

Philosophy 2²3³: Intuitions and Philosophy
Fall 2009
Tuesdays and Thursdays, 1pm - 2:15pm
Library 209

Hamilton College
Russell Marcus
Office: 210 College Hill Road, Room 201
email: rmarcus1@hamilton.edu

Class 17 - Prospect Theory and the Doctrine of Doing and Allowing

I. Deontology, consequences, acts and omissions

Moving on from the problem of free will, which bridged ethics and metaphysics, we are going to spend some time in the next couple of weeks on more ethical and practical topics, turning to some results concerning ethical theory, and practice.

Horowitz's article applies results of prospect theory, which we first saw in Shafir's essay on cognitive mechanisms, to a debate over the ethical doctrine of doing and allowing.

The doctrine of doing and allowing, also called the doctrine of acts and omissions, is essentially a deontological theory.

Deontology is opposed to consequentialism.

Both terms are labels for classes of action-guiding ethical theories.

Deontological theories, strictly speaking, are duty-based theories.

Kant's moral theory is deontological.

Deontological theories classify actions as morally permissible or morally prohibited (or, perhaps, morally beneficent) according to whether they are in accordance with one's ethical duties.

For deontologists, certain kinds of acts, like murder or lying, are naturally wrong, independently of the consequences that result from them.

In contrast, consequentialist theories, like utilitarianism, classify actions according to their outcomes.

For a consequentialist, all that matters in the ethical evaluation of one's actions are their results.

A consequentialist may perform any kind of action, as long as the benefits of that action outweigh the negative outcomes.

Oversimply, deontological theories look backwards and consequentialist theories look forwards.

The doctrine of acts and omissions says that there can be a moral difference between performing an action and failing to perform an action.

It entails that two acts which have identical consequences can have different moral worth.

Thus, acts and omissions is essentially a deontological doctrine.

In work we have not read, Warren Quinn defends the doctrine of acts and omissions on the basis of intuitive arguments.

Horowitz objects both to the doctrine of acts and omissions and to Quinn's arguments for the theory.

She adduces experimental results in psychology to explain Quinn's intuitions without appealing to the doctrine of acts and omissions.

II. Quinn's argument

Quinn first presents a thought experiment structurally similar to a standard trolley case.

The thought experiment consists of two scenarios.

In the first scenario, we must choose between saving one person and saving five others; we naturally choose to save the five people, even though the other one person will die.

In the second scenario, we are asked whether we are willing to save five people when it means actively killing a sixth.

In both cases, the best outcome, as the utilitarian calculates, would be to save the five people. But, in the first scenario, we merely let the sixth person die, while in the second scenario we must kill the sixth.

The two scenarios thus have the same outcome, but with different actions leading to those outcomes. The consequentialist, Quinn argues, has no explanation for evaluating the two scenarios differently. If we think that our actions in the two scenarios have different moral worth, then we must favor the doctrine of acts and omissions over consequentialism.

Quinn's intuitions are that our decision in the first case is clear, while the second case is more difficult. Notice that Quinn does not argue that we should not kill the one to save the five.

He merely argues that our decision in this case is, unlike that in the first scenario, unclear.

Even this weak conclusion suffices to drive a wedge between the consequentialist and the deontologist. There is nothing in the consequences of the two scenarios that could engender or account for the differential reactions.

Thus, our moral reasoning must have at least some deontological component; there must be a difference between acting and failing to act.

III. Against acts and omissions, from intuition

One standard response to intuitive arguments for the doctrine of acts and omissions is to present different thought experiments against the doctrine.

If there are compelling intuitions against the doctrine as well as in favor of the doctrine, at least some of those intuitions must be defective, on pain of consistency.

Then, since we have to cede some of our intuitions, and since our reasoning seems so clearly consequentialist in many other cases, we might as well give up the deontological reasons, and abandon acts and omissions.

Deontological theories are generally thought to be more contentious than consequentialist theories.

So, the burden of proof would be on the deontologist.

Here one such thought experiment, from James Rachels, against acts and omissions.

Imagine two scenarios, both of which concern a greedy young adult, let's call him Reed, who resents his baby cousin.

The cousin is in line to inherit a fortune.

Reed himself would inherit the fortune if his cousin were to die.

