

Consequentialism and Indeterminacy

Abstract

Many actions have indeterminate consequences. There is no fact of the matter about precisely what would have happened if they had not been taken. How should act consequentialists think about their objective moral status? This paper suggests a way for act consequentialists to think, and draws attention to three remaining problems for act consequentialists to solve.

Keywords: Conditionals, Conditional Indeterminacy, Consequentialism, Outcomes.

1. Outcomes

Act consequentialists say that the objective moral status of any action is determined by whether its outcome is relevantly better or worse than the outcomes of alternatives to it.

That's well and good for a rough sketch of a moral theory, but there are many details to fill in. What are *actions*? What is *objective moral status*? For any given action, what is its *outcome*? What is it for one outcome to be *relevantly better or worse* than another? For any given action, what are the *alternatives to it*? How, exactly, is the objective moral status of an action *determined* by these things?

I will focus on one little bit of the detailed picture here. For any given action, what is its *outcome*?

Back in the day some people working in the tradition then known as 'utilitarian', now known as 'consequentialist', may have thought of the outcome of an action as including all

and only things caused by the action.¹ So yesterday afternoon, the afternoon of July 22 2019, I scratched my nose. Subsequent to that, many things happened. I felt a pulse of sweet relief. Boris Johnson became Prime Minister of the United Kingdom. The pulse of relief was part of the outcome of my action, because it was caused by my action. But Johnson's ascension was not, because it wasn't caused by my action.

Nowadays consequentialists tend to think of the outcome of an action as including all and only the things that would happen if the action were taken. On this way of thinking the outcome of my scratching my nose includes my pulse of relief and Johnson's ascension, along with everything that has happened, is happening or will happen: the Big Bang, the evolution of the opposable thumb, the French Revolution, the evaporation of the waters of the Earth, Mercury's being swallowed up by the Sun... *everything*.

There are at least three reasons to prefer the modern approach. First, from the point of view of theoretical economy, it enables us to avoid an unwieldy sort of double comparison. On the old-fashioned approach, to establish what the outcome of an action is I must establish what it caused to happen, which effectively requires me to compare the action to its alternatives (with respect to what would happen if they were taken). Then, to establish what the moral status of an action is, I must compare the action to its alternatives a second time (with respect to whether it or its alternatives has a relevantly better outcome). On the modern approach there's just one comparison to be made: outcome to outcomes.

Second, the modern approach enables act consequentialism to accommodate more complex evaluative judgments without stretching the notion of causal consequence beyond

¹ This is what Eric Carlson (in Carlson 1995) calls the 'Principle of Causal Outcomes.' Carlson attributes the principle to GE Moore in Moore (1912), though Carlson acknowledges that the strength of Moore's commitment to the principle is somewhat unclear.

its natural limit. An act consequentialist may want to say, for example, that sometimes whether you ought to punish someone depends on whether they committed a crime. Other things being equal, it is bad to punish someone for a crime they did not commit, good to punish someone for a crime they committed. So sometimes events uncaused by an action may have a bearing on its moral status. To accommodate this on the old fashioned approach we must say that one of the things you cause, when you punish the innocent person, is *somebody being punished for a crime they did not commit*. Facts about events you do not cause (in this case *the person not having committed the crime*) are so-to-speak baked in to events that you cause. But this is inelegant.

Third, the modern approach more cleanly captures the motivating idea behind act consequentialism. What fundamentally matters is *all that ever happens*. So the important question with respect to an action is this: ‘what is all that would ever happen if this action were taken?’

Well and good, but there’s a problem. What if, sometimes, for some actions, there is no fact of the matter about precisely what would happen if they were taken? What if *conditional indeterminacy* is a real thing?

2. Conditional Indeterminacy

Sometimes, when we ask what would happen (or would have happened) if a certain condition obtains (or has obtained), we get a curious pattern of truths. A proposition of this form is true:

If A then it would be that B or C

But neither proposition of these forms is true:

If A then it would be that B

If A then it would be that C

One situation in which this pattern can arise involves indeterministic laws. Consider

The Truncated Quantum Mechanics Lecture

I am teaching a class on quantum mechanics. As a final demonstration to the class, I plan to fire a single photon through a narrow slit and measure its trajectory. The laws of quantum mechanics determine that it will angle up or angle down, but they do not determine which. Sadly I run out of time. I never perform my demonstration.

