Strict Pluralism

Abstract: A deontic dilemma is a case in which an agent has multiple incompatible final obligations, such that the agent will unavoidably violate at least one final obligation. Many philosophers have defended the possibility and the actuality of deontic dilemmas. But defenders of deontic dilemmas have not offered a general theory, i.e., a system of moral principles, to predict and explain deontic dilemmas. The purpose of this paper is to consider what such a theory might look like. The paper introduces strict pluralism, a family of dilemmas-entailing theories. After introducing and motivating strict pluralism, the paper argues that strict pluralism has three principal theoretic virtues: it is more intuitively sensitive, simple, and informative than rival theories. Then the paper discusses problems for strict pluralism.

1. The trolley problem and the equivalence solution

Here's an inconsistent tetrad:

> **I1**: The agent in Switch (one on the side track; five on the main track; the agent can save the five only by throwing the switch, redirecting the trolley toward and thus killing the one) has a final (i.e., non-overridden) moral obligation to throw the switch.
>
> **I2**: The agent in Push (five on the main track; the agent can save the five only by pushing the large man from the footbridge into the way of the trolley) has a final moral obligation not to push the football player.
>
> **Difference Requirement**: I1 and I2 are both true only if there is a morally relevant difference between Switch and Push.
>
> **No Difference**: There is no morally relevant difference between Switch and Push.

Each of these claims is to some degree attractive in its own right, but we cannot coherently endorse all of them. This is a version of the trolley problem.

There are three standard ways to answer this problem. First, one can reject I1.[1] But rejecting I1 is costly. Consider the following N-victim version of Switch: one on the side track; N on the main track; the agent can save the N only by throwing the switch, redirecting the trolley toward and thus killing the one. When N is very large (say, N=one hundred thousand people) it seems that it is not just permissible but finally obligatory to throw the switch. If we combine this judgment with the negation of I1 (i.e., the view that it is *not* finally obligatory to throw the switch in the standard Switch case (where N=5)), then we have to say that there is some number X between 5 and 100,000 such that when N is less than X, it is *not* finally obligatory to throw the switch—but when N is equal to or greater than X, it *is* finally obligatory to throw the switch. Whatever number we choose for X, we'll have a problem: we'll need to explain why that number is not arbitrary. We can avoid this problem if we say that it is never finally obligatory to throw the switch no matter how large N is, but this is counterintuitive.

Second, one can reject I2. Many philosophers reject I2, and I have nothing to offer those who are comfortable with this move.[2] But this move is widely regarded as deeply counterintuitive and costly.[3]

1 One way to reject I1 is to say that it is wrong to throw the switch (i.e., finally obligatory to refrain from throwing the switch). See Judith Thomson, "Turning the Trolley," *Philosophy & Public Affairs* 36 (2008). Another way to reject I1 is to say that throwing the switch is simply permissible—neither wrong nor finally obligatory. All of the various ways of rejecting I1 run into the line-drawing problem described below.

2 For a useful version of this line, see Neil Sinhababu, "Unequal Vividness and Double Effect," *Utilitas* 25 (2013).

3 In experiments that I have conducted with [redacted] and are reported in our paper [redacted], we found that judgments in line with I2 are reported by ordinary people even after sustained reflection and co-deliberation with their peers over the course of many hours of conversation about trolley cases. These findings suggest that intuitions in favor of I2 are stable and withstand at least some significant degree of scrutiny; these intuitions are not unreflective.

Third, one can reject No Difference. To pursue this option, we need to be able to identify a morally relevant difference between Switch and Push. Unfortunately, all ways of doing this are controversial and costly.[4]

Given that all three of the usual ways of addressing the trolley problem are costly, we should consider a fourth possible solution: rejecting the Difference Requirement. Yet the Difference Requirement has almost never been seriously questioned.

To see how to go about rejecting the Difference Requirement, consider what I'll call the *equivalence solution* to the trolley problem. According to the equivalence solution, all of the following judgments are true:

I1: The agent in Switch has a final moral obligation to throw the switch.

S1: The agent in Switch has a final moral obligation not to throw the switch.

I2: The agent in Push has a final moral obligation not to push the football player.

S2: The agent in Push has a final moral obligation to push the football player.

[4] When the trolley problem was first introduced it was believed that the moral difference between the cases (if any) would have something to do with the principle of double effect (according to which harmfully using someone as a means, which occurs in Push, is morally worse than harming someone as a merely foreseen side-effect, which occurs in Switch). See Philippa Foot, "The Problem of Abortion and the Doctrine of Double Effect," *Oxford Review* 5 (1967). This is still the most widely discussed hypothesis about the difference between the cases. But explaining the moral difference between Push and Switch in terms of the principle of double effect is problematic and controversial for at least two reasons. First, there are many cases where harmfully using someone as a means and harming someone as a foreseen side-effect seem to be morally similar (see the widely discussed "loop" version of the trolley case; for discussion, see Michael Otsuka, "Double Effect, Triple Effect, and the Trolley Problem," *Utilitas* 20 (2008)). Defenders of the principle of double effect have to undertake the difficult task of explaining why that is so. Second, when stated abstractly, the principle of double effect is very counterintuitive: the idea that foreseen harm somehow becomes morally worse *by virtue of its being put to good use*, even where this makes no difference to the magnitude of the harm or any other relevant feature of the case, is intuitively puzzling. See Ralph Wedgwood, "Defending Double Effect," *Ratio* 24 (2011). Of course, the principle of double effect is not the only way of explaining the difference. For example, it is sometimes thought that a moral distinction between doing and allowing can account for the difference (but see Fiona Woollard, *Doing and Allowing Harm* (Oxford, 2015), p. 15, for grounds for skepticism on this point). I cannot here discuss all of the many possible ways of trying to identify a morally relevant difference between Switch and Push; I can only report that no one has identified a difference between these cases that is uncontroversially relevant.