So, he prepares to drown his cousin in the bathtub.

In the first scenario, just as Reed prepares to push his cousin under the water, the child slips in the tub, falling face-down into the water.

Reed watches the child die.

In the second scenario, Reed holds the child's head under the water, as planned, until the child is dead.

Rachels argues that Reed, intuitively, is equally morally culpable in both scenarios.

All factors, other than acts and omissions, have been controlled in this thought experiment.

The two cases differ only in that Reed acts, in the second case, and omits an act, in the first.

If Reed is equally culpable in both scenarios, then there is no moral difference between acts and omissions.

We have intuitions on both sides, and so the burden of proof shifts to the more contentious deontological theory.

IV. Prospect theory and Quinn's case

Returning to Quinn's case, Horowitz takes a different approach than Rachels.

She looks at empirical data, rather than armchair results.

Quinn argued that our intuitions in his two scenarios were best explained by our preference for acts and omissions over consequentialism.

He did not deny that a consequentialist analysis could explain our intuitions in his first scenario.

He argued that consequentialist reasoning could not explain the differences in our intuitions between the two scenarios.

Horowitz argues that a sophisticated version of consequentialism, one that abandons expected utility in favor of prospect theory, could differentiate the two cases.

According to traditional consequentialism, to calculate the harm done by an action, we aggregate the consequences using standard expected-utility functions.

Our decisions whether to act or not depend on how we perceive, or calculate the results of, this aggregation.

As we know from the law of diminishing returns, though, our interests do not precisely track simple additive aggregations: the eighth cold beverage on a hot day is not nearly as valuable to us as the first.

Kahneman and Tversky show that other psychological factors undermine simple expected-utility aggregations.

Horowitz presents two examples, from Kahneman and Tversky, of differential preferences for the same outcomes which do not depend on any controversial doctrine like acts and omissions.

In the first example, subjects were given two scenarios in which they were asked to choose whether to bet or take a sure result.

Subjects were either risk-averse or risk-seeking with the same gamble, depending on whether the bet was portrayed as a possible win or a possible loss.

In the first scenario, subjects were told to imagine that they were given \$300, and that they had to choose between either a further \$100 gain, or a fifty-fifty chance of gaining either a further \$200 or nothing more.

In the second scenario, subjects were told to imagine that they were given \$500, and then had to choose between a \$100 loss or a fifty-fifty chance of losing either \$200 or nothing more.

In both cases, a risk-seeking subject will take the bet; a risk-averse subject will not.

The outcomes of the scenarios are identical.

But, people tended to be risk-averse in the first scenario and risk-seeking in the second.

Expected utility fails to explain the difference.

The second example is structurally similar, but uses disease, rather than money, to determine preferences.

In the second example, people are told about a program to fight a disease.

In one scenario, they are told about saving lives; in the other they are told about the lives that will be lost.

Again, the scenarios have precisely the same expected values, in this case expected survival rates.

Again, people tended to be risk-averse in the positive case and risk-seeking in the negative case.

In the Kahneman and Tversky cases, we need an explanation for the differential preferences, and an appeal to a doctrine of acts and omissions would be irrelevant.

Instead, Kahneman and Tversky plot their findings, and develop a function that describes people's preferences given various prospects, or possible actions.

It turns out that people tend to be consistent over cases like the two examples that Horowitz presents.

Many studies have confirmed that value functions over possible outcomes generally conform to the following pattern: they are concave for gains, convex for losses, and steeper for losses than for gains. People tend to be risk averse when it comes to gains, risk seeking when it comes to losses, and their response to losses tends to be more extreme than their response to gains (147).

Given differential responses to gains and losses, prospect theory predicts subjects' responses to such examples.

People in the different scenarios first choose different neutral points.

This first step in the decision process is called editing.

Various factors affect one's choice of a neutral point.

One factor is what the status quo is.

Other factors may include ordering and wishful thinking.

While prospect theory seems like a descriptive psychological theory, there is room in our choices of neutral points for a normative element.

One could debate which neutral points are appropriate, or correct.

Then, subjects evaluate the different alternatives from their neutral points.

If prospect theory were a correct account of the way in which we make decisions, it could be amended to be a prescriptive theory of how we should make decisions by adding criteria for correct choices of neutral starting points and for differential evaluations.