Now this is true:

If I had fired the photon, then it would have angled up or angled down.

But neither of these is true:

If I had fired the photon then it would have angled up.

If I had fired the photon then it would have angled down.

Why? Well, the short version of the answer is this: Modal propositions (those that have to do with what might be the case, could be the case, or would be the case) are made true or false by non-modal propositions (those that have to do with what is, has or will be the case). But there are no non-modal propositions that make one, but not the other, of the above pair true. It is consistent with the way our world is, and the laws that govern our world, that the photon would have angled up, and it is consistent with the way our world is, and the laws that govern our world, that the photon would have angled down.

The longer version of the answer is this: On the now standard approach to evaluating the truth of conditionals,² a proposition of the form

² Canonically sourced to David Lewis and Robert Stalnaker in Lewis (1973) and Stalnaker (1984).

If it had been that A, then it would have been that B.

is true just in case some possible world (where possible worlds represent fully determinate ways for the world to be, non-modally speaking) in which A and B are both true is closer (where distance is a measure of relevant similarity) to the actual world than any possible worlds in which A is true but B is not. Applying the standard approach to the three conditionals above, some world in which I fire the photon and it either angles up or angles down is closer to the actual world than all worlds in which I fire the photon and it does not angle up or down, so the first conditional comes out true. But no world in which I fire the photon and it angles up is closer to the actual worlds than all worlds in which I fire the photon and it angles down, and vice-versa, so neither of the second and third conditionals comes out true.³

In situations like this I will say that there is *no fact of the matter* about whether the situation described by the consequent (e.g. the photon angling up) would obtain. And I will refer to the general phenomenon as *conditional indeterminacy*.

Conditional indeterminacy need not arise from indeterminacy in physical laws.

Consider

The Balance of Justice

You stand before a scale-of-justice-style balancing scale, with two equal weights nicely balanced on either side.

Now this is true

³ I should note that on Lewis' version of the approach the second and third conditionals are false, while on Stalnaker's version of the approach they each have indeterminate truth value. I prefer Lewis' version, for reasons I explain in Hare (2011) and Hare (manuscript), but this won't matter as we move on.

If one of the weights were substantially heavier than the other, then the scale would tip left or tip right.

but neither of these is true

If one of the weights were substantially heavier than the other, then the scale would tip left.

If one of the weights were substantially heavier than the other, then the scale would tip right.

No non-modal propositions make one, but not the other, of the second and third modal propositions true. No possible world in which the left weight is substantially heavier, and the scale tips left, is relevantly more similar to the actual world than all worlds in which the right weight is substantially heavier, and the scale tips right. This is not because of indeterministic physical laws. It is because the common antecedent of the second and third conditionals, 'If one of the weights were substantially heavier than the other...', is underspecified in relation to their consequents. If the left weight were heavier then which way would the scale tip? There's a good answer to that question – it would tip left. If the right weight were heavier then which way would the scale tip? There's a good answer to that question – it would tip right. If one of the weights were heavier then which way would the scale tip? There's no good answer to that question.

Can conditional indeterminacy due to conditional underspecification arise when the antecedent of the conditional describes a human action? It would appear so. Consider

The Sensitive, Accurate Kitchen Scale

There's a kitchen-style electronic scale before me. It gives an accurate reading in newtons, to 4 decimal places. I mull over whether to press my thumb down on the scale, but decide not to. I leave the scale entirely alone.

Now this is true

If I had pressed my thumb down on the scale, then the final digit of its maximum reading would have been even or odd.

but neither of these is true

If I had pressed my thumb down on the scale, then the final digit of its maximum reading would have been even.

If I had pressed my thumb down on the scale, then the final digit of its maximum reading would have been odd.