The equivalence solution implies that Switch and Push are both *deontic dilemmas*—because it implies that in each case, the agent has two final moral obligations that cannot both be satisfied, and that is just what a deontic dilemma is.[5]

If the equivalence solution is true, then I1 and I2 are both true. The equivalence solution also allows No Difference to be true: there needn't be any morally relevant difference between Switch and Push because, according to the equivalence solution, Switch and Push are morally similar: they are both deontic dilemmas. The equivalence solution thus allows us to endorse I1 and I2 while at the same time enabling us to call off the search for a morally relevant difference between Switch and Push.

The cost of this solution is that we have to diagnose Switch and Push as deontic dilemmas. I'll discuss whether this cost is bearable below. What I want to emphasize right now is that the *mere* fact that the equivalence solution is costly in this way is not a decisive count against it. As we've seen, the standard solutions to the trolley problem are also costly.

The equivalence solution seems to call for some sort of ethical pluralism—but not Rossian pluralism. Rossian pluralism is committed to a view succinctly expressed in a

5 In entailing the possibility of deontic dilemmas, is the equivalence solution committed to logical incoherence? No. A deontic dilemma is a case where O(X)&O(Y)&~C(X&Y) (the agent is finally obligated to do X, and finally obligated to do Y, but the agent cannot do both). This is not an overt contradiction. To get to a contradiction, we need additional principles. To illustrate, consider the most widely discussed path from deontic dilemma to contradiction. This path relies on two principles: Agglomeration: If O(X)&O(Y) then O(X&Y); and OIC ('ought implies can'): If O(X&Y) then C(X&Y). It can be seen that these principles, together with the assumption of a deontic dilemma, can produce a contradiction. But both of these principles can be coherently rejected and thus contradiction is avoided. Lisa Tessman is a very influential defender of deontic dilemmas who rejects OIC (see her *Moral Failure* (Oxford, 2014)). By contrast, my own preference is to uphold OIC and to reject Agglomeration. The important point for present purposes is that strict pluralists and other defenders of deontic dilemmas have multiple ways of rendering their views logically coherent. For discussion of these points see John Holbo, "Moral Dilemmas and the Logic of Obligation," *American Philosophical Quarterly* 39 (2002). For a valuable overview of ways to have deontic dilemmas without contradiction, see Lou Goble, "A Logic for Deontic Dilemmas," *Journal of Applied Logic* 3 (2005).

church placard observed by R.M. Hare: "If you have conflicting duties, one of them isn't your duty."[6] According to Rossian pluralism, whenever there is a conflict between obligations, one of them overrides the other (except in the event of a tie, in which case the Rossian says the agent has two permissible options).[7] And an overridden obligation is not a *final* obligation, by definition. Thus, there are no genuine deontic dilemmas in the Rossian picture, and so there can be no equivalence solution.

To see what a non-Rossian pluralistic theory explaining the equivalence solution might look like, consider the following toy theory. Suppose there are two general moral principles, as follows:

> **Schematic utilitarianism**: In any given situation, we always have a final moral obligation to do whatever will maximally serve everyone's interests.

> **Schematic pacifism**: In any given situation, we always have a final moral obligation to refrain from inflicting grave undeserved harm on others.

If these two principles were true, they would explain the truth of the equivalence solution: schematic utilitarianism would explain I1 and S2; schematic pacifism would explain I2 and S1.

Let's use SP-UP as our name for the theory that contains just these two principles. SP-UP is a member of a general family of theories that I want to call *strict pluralism*.

6 R.M Hare, *Moral Thinking: Its Levels, Method, and Point* (Oxford, 1981) p. 26.
7 W.D. Ross, *The Right and the Good* (Oxford, 1930), ch 2.

Theories in this family posit multiple conflicting principles; in that respect, these theories resemble Rossian pluralism. The difference is that the strict pluralist holds that these multiple conflicting principles yield final obligations, not mere pro tanto obligations. There are many ways to be a strict pluralist; SP-UP is an illustrative example of such a theory.

The foregoing considerations provide the following initial rationale for strict pluralism. There are three traditional ways to solve the trolley problem: reject I1, or reject I2, or reject No Difference. But all of those options are costly. In hope of avoiding these costs, we should consider whether we can reject the Difference Requirement. The equivalence solution gives us a way to do this. This puts us in the market for a way to uphold the equivalence solution. Strict pluralism provides a straightforward theoretical basis for the equivalence solution. So, we ought to consider whether some form of strict pluralism—perhaps SP-UP, or perhaps strict pluralism in some other form—can be defensible. These considerations obviously do not establish the truth of strict pluralism. But they show that strict pluralism deserves our attention and that investigating the advantages and disadvantages of strict pluralism may be worth our time.

Strict pluralism is almost completely uncharted territory. As far as I know, the only precedent in the literature for this type of view is provided in a decades-old paper by Gerald Paske.[8] Paske's argument for a view in the neighborhood of strict pluralism appeals to the seemingly intractable theoretical disaagreement between consequentialists and Kantians. His argument is not (on its own) rationally persuasive, as his critics have

8 Gerald Paske, "Genuine Moral Dilemmas and the Containment of Incoherence," *Journal of Value Inquiry* 24 (1990). Although Paske's view is closer to strict pluralism than any other view of which I'm aware, the core idea that conflicts between different types of value are in some way irresolvable is not unusual; see, e.g., Thomas Nagel, *Mortal Questions* (Cambridge, 1979), especially chapter 9, for a prominent example.

argued.[9] But this does not mean that strict pluralism or some theory in the neighborhood of strict pluralism *cannot* be defended persuasively.

