

⁵ In some variants, the claim is that the second principle should be understood as a principle about the avoidance of blameworthiness, rather than as a principle about some form of rightness. I discuss this variant later in the paper.

⁶ Timmons says that to be successful, a moral theory has to satisfy "the standard of applicability". 'According to this standard, a moral theory ... ought to specify a procedure that human beings, with their various limitations, can actually use in moral deliberation. ..., the utilitarian theory fails to do this' (2002, p. 122). The problem is that we often cannot know the total actual consequences of our alternatives. He then goes on to introduce expected utility

utilitarianism (PCU) and he says ‘Since expected utility is something we can apparently calculate (at least in many cases), PCU, unlike ACU [actual utility utilitarianism], represents a useful version of the basic utilitarian principle *that satisfies the applicability standard*’ (p. 124, emphasis added). Timmons goes on to suggest that the utilitarian can keep ACU as the criterion of objective rightness, but use PCU as the criterion of subjective rightness.

⁷ See pp. 46-7 of *Utilitarianism: For and Against*. In this passage Smart distinguishes between two kinds of evaluation. He says that the utilitarian can say that an act is *right* if it ‘in does *in fact* produce the best results’. On the other hand, we can say that an act is *rational* if it is the action which is, on the evidence available to the agent, *likely* to produce the best results...’ He immediately corrects this by saying, ‘For a more accurate formulation we should have to weight the goodness of the results with their probabilities’ (p. 47).

⁸ See, for example, Michael Zimmerman, ‘The Relevance of Risk to Wrongdoing’, currently unpublished; Elinor Mason 2003; Graham Oddie and Peter Menzies 1992. Each of these philosophers makes use of some slight variant of the pattern of reasoning described. The views of Oddie and Menzies are more complicated. I discuss them in greater detail in an appendix.

⁹ I should mention here at the outset that I am not convinced that the objection has any merit in the first place. I bypass that question here, focusing instead on the question whether the move to expected utility provides an answer. My own view is that R. Eugene Bales (1971) gave a satisfactory response to the objection a long time ago.

¹⁰ Jackson (1991), p. 463.

¹¹ Some commentators have suggested that when you don’t know what would actually be best, you should do the thing that, given your evidence, seems to you to be best. But Jackson’s example shows that this approach is mistaken. Consider act a2. Jill’s currently available evidence *guarantees* that it does not have the best effects. According to that evidence, either a1 or a3 is best; a2 is guaranteed to be second-best. So the available evidence does not support the belief that a2 has the best results. Thus, if we are impressed by the epistemic objection, and we want to reformulate consequentialism so as to make it yield the desired results, we cannot revise it to say that we should do what maximizes *likely value*, where the likely value of an outcome is the value such that your evidence supports the view that it will have that value.

¹² This account of expected utility is incomplete. I have not said anything about the constraints on the assignments of probabilities. I have not said anything about the relevant conception of probability, either. I discuss these questions in slightly greater detail below in the Appendix.

¹³ I write here as if there is only one way to determine which alternative maximizes expected utility. Of course, that is a simplification. There are many other ways. For example, you might have access to a friend who is capable of performing the calculations. You could gain the relevant information from his testimony. In ordinary circumstances, however, no such source of information is readily available.

¹⁴ Talk of “writing down an alternative” is a bit messy. I assume that the decision procedure would require the writing down of a brief description of the proposed alternative. Perhaps it would be sufficient for Dr. Jill to write down ‘prescribe 325 milligrams of aspirin to be taken twice daily.’

¹⁵ Again there is a problem with the notion of “writing down an outcome”. I take this to mean that anyone following the decision procedure would have to write down a brief but identifying description of the outcome. Perhaps it would be sufficient to write something like ‘Jack is cured; he lives to age 90 in good health; and then dies peacefully in his sleep. In every other respect things turn out as they do in the actual world.’

¹⁶ In fact, it has been proposed. Michael Zimmerman defines the “expectable value” of an alternative as the sum of the products, for all outcomes of that alternative, of the outcome’s probability and the amount of value such that the agent’s evidence supports the conclusion that the outcome would have that much value. See Zimmerman “Is Moral Obligation Objective or Subjective?”, pp. 15-16.