There are many nearby possible worlds in which I press down on the scale and get different maximum readings. In some the final digit is even, in others odd. But no world in which I press and get an even final digit is relevantly more similar to the actual world than all worlds in which I press and get an odd final digit, and vice versa. The antecedent of the pair of conditionals above, 'If I had pressed my thumb down on the scale...', is underspecified relative to their consequents. If I had pressed down with exactly 3.4068 N of force, would the final digit have been even or odd? There's a good answer to that question. It would have been even. If I had pressed down with exactly 3.4067 N of force, would the final digit have been even or odd? There's a good answer to that question. It would have been odd. If I had pressed down would the final digit have been even or odd? There's no good answer to that question. Whether the final digit would have been even or odd depends on precisely how I

would have pressed down, if I had pressed down. But there is no fact of the matter about precisely how I would have pressed down, if I had pressed down.

How widespread is this phenomenon? Two further examples should illustrate that it is very widespread indeed. Consider:

The Unflipped Coin

You balance a coin on your thumb, think about whether to flip it, decide against doing so, and return it to your pocket.

Now this is true

If you had flipped the coin, then it would have landed heads or landed tails.

But neither of these is true

If you had flipped the coin, then it would have landed heads.

If you had flipped the coin, then it would have landed tails.

There are many nearby possible worlds in which you flip the coin at many different heights, many different angular velocities, many different rates of rotation. In some you get heads, in others you get tails. But no world in which you flip the coin and get a head is relevantly more similar to the actual world than all worlds in which you flip the coin and get a tail, and vice-versa. The antecedent of the conditionals above, 'If you had flipped the coin...', is underspecified relative to their consequents. Whether or not you would have gotten a head or tail depends on precisely how you would have flipped the coin, if you had flipped the coin. But there is no fact of the matter about precisely how you would have flipped the coin, if you had flipped the coin.

Finally consider

The Unwaved Hand

Shall I wave at the marching band as they process by me? I don't do it. And exactly 100 days later it rains in Cambridge MA.

Now this is true.

If I had waved my hand then 100 days later it would have rained or not-rained in Cambridge MA.

but chaotic behavior in our best models of the atmosphere suggests that neither of these is true⁴

If I had waved my hand then 100 days later it would have rained in Cambridge MA.

If I had waved my hand then 100 days later it would not have rained in Cambridge MA.

The antecedent of these conditionals, 'If I had waved my hand...' is underspecified relative to their consequents. Whether or not it would have rained in Cambridge MA depends on precisely how I would have waved my hand, if I had waved my hand. But there is no fact of the matter about precisely how I would have waved my hand, if I had waved my hand.

3. Accommodating Conditional Indeterminacy

What can an act consequentialist do about conditional indeterminacy? One approach might be to allow the indeterminacy to infect the objective moral status of actions. When there is no matter fact of the matter about would have happened if an action had been taken,

⁴ See Hare (manuscript).

then it is indeterminate what the outcome of the action is. When it is indeterminate what the outcome of the action is, it may be indeterminate whether its outcome is in the relevant way better than the outcomes of alternatives to it, hence indeterminate what its moral status is. That's just how it goes, sometimes. So, for example, in the Unflipped Coin case it is indeterminate whether the outcome of your flipping the coin is one in which the coin lands heads or one in which the coin lands tails. And supposing that it would be relevantly *wonderful* if you were to flip the coin and it land heads, and relevantly *terrible* if you were to flip the coin and it land tails, then it is indeterminate what the objective moral status of your flipping the coin is.

But this is not so satisfactory. It suggests that if chaos theorists are right about the weather then pretty much all actions have indeterminate objective moral status. For pretty much anything I might do now, in 2019, it is indeterminate whether the outcome of that action includes tornados sweeping through the city of Tulsa in September 2020. So (on the reasonable assumption that, other things being equal, outcomes in which tornados sweep through Tulsa are relevantly much worse than outcomes that don't) it is indeterminate what the objective moral status of the action is. But we moral theorists hope to say something more discriminate than that 'everything we ever do is indeterminately right, indeterminately wrong, indeterminately such that it ought to be done, indeterminately such that it ought not to be done.' If we are forced to say that then so be it, but let's see if there is something else to say.

Another approach might be to let the outcome of an action be a collection of fully specific world histories (think of these things as Lewisian possible worlds, if you are familiar with such things), each one representing a way things might go if the action were to be taken. So, given that it is determinate that the coin would travel through the air, if flipped,

but indeterminate whether it would land heads or tails, if flipped, all worlds in the outcome of flipping the coin are worlds in which the coin travels through the air, some worlds in the outcome are worlds in which it lands heads, some worlds in the outcome are worlds in which it lands tails.