Strict pluralism represents a non-traditional way to defend and explain the existence of deontic dilemmas. A traditional line of argument for deontic dilemmas appeals to certain cases, such as Sophie's Choice,[10] that intuitively *seem* (to some people) to be moral dilemmas. The argument is that because these cases seem like moral dilemmas, they must really be moral dilemmas. Lisa Tessman, for instance, provides nuanced and deeply insightful discussions of fictional and real-life cases from the Holocaust in the course of an argument along these lines.[11] By contrast, the strict pluralist case for deontic dilemmas does not appeal to intuitions about *those* sorts of cases. Rather, it appeals to intuitions that are straightforwardly about final obligations; deontic dilemmas emerge as a strategy for vindicating those intuitions, not as part of the content of the intuitions themselves.

Strict pluralism has a lot of serious disadvantages. Its prospects will depend heavily on whether its disadvantages can be overcome. I'll discuss some of those disadvantages in due course. But strict pluralism also has at least three important advantages, and I want to be very clear about those advantages before I move on to the disadvantages.

9 See Daniel Statman, "A New Argument for Genuine Moral Dilemmas?" *The Journal of Value Inquiry* 26 (1992), and Terrance McConnell, "Dilemmas and Incommesurateness," *The Journal of Value Inquiry* 27 (1993).

10 A mother has to decide which of her children will be killed or else both will be killed. The case comes from the William Styron novel entitled *Sophie's Choice*.

11 Tessman, *Moral Failure*, chapter 4.

2. Intuitive sensitivity

Let's say that a theory T is maximally *intuitively sensitive* iff: Whenever a judgment J seems to be true, then T implies that J is true. And a theory T is maximally *intuitively specific* iff: Whenever a judgment J *does not* seem to be true, then T implies that J is *not* true.

Intuitive sensitivity and intuitive specificity should be understood as matters of degree. A theory possesses a high degree of intuitive sensitivity iff: When a judgment J seems to be true, then T *usually* implies that J is true. A theory possesses a high degree of intuitive specificity iff: When a judgment J *does not* seem to be true, then T *usually* implies that J is *not* true.

If a given theory possesses a high degree of intuitive sensitivity, is this an advantage (all else equal) for that theory? Yes. Whether we accept any of the epistemological views that go by the name of intuitionism, philosophers almost uniformly accept a kind of *methodological* intuitionism. That is, we theorize as if we believe that our moral intuitions are reliable, and we show a strong preference for theories that can account for and confirm the bulk of our intuitions (especially those intuitions that are stable and widely shared). And intuitive sensitivity is, quite straightforwardly, a measure of the degree to which a theory is capable of providing such an account. So, most of us will want a moral theory that is intuitively sensitive.

What about intuitive specificity? In order for a theory to be intuitively specific, it must not *go much beyond* our intuitions—a maximally intuitively specific theory does not tell us anything that our intuitions do not already tell us. This is not an obviously desirable feature in a theory. On the contrary, it seems to me, we should want our moral

theory to expand our moral knowledge beyond what is available via intuition alone. Otherwise, what would be the point of devising a moral theory in the first place?

So, intuitive *sensitivity* is likely to be pretty important (at least from a broadly intuitionist perspective), whereas intuitive *specificity* should be regarded as a relatively low priority (and perhaps even sometimes a disadvantage). And that's good news for strict pluralism, because strict pluralism has a high degree of intuitive sensitivity but a low degree of intuitive specificity.

To see an illustration of this, return to our trolley cases. I1 and I2 both *seem* true and SP-UP implies that they are both true. So, it appears, at least one version of strict pluralism, namely SP-UP, is able to affirm our intuitions about these cases, thus demonstrating its intuitive sensitivity (at least with respect to this pair of cases).

At the same time, of course, SP-UP *also* implies judgments about Switch and Push that do *not* show up on our intuitive radar screen: namely, S1 and S2. It does not seem to most of us that S1 and S2 are true. In this way, SP-UP goes beyond intuition, and thus demonstrates a lack of intuitive specificity (at least with respect to this pair of cases).

SP-UP is intuitively sensitive with regard to Switch and Push, but not with regard to other cases. Consider this case:

> **Spouse-Switch**: The agent's spouse is on the main track; one stranger is on the side track; the only way for the agent to save her spouse is to throw the switch. The stranger has slightly more to live for than the spouse.

Both schematic pacifism and schematic utilitarianism imply that it is wrong to throw the switch. But many people have the intuition that the agent has a final moral obligation to throw the switch (to save the spouse rather than the stranger). This intuition about Spouse-Switch lowers SP-UP's intuitive sensitivity.

To deal with this intuition, consider:

> **Schematic relationism**: In any given situation, we always have a final moral obligation to protect the interests of those with whom we are in special relationships.[12]

Schematic relationism implies and explains our intuition about Spouse-Switch. If we construct a new version of strict pluralism that combines schematic utilitarianism, schematic pacifism, and schematic relationism, this new theory will have a higher degree of intuitive sensitivity than SP-UP.[13] Call the resulting theory SP-UPR.

In building a version of strict pluralism, we could continue in this manner: begin with some version of strict pluralism such as SP-UP or SP-UPR; consider cases in sequence; do our best to explain our intuitions about each in terms of principles already part of our current version of strict pluralism; and when we encounter a case that defies

12 For an extremely valuable development of a relationist view, see Simon Keller, *Partiality* (Oxford, 2013).

13 It is useful to compare Spouse-Switch to

> **Spouse-Push**: The agent's spouse is on the main track. The only way for the agent to save her spouse is to push the large man into the way of the trolley. The stranger has slightly more to live for than the spouse.