Returning to the Horowitz/Kahneman and Tversky cases, from the neutral point of gaining \$500, both options look like losses, and so people will tend to be risk-seeking.

The concavity of the negative value function means that the greater loss doesn't look nearly twice as bad as the loss of \$100, so it's worth the risk.

In contrast, from the neutral point of gaining \$300, both options look like gains, and so people will tend to be risk-avoiding.

The gain of \$200 looks like less than twice the gain of \$100, so is not worth the bet.

Here, we have a case precisely like the law of diminishing returns.

\$200 is twice as much as \$100, but it is not perceived as twice as valuable.

Similarly, 200 lives lost are twice as many as 100 lost, but it doesn't seem twice as bad.

In deciding which actions are preferable to which other actions, agents do not simply multiply the value of outcomes by the probability of those outcomes. Some probabilities are factored in at more than their face value (146).

The factor by which prospect theory differs from expected utility in evaluations is called the decision weight.

Horowitz claims that the existence of decision weights, along with some natural assumptions about how to extend Kahneman and Tversky's work to more complex scenarios, can explain our differential responses to Quinn's examples without committing us to a deontological doctrine of acts and omissions. She begins by considering a simpler case, contrasting whether to kill someone who would otherwise be safe with whether to let someone die whom you could save.

Different responses to the two cases can be explained by different choices of neutral outcomes, in consequentialist spirit, rather than by acts and omissions.

We can turn to Prospect Theory for an account of the differing intuitive responses between the two decision problems in this simple case. There is a shift in choice of neutral outcome. In deciding whether to kill the person or leave the person alone one thinks of the person's being alive as the *status quo* and chooses this as the neutral outcome. Killing the person is regarded as a negative deviation, and its value is found in a correspondingly steep part of the *v*-curve. But in deciding whether to save a person who would otherwise die, the person being dead is the *status quo* and is selected as the neutral outcome (153).

Similarly, in the more complex case, we can explain a differential response in a manner that is consistent with consequentialism, i.e. without appealing to acts and omissions, by referring to the psychological processes underlying prospect theory.

The perceived difference between killing and not killing the one person, in the second scenario, is greater than the difference between the difference, in the first scenario, between letting the one person die and saving that person.

So, we are more puzzled in the second scenario than in the first one.

V. Some questions

One question for Horowitz's analysis is whether the prospect theoretic analysis of our reasoning actually tracks her explanation.

Horowitz presents prospect theory as a way of explaining how we might reason about the two cases differently.

But, she does not present experimental results showing that subjects actually do reason this way.

If Kahneman and Tversky's results hold, though, such an argument seems more or less forthcoming, and would have substantial empirical support.

If this is right, then one's intuition that there is this difference in the force of the reasons should not be explained in terms of a perceived difference between action and inaction, but rather in terms of differing responses to gains and to losses. The advantage of this explanation is that it rests on a psychological theory that predicts fairly well in a wide variety of decision-making situations. It is not clear that a theory can be formulated, turning on perceived differences between action and inaction, that has comparable empirical authority. Certainly, none with comparable generality (154).

Still, even if we do reason about cases such as Quinn's by using the psychological mechanisms that underlie prospect theory, it is not clear that appeals to these psychological processes explain fully the different reactions we have to the two scenarios.

Does the difference in the steepness of the negative valuation curve match our hesitance to kill someone? More specifically, how can prospect theory explain that our puzzlement remains whether we are killing one person to save five others or fifty others?

If we were really reasoning consequentially, wouldn't our hesitance to kill one to save a number of others vary with that number of others?

It does not seem to fluctuate according to the curves presented in Horowitz's article.

Further, to repeat the descriptive/normative distinction: the fact that we reason according to the prospect-theoretic descriptions Horowitz presents does not entail that we should reason this way.

Horowitz recognizes this point.

In fact, she emphasizes it.

I have *not* made the claim that Prospect Theory provides a distinction among Quinn's Rescue cases *that is morally significant*. I do not see why anyone would think the distinction is morally significant... My contention is that when Quinn, or anyone else, judges that there is a difference in what it is permissible to do in the two Rescue Dilemmas, they are mistaken in thinking that they are making a moral judgment at all(155-6).