On a crude implementation of this approach we leave it at that. On a more sophisticated implementation we go on to associate probabilities with classes of possibilities in the outcome. So, in the Sensitive, Accurate Kitchen Scale case, though there is no fact of the matter about whether the final digit of the scale's reading would have been even or odd, if I had pressed it, the class of worlds in which I press and the final digit comes out even has the number 0.5 associated with it – representing the fact that the conditional probability of the final digit being even, were I to press it, is 0.5.

With these conditional probabilities in place, we can use the formidable resources of expected value theory to assign values to outcomes. First we divide the outcome into classes of similarly valued possibilities. Next we assign a number to each class, representing the value of each of the possibilities in it. We weight that number by the conditional probability associated with the class. We sum. We get the value of the outcome.

Are act consequentialists home free? Not quite. In what remains of this chapter I will describe three remaining problems for the act consequentialist. I have ideas about how to solve the first problem. The second and third problems I will leave to you.

4. What if Causation Matters?

The first problem has to do with incorporating conditional indeterminacy into a ranking of possibilities. Some act consequentialists may want to say that causal facts can have a bearing on whether one possibility is relevantly better than another. Sometimes, in order to

know which outcome is relevantly better, we need to need to know more than what happens in each, we need to know what causes what to happen in each. They may want to say, for example, that it is relevantly better, other things being equal, that human success be caused by human effort. It is relevantly better that you study hard and pass your exam as a result of your study, than that you study hard and pass your exam by chance.⁵ They may want to say, for another example, that it is relevantly worse, other things being equal, that human misery be caused by malicious human action. It is relevantly worse that you be maimed by a boulder dislodged from on high by an enemy than that you be maimed by a boulder dislodged from on high by wind and hail.

Given that causation is very important to this sort of act consequentialist (call her a *sensitive-to-causes consequentialist*, or *STC consequentialist*, for short), she now owes us an account of how causation works when conditionals are indeterminate. Consider, for example:

The Stolen Coin

For reasons too tedious to itemize here (embellish the story as you like so as to make this all true) nature has conspired to put me in a terrible situation: I am heading precipitously towards death. My only hope of salvation lies with a coin in my pocket. If I flip it and get a head then I will live. Otherwise I will die. But my enemy takes the coin from me before I get a chance to flip it. I die.

The STC consequentialist wants to say that my death, relevantly bad in any case, is relevantly worse if my enemy caused it to happen by taking the coin. But did my enemy cause my death to happen by taking the coin? I died. There is no fact of the matter about whether I would

⁵ This is explicit, for example, in David Velleman's theory of the value of a life. See Velleman (1991).

have lived or died if my enemy had not taken the coin. The conditional probability of my dying if my enemy were not to have taken the coin is $\frac{1}{2}$. Does all that make for causation?

The STC consequentialist might want look to the philosophical literature on causation for help with this question.

On David Lewis' theory of causation in what he called a 'chancy world', we cause things to happen by significantly raising the chances of their happening – with what counts as 'significant' varying from context to context.⁶ In this case, after my enemy stole the coin I was certain to die, but if my enemy had not stolen the coin then I would have had a $\frac{1}{2}$ chance of dying. By stealing the coin my enemy doubled the chances of my dying. Is that significant? It is up to the sensitive-to-causes consequentialist who adopts Lewis' theory to tell us whether it is. If it is, then mine is the worse sort of death, the sort of death caused by human action.

Johan Frick has recently criticized Lewis' theory on structural grounds (e.g. sometimes, on Lewis' account, two people cause something to happen without either person causing the something to happen – Frick says that cannot be).⁷ Frick proposes instead that we cause things to happen by to-any-degree raising the chances of their happening. In this case you most definitely did to-some-degree raise the chances of my dying by stealing the coin, so you caused me to die. So an STC consequentialist who adopts Frick's theory will again say that this is the worse sort of death, the sort of death caused by human action.

But STC consequentialists should not adopt either of Lewis or Frick's theories.

Consider:

⁶ See Lewis (1986).

⁷ See Frick (manuscript).