If our intuition about this case is pacifist, as I expect it will be, then our divergent intuitions about Spouse-Push and Spouse-Switch could be used to provide an alternative rationale for strict pluralism along analogous lines as the rationale discussed at the outset of this paper.

such explanation, then introduce new principles (or modify the earlier principles) to deal with the case. The resulting theory is by no means guaranteed to be defensible or plausible, but it is guaranteed to have a high degree of intuitive sensitivity.

3. Elegance and informativeness

I've argued that strict pluralism, when fleshed out in the right way, can possess a high degree of intuitive sensitivity (and a low degree of intuitive specificity), and that this is a virtue of the theory. Now I want to discuss two further virtues: elegance and informativeness.

Elegance is a measure of the simplicity of an ethical theory. According to one view, the size of the shortest available synopsis of any given theory either is that theory's elegance or is a good enough heuristic for that theory's elegance. (Elegance in this context can be distinguished from parsimony—a measure of the simplicity of a theory's ontological commitments.) So, for example, traditional utilitarianism is a very elegant theory because its gist can be expressed in a single sentence: In any given situation, our sole final moral obligation is to do whatever will maximally serve everyone's interests. By contrast, it would be quite a challenge to produce an accurate one-sentence synopsis of more complex theories such as Kant's theory. Traditional hedonistic act-utilitarianism is more elegant than most of its main rivals.[14]

14 Many philosophers proceed as if elegance is valuable in ethics (see below). But it is not obvious that they are right to proceed in that way. There are two main ways to argue that elegance is a theoretic virtue in ethics. First, one can argue by analogy with science: scientific theories are better when they are more elegant; therefore, ethical theories ought to be better when they are more elegant too. A (perhaps surmountable) problem for this strategy is that there is some debate about whether elegance really is a virtue in science (Elliott Sober, *Ockham's Razors: A User's Manual* (Cambridge, 2015)). The second strategy is to argue that there are features of ethics that make elegance a virtue in ethical theorizing *regardless* of whether elegance is valuable in other areas such as science.

The *informativeness* of a theory, for my purpose here, is just the degree to which it is useful for predicting and explaining our final moral obligations. So, for instance, if a given theory can tell me how much of my income I am finally obligated to donate to charity, and can explain why it's that amount (rather than some other amount), then the theory demonstrates informativeness.

Suppose for the moment that intuitive sensitivity, elegance, and informativeness are theoretic virtues (i.e., a theory that possesses a high degree of intuitive sensitivity, or of elegance, or of informativeness, is for that reason preferable, all else equal). I'm going to argue that it's relatively easy to come by a theory that possesses an ample supply of *any two* of those virtues, but strict pluralism is very unusual—perhaps unique—among moral theories in that it can possess an abundance of all three virtues.

To make my argument, I'll want to discuss a series of six cases. We've already talked about three of the cases that I want to discuss: Switch, Push, and Spouse-Switch. Our intuitions about Switch are utilitarian (that is, they can be explained by schematic utilitarianism); our intuitions about Push are pacifist; and our intuitions about Spouse-Switch are relationist. Here are three more cases:

> **Transplant:** A doctor must decide whether to kill one healthy person (who does not want to die) in order to transplant her organs into two others, or allow the two others to die.
>
> **Fighting Nazis**: A politician must decide whether to go to war against the Nazis, which will inevitably result in the death of some innocent civilians, or allow the Nazis to run rampant, which would result in an even larger number of civilian deaths.

Education: A mother has to decide whether to send her child, who is academically brilliant, to a private school, in which case her child will get a very good education—or to a public school, in which case her child will receive a merely adequate education but will significantly improve the education of less advantaged classmates.

We tend to have pacifist intuitions about Transplant (that is, we tend to think that the main moral concern in Transplant is to refrain from doing harm to an innocent person), utilitarian intuitions about Fighting Nazis (we tend to think that the main moral concern in Fighting Nazis is to minimize overall suffering), and relationist intuitions about Education (we tend to think that the main moral concern in Education is for the mother to fulfill duties arising from her relationship to her child). So, here's a diagram to illustrate our intuitions about the above six cases:

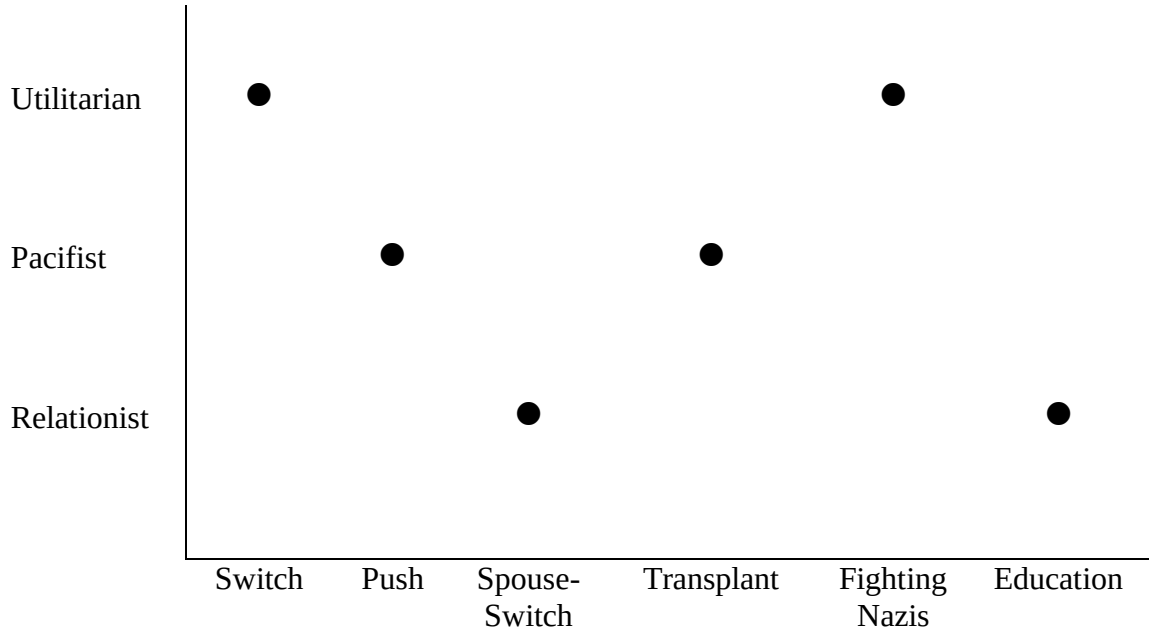

(Types of intuitions appear on the vertical axis; cases appear on the horizontal axis.)

Traditional utilitarianism captures some, but not all (and not even most) of our intuitions about these cases. Here's a diagram to illustrate the relationship of traditional utilitarianism to our intuitions:

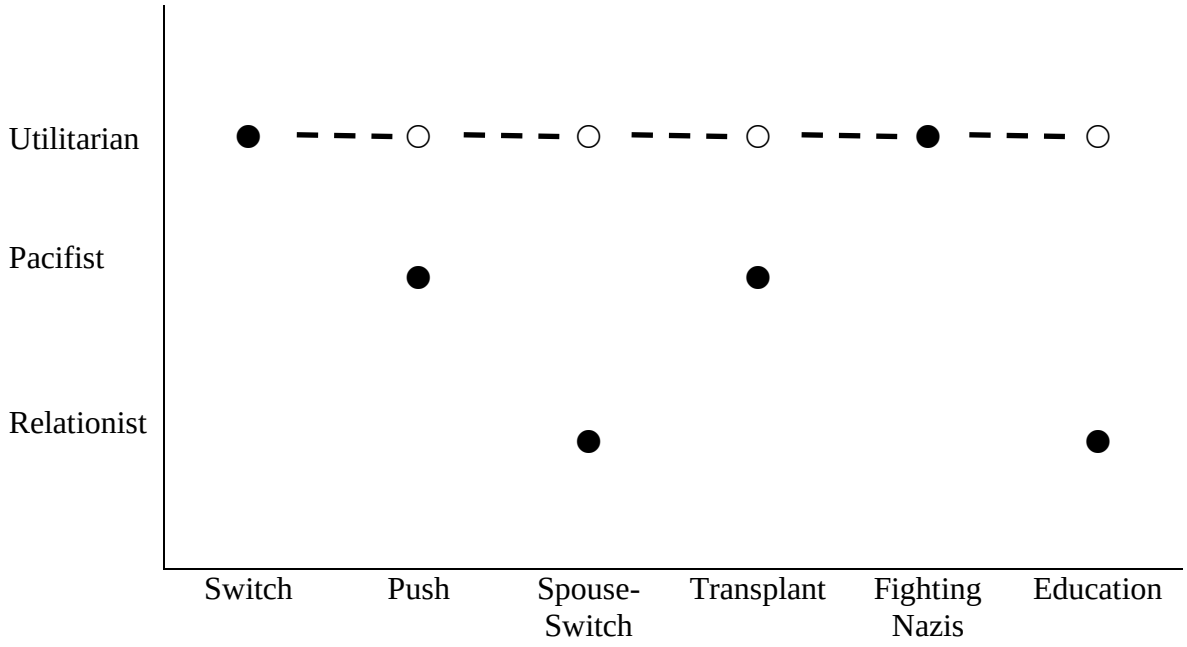

(The black points indicate intuitions; the dotted line represents the traditional utilitarian's moral theory; the open points indicate traditional utilitarianism's non-intuitive verdicts.)

Traditional utilitarianism exemplifies what I'll call the *single-class approach* in designing a moral theory. The single-class approach takes a single class of intuitions and then crafts a theory to capture intuitions of that class. Traditional utilitarianism is, of course, not the only option for the single-class theorist; an alternative would be to devise a theory that captures all and only our relationist intuitions, or one that captures all and only our pacifist intuitions. (And if we are subject to a fourth type of intuition, e.g., fairness intuitions, those could also be the basis for a single-class theory as well.)

Traditional utilitarianism is elegant, as we've already seen. And it's informative: it predicts a large number of moral obligations, and can explain each of the obligations

that it predicts. But it is not intuitively sensitive, precisely because it is a single-class theory.

If we want a theory that is more intuitively sensitive than traditional utilitarianism and other single-class theories tend to be, then we can craft a theory that will capture all or most of our intuitions, as in the below diagram. This is what I call the *curve-fitting approach.*