¹⁷ In his commentary on my paper in Bellingham, Gustaf Arrhenius suggested that the relevant value is the “ought to believe” value. He understood this, roughly, as the value that the agent would assign to the outcome if she were to be believing all the things that she ought to be believing – that is, all the things she would be believing if she were an epistemically responsible believer. It seems to me that this makes Dr. Jill’s project even more daunting. For now it is not enough for her to know the value that she in fact assigns to On; now she has to know the value that she would assign to On if she were to be believing all the things that she, if she were a responsible believer, would be believing. This, perhaps, is a specific number. But the notion that a real-life, epistemically ordinary person could ever determine such a number in any case is mind-boggling. There is no number, *n*, such that Dr. Jill could determine that if she believed everything she ought to believe, she would believe that the value of On = *n*.

¹⁸ There are some constraints on the assignments of probabilities. One crucial constraint is that the probabilities must be assigned in such a way that the sum of the probabilities of all the outcomes of a given alternative = 1. She cannot just assign at random.

¹⁹ In her comments on my paper at the Bellingham Conference, Elizabeth Harman said, ‘...we might know which option has greatest expected utility without actually performing a calculation: we might be good at intuiting which option has greatest expected utility given our beliefs about the values of outcomes and their likelihoods. In general, one can come to know that p without carefully deriving p from all the particular facts in virtue of which p is true.’ Although it is not entirely clear to me that this is precisely what Harman meant to suggest, the remarks seem to indicate a decision procedure like this: when in a moral quandary, we could first list all the (purely subjective) values and (purely subjective) probabilities of our alternatives, and then we could contemplate the list of numbers and simply intuit, with respect to one of the alternatives, that it maximizes expected utility. Then we could perform the action that maximizes expected utility. In the quoted passage, Harman suggests that we could do this without having to go through the laborious process of calculating the expected utility of all the alternatives. This seems to me to presuppose a seriously implausible assumption about our powers of intuition. Perhaps some mathematical savants have the cited power (I think here of the character Raymond played by Dustin Hoffman in the movie *Rain Man*). I am certain I don’t have any such power.

²⁰ Elinor Mason makes a similar suggestion in “Consequentialism and the ‘ought implies can’ principle”. Michael Zimmerman has indicated (in private correspondence) that he would endorse this sort of approach, too.

²¹ Someone might prefer to formulate utilitarianism in terms of expected utility for a different reason. He might reject the idea that for every possible action, there is some outcome that would occur if the action were performed. He might think instead that there is real indeterminacy in nature, and that the only facts of the matter are that various outcomes are possible to various degrees. He might claim, then, that ACU is based on a faulty metaphysical picture. For this reason he might advocate the move to PCU. I take this to be a completely different motivation for moving to expected utility. I do not discuss it in the paper. I discuss only the lines of reasoning displayed at the outset – lines of reasoning based on the claim that though each alternative in fact has an actual utility, often we don’t know these utilities.

²² Earlier I mentioned that it is sometimes possible to know which alternative maximizes expected utility even when one does not know the expected utility of any alternative. This could happen if a sufficiently trustworthy sage testified in favor of one of the alternatives. But this does not offer the advocate of PCU much of an advantage over the advocate of ACU, since precisely the same possibilities may arise concerning actual utility. A wise and trusted sage might turn up. He might say, “Trust me. Choose alternative a3. It maximizes actual utility.” If the sage were known to be sufficiently wise, someone might in this way come to know what maximizes actual utility by his testimony.

²³ Gruzalski claims that one important advantage of foreseeable consequence utilitarianism is that it reinforces the link between right action and moral praise and blame. He imagines a case in which someone performs an act that maximizes expected utility but turns out to have low actual utility. He says, ‘... the agent is not to be blamed or to feel guilty for so acting even if the consequences which occur are other than those which one would have expected’ (Gruzalski 1981, p. 174).

²⁴ A more restricted principle would say that if a person chooses a certain alternative *knowing* that it maximizes expected utility, then he cannot be blamed for making that choice. (Note that this principle suggests only a *sufficient condition* for blamelessness.) It might be harder to provide a convincing counterexample to this principle. However, it seems clear that this does not affect the main point I am trying to make in this paper. The problem, obviously enough, is that in real life there are no interesting cases in which an ordinary human being knows, with respect to one of his alternatives, that it maximizes expected utility. The proposed principle has no application to real world practical moral problems.

²⁵ Although I do not pretend to understand Gibbard’s view entirely, there are passages in which he seems to be saying that in order to remain above legitimate reproach, one should *always* do one’s subjective obligation; and there are passages in which he seems to be saying that one’s subjective obligation is *always* to do the act that maximizes expected utility. I have indicated my reasons for rejecting this pair of claims. I am not sure that they are claims that Gibbard actually meant to endorse.

²⁶ I think there is a typo in the paper here. They left out the left parenthesis, which I have inserted.

²⁷ My impression is that there should be infinitely many such K_i ’s – one for each amount of probability.