Two Deaths, One Stolen Coin

Nature has conspired to put you and me in a terrible situation: We are heading precipitously towards death. Our only hope of salvation lies with a coin in my pocket. If I flip it and get a head then I will live and you will die. If I flip it and get a tail then you will live and I will die. Otherwise we both will die. But my enemy takes the coin from me before I get a chance to flip it. We both die.

Two Deaths, One Stolen Envelope

Again you and I are heading towards death. Our hope this time around lies with a sealed envelope. One of our names is written inside. If I open the envelope then the named person will live and the unnamed person will die. Otherwise we both will die. But my enemy takes the envelope from me before I get a chance to open it. We both die. Though we never know it, as a matter of fact my name was written inside the envelope.

On Lewis' and Frick's theories, in Two Deaths, One Stolen Coin my enemy causes two deaths, while in Two Deaths, One Stolen Envelope my enemy causes just one death. So, given that other things are relevantly equal, STC consequentialists who adopt Lewis or Frick's theory must say that the outcome of Two Deaths, One Stolen Coin is worse than the outcome of Two Deaths, One Stolen Envelope. But they should not say this. They would thereby be saying that it matters whether conditionals are determinate or indeterminate. But it doesn't matter in this way whether conditionals are determinate or indeterminate.

If this isn't obvious to you, imagine yourself facing an awful choice.

Stealing an Envelope vs. Stealing a Coin

You can prevent the apocalypse in either of two ways: You can steal a coin that will otherwise save one of two unknown-you-people (in the manner above). You can steal an envelope that will otherwise save one of two different unknown-to-you people (in the manner above).

An STC consequentialist who adopts Lewis or Frick's theory must say that you ought to steal the envelope. If you steal the coin then you cause the deaths of two people. If you steal the envelope then you cause the death of one person. It is worse, other things being equal, that you cause two deaths than that you cause one death. But I submit to you that that's not right. It's fine for you to go either way here.

So what theory of chancy causation should an STC consequentialist adopt? That's our first problem.

Here's a proposal: She should say that the important sort of causal relations come in degrees. She should say that in Two Deaths, One Stolen Coin, my enemy causes my death to degree 0.5, while in Two Deaths, One Stolen Envelope, my enemy causes my death to degree 1. It is bad that people cause deaths to any degree. The badness varies by degree.

How exactly does the badness of causing-to-a-degree aggregate? Other things being equal (in particular, the number of deaths being the same) is it better or worse, for example, that three deaths be caused to degree 0.25 than that one death be caused to degree 0.8? Questions about value aggregation are notoriously difficult to answer, but we have a guide to how to answer this one. If the outcome of Two Deaths, One Stolen Coin is exactly as bad as the outcome of Two Deaths, One Stolen Envelope, then it must be exactly as bad that two deaths be caused to degree 0.5 as that one death be caused to degree 1. More generally, if it

does not matter whether conditionals are determinate or indeterminate, then it must be worse, other things being equal, that m deaths be caused to degree n than that i deaths be caused to degree k , just in case mn is greater than ik .

5. What Makes Some Worlds Closer Than Others?

The second problem has to do with relevant closeness. Sometimes it is not obvious how to assess conditionals. Consider, for example:

If kangaroos had no tails then they would topple over.⁸

If kangaroos had no tails then they would not topple over

On one way of thinking the first is true. If all the kangaroos in the world were to lose their tails, through some terrible, coordinated mishap, then they would become unbalanced, and topple forward. On another way of thinking the second is true. If kangaroos had not evolved powerful, weighty tails, then they would have evolved different forelegs, and remained nicely balanced. On yet another way of thinking neither is true. This is a case of conditional indeterminacy due to conditional underspecification.