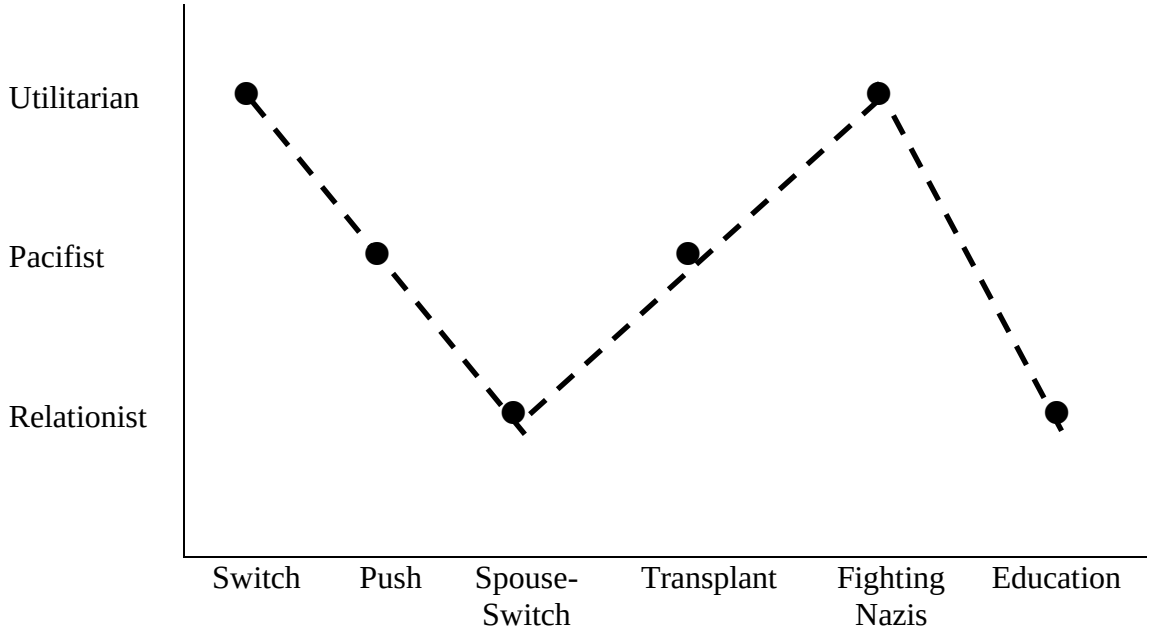

(The dotted line indicates a curve-fitter's theory.)

Of course, any theorist who wants to pursue the curve-fitting approach in a serious way would need to take account of all (or a sizeable fraction) of our intuitions, not just the above six. Imagine an expanded version of the above diagram that depicts our intuitions about many thousands and thousands of different cases. The curve-fitter's task is to produce a theory that captures all (or many) of those thousands and thousands of

intuitions. The resulting theory would be both intuitively sensitive and intuitively specific.

It seems to me that there are two main ways to pursue the curve-fitting approach. Rossianism is the first way. For the Rossian theorist, an agent's final moral obligation, in any given case, is to fulfill her weightiest prima facie duty. So, for instance, if one asks the Rossian curve-fitter why the agent in Push ought to refrain from killing the one, the answer must be: "In Push, it so happens that the agent's weightiest prima facie duty requires her to refrain from killing the one." And if one asks the Rossian curve-fitter why the agent in Switch ought to kill the one, the answer must be: "In Switch, it so happens that the agent's weightiest prima facie duty requires her to kill the one." But if we then ask *why* a given prima facie duty is weightier than other prima facie duties in a given case, the Rossian cannot give a truly informative answer. For, according to Ross, there is no general, describable principle or method—at least none discoverable by us—that can be used to compare the weights of different prima facie duties. Rather, in order to find the weighiest duty in any given case, one must simply reflect carefully about that case. So, for the Rossian, there is no substantive explanation of *why* a given prima facie duty is weightier than others; and consequently there is no substantive explanation of the content of our final moral obligations. The upshot is that the Rossian curve-fitter does not even attempt to informatively explain the shape of the line in the above diagram. And so, although the Rossian curve-fitting theory is arguably quite elegant and is intuitively sensitive, it is far less informative than a theory such as traditional utilitarianism.

The second way to pursue the curve-fitting approach is to make a serious attempt to produce a system of substantive general principles that can be used to derive final

moral obligations corresponding to our moral intuitions. Frances Kamm, for instance, is a curve-fitter of that school. Observing that Ross's theory does not provide any substantive rule by which to rank conflicting prima facie duties, she writes, "[w]e want something more helpful than a principle [such as Rossianism] that merely says, sometimes the duty to not harm takes precedence over aiding and sometimes it does not." When it comes to cases such as Switch and Transplant, she says, "[w]e are looking for a principle and its justification that explains why it is permissible to help some people by redirecting a fatal threat so that it kills one other person, and yet it would be impermissible to kill one person in order to harvest his organs to save others."[15] In other words, Kamm is looking for a genuinely *informative* curve-fitting theory.

And so Kamm sets out to develop such a theory. She ends up with a theory that is infamously complex; some of her principles take up a full page. This is not because Kamm prefers complexity over elegance. Rather, it appears, Kamm *has* to develop a very complex theory, because there just isn't any simple and informative theory that follows the path of the jagged line in the chart of our intuitions. In other words, Kamm's work makes it seem that any system of substantive, non-conflicting general principles will have to be quite complex if it is to be both informative and curve-fitting.

To sum up: If we take the single-class approach, we can produce a theory that is elegant and informative but not intuitively sensitive (e.g., traditional utilitarianism), while if we take the curve-fitting approach, we can produce a theory that is either elegant and intuitively sensitive but not very informative (e.g., Ross's theory), or one that is intuitively sensitive and informative but very complex (e.g., Kamm's theory). It may also

15 Frances Kamm, *Intricate Ethics* (Oxford, 2007), pp. 23-24.

be possible to pursue a hybrid of these different approaches. Robert Audi, for instance, is a prominent theorist whose approach seems to combine some of the advantages (and disadvantages) of both the Rossian and the Kammian approach.[16]

But is it possible to avoid these trade-offs between elegance, informativeness, and intuitive sensitivity? That is, can we produce a theory that is elegant *and* informative *and* intuitively sensitive? Well, versions of strict pluralism—such as SP-UPR—seem to have all three of those virtues. Here's SP-UPR represented in a diagram:

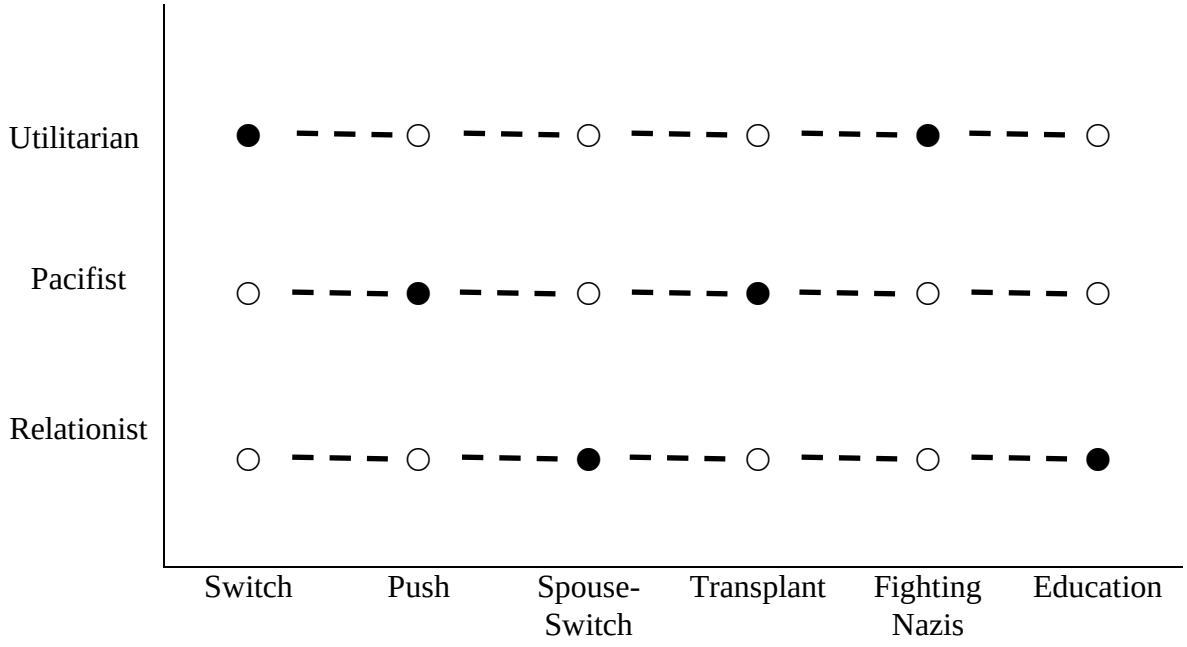

(The black points indicate intuitions; the dotted lines represent the three principles of SP-UPR; the open points indicate SP-UPR's non-intuitive verdicts.)

16 Audi (*The Good in the Right* (Oxford, 2004)) defends a series of Rossian principles, but also offers a version of the Kantian Categorical Imperative that may be used derive the Rossian principles and to guide us in negotiating conflicts between them. This allows Audi's view to be more informative than Ross's view (because Audi can provide *some* articulable explanation, via the Categorical Imperative, of why certain prima facie duties are overriding in specific cases), but less informative than Kamm's view (because the correct application of the Categorical Imperative in any given case is often highly underdetermined, unlike Kamm's rigorously elaborate principles). Similarly, Audi's view is more elegant than Kamm's view, but I think more complex than Ross's view (because Audi presents not just Rossian principles but also the Kantian Categorical Imperative).

As the above diagram illustrates (and as I have explained earlier) SP-UPR is intuitively sensitive. SP-UPR is also elegant. Perhaps it's not *quite* as elegant as traditional utilitarianism, but it's nearly so: traditional utilitarianism has just one simple principle, while SP-UPR has three simple principles. In any event, SP-UPR is far more elegant than, e.g., Kamm's curve-fitting theory. And SP-UPR is informative, since SP-UPR (unlike Rossian pluralism) offers substantive principles that can be used to predict and explain our final moral obligations. SP-UPR has these three virtues because it draws straight lines through each of three separate patterns in our moral experience, as illustrated in the diagram above.

So, at least one version of strict pluralism, namely SP-UPR, possesses a rare mix of virtues—elegance, informativeness, and intuitive sensitivity. So, *if* we want a theory that has these virtues, we ought to consider developing a version of strict pluralism. But these virtues do not come without serious costs, as I'll discuss next.

4. Costs of strict pluralism

Let's say a *shadow obligation* is a final moral obligation that is not revealed by intuition—an obligation to which we are intuitively blind. Any version of strict pluralism, whether it is SP-UP or SP-UPR or some other version, will imply that we have *a lot* of shadow obligations. That's just another way of saying that strict pluralism in any form has a very low degree of intuitive specificity. I've claimed that a low degree of intuitive specificity isn't a serious defect in a theory, but my opponents might still remain unconvinced. Strict pluralism's intuitive specificity isn't just low—it's *very* low. Even those who agree that we do not want a theory to be *maximally* intuitively specific might

still claim that we require a minimum degree of intuitive specificity, and might think that strict pluralism falls below that minimum.