Which way of thinking is right? On the standard theory that depends on whether worlds in which kangaroos lose their tails and topple forward (call these the *topple-worlds*) are more relevantly similar to the actual world than worlds in which kangaroos evolve no tails and remain nicely balanced (call these the *no-topple worlds*). That depends on what ‘relevant similarity’ is. And that changes from conversational context to conversational context. When you are amongst mechanical engineers, talking about mechanical engineering, you are right to say that tail-less kangaroos would topple over. In that context the topple-worlds are

⁸ This is the first example of a conditional that Lewis gives in Lewis (1973). He does not use it to illustrate the phenomenon of underspecification, but it serves that purpose well.

relevantly more similar to the actual worlds than the no-topple worlds. When you are amongst evolutionary biologists, talking about evolutionary biology, you are right to say that tail-less kangaroos would be nicely balanced. In that context the no-topple worlds are relevantly more similar to the actual world than the topple worlds. When you are with your friends, talking about nothing in particular, you are wrong to say that tail-less kangaroos would topple over, wrong to say that tail-less kangaroos would be nicely balanced. In that context neither topple nor no-topple worlds are relevantly more similar to the actual world.

Consider, for another example, a pair of conditionals having to do with the men's long jump at the Rio Olympics (actually won by Jeff Henderson, with a jump of 8.38 meters):

If I had entered the men's long jump at the Rio Olympics and jumped between 8.20 and 8.56 meters, then I would have won gold.

If I had entered the men's long jump at the Rio Olympics and jumped between 8.20 and 8.56 meters, then I would not have won gold.

On one way of thinking neither conditional is true. Their antecedents are underspecified. On another way of thinking the second is true. Wildly unrealistic as it may be, my jumping and just losing is a little closer to reality than my jumping and just winning.

Which way of thinking is right? On the standard theory that depends on whether distant worlds in which I jump between 8.20 and 8.31 meters are, other things equal, relevantly just a bit closer to the actual world than distant worlds in which I jump between 8.32 and 8.50 meters. That depends on what relevant closeness is, and that may change from conversational context to conversational context.

The general phenomenon is this: Different groups of people are considering whether or not something, call it E, would happen if something else were to happen. When talking

amongst themselves, one group says that E would happen, another group says that E would not happen, and yet another group says that there is no fact of the matter about whether E would happen. They are all right, because they are correctly applying different standards of evaluation to the counterfactuals, using different, equally legitimate measures of similarity.

In light of this phenomenon, we might want to ask: When we do moral philosophy is there one right way to evaluate conditionals whose antecedent is a human action? The act consequentialist is in a bind here. She must either say there is or say there is not one similarity relation that plays a special role in determining the moral status of an action. If there is not one special similarity relation, then she must say that actions are merely objectively right or wrong relative to similarity relations. If there is a special similarity relation, then she owes us an account of what it is. What is it? I know of no very successful efforts to address this question.

6. What If There Are No Conditional Probabilities?

The third problem has to do with conditional probabilities. The act consequentialist treatment of conditional indeterminacy that I sketched in section 3 relied on the fact that sometimes, when events might have occurred, and there's no fact of the matter about whether they would have occurred, we can associate conditional probabilities with their occurring. How exactly does this work? The idea is that, when no world in which A and B is relevantly closer than all worlds in which A and not-B, and vice versa, we may nonetheless be able to say that a certain percentage of suitably nearby A-worlds are B worlds. Then we can say that the conditional probability of B, if it were that A, is that percentage. For example, in the Sensitive, Accurate Kitchen Scale case, though no world in which I press and get an even final digit is closer than all worlds in which I press and get an odd final digit, we

can say that 50% of suitably nearby worlds in which I press are worlds in which I get an even final digit. So the conditional probability of my getting an even final digit, if I were to press, is 0.5.

But there are infinitely many suitably nearby worlds in which I press! How can it be that 50% of them are worlds in which I press and get an even final digit? What's 50% of infinity? The idea, roughly, is that though the set of suitably nearby worlds is infinite, it can be partitioned into finitely many cells. Some finite partitions are more *natural* than others – for example the partition that divides the set into two cells, one containing those worlds in which I press down with more than 4 Newtons of force, and the other containing those worlds in which I press down with less than 4 Newtons of force, is more natural than the partition that divides the set into two cells, one containing those worlds in which I press down with between 4 and 4.1 Newtons of force, and the other containing those worlds in which I don't. And some finite partitions are more *fine-grained* than others – for example a partition that divides the set into 100 cells is more fine-grained than a partition that divides the set into 2 cells. For any sufficiently natural, sufficiently fine-grained finite partition of the set of nearby worlds in which I press, 50% of the cells contain only worlds in which I get an even final digit. That is what we mean by saying that 50% of the cells in the infinite set are worlds in which I get an even final digit.

Well and good. But sometimes, for some instances of conditional indeterminacy, it just isn't possible to associate conditional probabilities with events occurring. Consider

The Absent Cube

No cube of gold ever enters my room.

Now this is true:

If a cube of gold, of side length between 0 and 2 meters, had materialized in my room (violating actual laws of physics), then its mass would have been above or below 19,300 kilograms.

And, due to conditional underspecification, neither of these is true:

(Above) If a cube of gold, of side-length between 0 and 2 meters, had materialized in my room, then its mass would have been above 19,300 kilograms.

(Below) If a cube of gold, of side-length between 0 and 2 meters, had materialized in my room, then its mass would have been below 19,300 kilograms.

A cubic meter of gold has mass 19,300 kilograms. So whether the cube would have had mass above or below 19,300 kilograms depends on whether it would have had side-length above or below 1 meter. But there is no fact of the matter about whether it would have had side-length above or below 1 meter, if it had had side length between 0 and 2 meters.

But this time we cannot easily associate conditional probabilities with the pair. Here are three simple finite partitions of the set of suitably nearby worlds in which their antecedent holds

P1 A partition that divides the set into worlds in which the cube has side-length between 0 and 1 meters, and worlds in which the cube has side-length between 1 and 2 meters.

P2 A partition that divides the set into worlds in which the cube has face-area between 0 and 2 square meters, and worlds in which the cube has face-area between 2 and 4 square meters.

P3 A partition that divides the set into worlds in which the cube has volume between 0 and 4 cubic meters, and worlds in which the cube has volume between 4 and 8 cubic meters.

If partitions like P1 are more natural than partitions like P2 and partitions like P3 then the conditional probability of (Above) is $1/2$. If partitions like P2 are most natural then the conditional probability of (Above) is $3/4$. If partitions like P3 are most natural then the conditional probability of (Above) is $7/8$. But really none of these partitions are more or less natural. So (Above) has no conditional probability.

Call a conditional like (Above), with no conditional probability, *deeply indeterminate*.

Can conditionals whose antecedents are actions available to an agent, and whose consequents are matters of relevant significance, be deeply indeterminate? If so, then, when applied to these actions and their alternatives, act consequentialism looks to be in deep trouble. I see no good act consequentialist account of the objective moral status of such actions. The most an act consequentialist can say is: 'If you were to do this then a bunch of important things might or might not happen, with no probability. If you were to do that then a bunch of important things might or might not happen, with no probability.' We must either bite the bullet, and accept that this class of actions has no objective moral status, or reject act consequentialism.

So the pressing problem, for an act consequentialist, is to tell us when conditionals are and are not deeply indeterminate, and to show that conditionals whose antecedents are the actions whose moral status we are interested in are not deeply indeterminate. Again, I know of no very successful efforts to address this problem.

8. Moving Forward

To summarize: Conditional indeterminacy is real. Once they have acknowledged this, act consequentialists have work to do.

You may want to say “This isn’t my work. I am not an act consequentialist. I don’t think that the objective moral status of any action is determined by whether its outcome is relevantly better or worse than the outcome of alternatives to it.” But be careful. So long as you think that consequences matter at least a bit, so long as you think that at least sometimes the moral status of an action is at least in part determined by whether its outcome is relevantly better or worse than the outcome of alternatives to it, this is your work too.

References

Carlson, E. (1995): *Consequentialism Reconsidered*, Kluwer Academic Publishers

Hare, Caspar (2011): ‘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, no. 1

Hare, Caspar (manuscript): forthcoming with Oxford University Press

Frick, Johann (manuscript): “Probabilistic Causation, Moral Responsibility, and the Problem of Aggregate Effects”

Lewis, David (1973): *Counterfactuals*, Oxford: Blackwell

Lewis, David (1986): “Postscripts to ‘Causation’” in his *Philosophical Papers: Volume II*, Oxford University Press

Moore, G.E. (1912): *Ethics*, New York: Henry Holt and Company

Stalnaker, Robert (1984): *Inquiry*, Cambridge, MA: MIT Press

Velleman, David (1991): ‘Well-Being and Time’, *Pacific Philosophical Quarterly* 72, no. 1, 48-77