This is a cost of strict pluralism, but this cost comes with some benefits. Consider the problem for intuitionism presented by order effects. A number of experiments have shown that if we consider Switch before considering Push, then we become more likely to issue a utilitarian verdict about Push, whereas if we consider Push before considering Switch, then we become more likely to issue a pacifist verdict about Switch.[17] This is disturbing to intuitionists, who want to view our intuitions as reliable, because order is a morally irrelevant factor and our intuitions cannot be reliable if they are significantly influenced by morally irrelevant factors. Thus, critics of intuitionism make much of order effects.[18]

But from a strict pluralist perspective, *the order effect does not necessarily lead us into error*. For, according to strict pluralism—more precisely, according to *some* versions of strict pluralism, such as SP-UP and SP-UPR—when we reach a pacifist verdict about Switch, we are making a true judgment; and when we reach a utilitarian verdict about Switch, we are making an equally true judgment. These verdicts are both correct; each merely reflects a different *part* of the moral truth about Switch. From the strict pluralist's perspective, such order effects merely change the direction of our attention from one class of moral truths to another. And so the order effect does not even slightly undermine the reliability of our intuitions. In this way, strict pluralism's implication that we can be

17 For example, Alex Wiegmann, et al., "Order Effects in Moral Judgment," *Philosophical Psychology* 25 (2012); Eric Schwitzgebel and Fiery Cushman, "Expertise in Moral Reasoning? Order Effects on Moral Judgment in Professional Philosophers and Non-Philosophers," *Mind & Language* 27 (2012).

18 Sinnott-Armstrong, et al., "Intention, Temporal Order, and Moral Judgments," *Mind & Language* (2008).

partially intuitively blind actually helps the moral intuitionist, by providing an intuitionism-friendly resolution of a problem presented by the phenomenon of order effects.

So, the intuitionist should see strict pluralism's shadow obligations as something of a mixed bag, rather than an unmitigated defect. I grant that strict pluralism's intuitive sensitivity may be lower than an intuitionist would consider to be ideal. But the upside is that strict pluralism's low intuitive sensitivity gives the intuitionist a way to respond to at least one vexing problem, namely the problem of order effects.

Let's now turn to another cost of strict pluralism. If strict pluralism is true, then it is not at all obvious that morality is useful for decision-making purposes. In Push, for example, versions of strict pluralism such as SP-UP and SP-UPR imply that there is a final moral obligation to push the football player and a final moral obligation to refrain from doing so. As moral advice, this seems quite unhelpful to say the least.

This is a serious cost of strict pluralism, but its magnitude mustn't overstated. To begin with, strict pluralism does provide univocal guidance in some cases. For, in some cases, all of strict pluralism's principles are in alignment. Call these *easy cases*. An easy case for SP-UPR would be a case in which the agent has a utilitarian obligation to φ, and a pacifist obligation to φ, and a relationist obligation to φ. If one is in such a case, SP-UPR is crystal clear about what to do.

It is true that strict pluralism doesn't give univocal guidance in *hard cases*—cases in which strict pluralism's principles in conflict. Switch and Push are hard cases, for instance, according to both SP-UP and SP-UPR.

However, even in hard cases, according to strict pluralism, we often can rely on moral considerations to narrow the options for us. Consider a variant of Switch in which the agent has a third option: the agent can press a button to kill everyone on the side-track *and* everyone on the main track. Such an action would violate all of the principles in SP-UP and so SP-UP can univocally rule that action out. In this way, SP-UP can inform an agent about what she univocally must *not* do, even if SP-UP also gives conflicting verdicts about what she positively *must* do.

A further cost of strict pluralism has to do with the *ubiquity* of deontic dilemmas in the strict pluralist picture. The problem isn't just that strict pluralism entails the existence of deontic dilemmas. There is excellent reason to believe that a theory can imply both the possibility and the existence of deontic dilemmas without incoherence (see note 4 above). But versions of strict pluralism that I've discussed—SP-UP and SP-UPR—entail not only that deontic dilemmas exist but that deontic dilemmas are exceedingly widespread and occur in places where philosophers least expect them. SP-UP and SP-UPR both entail, for instance, that Switch and Push are deontic dilemmas. This is likely to be contentious (and quite a bit more contentious than the view that, e.g., Sophie's Choice is a deontic dilemma). SP-UP and SP-UPR tell us that everyday life is *full* of deontic dilemmas. This may be seen as a very high cost of strict pluralism.

I want to grant that this cost is high. But how high? Most philosophers can report encounters with students in which the students sound as if they want to diagnose Switch and Push as deontic dilemmas (e.g., it is not uncommon for students to say "Both are wrong," or "Neither are right," when presented with two-option trolley cases). It's easy to dismiss these students as undisciplined, unfocused, too willing to lazily embrace

incoherence. But the strict pluralist view is that they're on to something. And there is something attractive about strict pluralism's ability to make sense of the intuitions reported by many students (and others not indoctrinated into the conventions of professional philosophy) upon first being presented with these sorts of cases.

So, I am not convinced that strict pluralism's implication that Switch and Push are deontic dilemmas is a very high cost. But I have to admit that the strict pluralist approach *does* have some serious costs in this area. To make the issue as vivid as possible, consider:

> **High Stakes:** An agent must either kill one innocent person or allow one hundred thousand innocent people to die gruesomely and painfully.

Schematic pacifism says that the agent must not kill the one in High Stakes, and schematic utilitarianism says that the agent must kill the one; so SP-UP and SP-UPR both imply that High Stakes is a deontic dilemma. If you think that High Stakes clearly isn't a genuine deontic dilemma, then you should think that SP-UP's and SP-UPR's implications for High Stakes are quite unsatisfactory.

So, SP-UP and SP-UPR have a cost: they imply that deontic dilemmas are widespread and occur unexpectedly. Whether or not this cost is high enough to destroy the prospects for defense of theories in the strict pluralist family is beyond the scope of this paper. This cost, along with other costs, has to be balanced against strict pluralism's attractions—its elegance, intuitive sensitivity, and informativeness. And we need to investigate whether these sorts of costs can be mitigated by refining strict pluralism: SP-UP and SP-UPR are, after all, only two of many theories in the strict pluralist family. My purpose in this paper has not been to show that strict pluralism is true or defensible, but

only to show that more investigation of strict pluralism's advantages and disadvantages is called for.