

Consider the Ostrich: Non-Utilitarians, Ex Ante Interests, and Burying Your Head in the Sand

RYAN DOODY

1 Introduction

Consider the following example:

Two Treatments. Twenty-six villagers in a remote area (Arden, Baldwin, Clay, ..., and Zephyr) have contracted a terrible, life-threatening disease. If nothing is done, they will all die. There are two possible treatments that could be administered. The first treatment (*Egalitrex*) ensures that all twenty-six will survive, but there's an unfortunate side-effect: paralysis of the left arm. The second treatment (*Utilicycline*), when it is effective, ensures a full recovery. You know that it will be effective for everyone but Arden, for whom (perhaps, let's say, owing to a genetic anomaly) it will be ineffective. The only way to administer the treatment is by mixing it into their water supply and, unfortunately, the two treatments have deleterious effects if taken together. Your only options, then, are the following:

Options	Results
<i>Do Nothing</i>	Arden, Baldwin, ..., and Zephyr all die.
<i>Administer Egalitrex</i> (E)	Arden, Baldwin, ..., and Zephyr all survive with paralysis in their left arms.
<i>Administer Utilicycline</i> (U)	Baldwin, Clay, ..., and Zephyr make a full recovery; Arden dies.

Many Non-Utilitarian moral views would regard it wrong to administer Utilicycline in this case. For while administering Egalitrex will make twenty-five out

of the twenty-six worse off than they would be if you administered Utilicycline instead, it does so at the cost of Arden's life. And it is not right to force a grave sacrifice on the few in order to secure comparatively smaller benefits for the many. This thought, or at least something in the neighborhood, is fairly common among Non-Utilitarians. Administering Utilicycline might generate a greater total amount of wellbeing than administering Egalitrex would, but, for many Non-Utilitarian moral views, this is not a decisive consideration. Some Non-Utilitarian views eschew the aggregation of harms and benefits across separate people entirely, attending instead to only the strongest complaints against the various alternatives [e.g., Scanlon, 1998, p. 229-30]. On other views, the sum total of wellbeing is morally relevant but must be weighed against how equally it is distributed (e.g., Pluralist Egalitarian Views).¹ And there are, of course, other views as well.²

Nothing especially turns on it being *Arden* (as opposed to Baldwin, or Clay, ..., or Zephyr) for whom Utilicycline will be ineffective. Were you instead to know that it would be ineffective for Baldwin (or for Clay, or for Doron, ..., or for Zephyr), the same verdict would hold: it is morally impermissible to administer Utilicycline when administering Egalitrex is an option.

But what if all you know is that administering Utilicycline will be ineffective for *someone*, but you don't know for whom in particular? Suppose, for example, that you know that the treatment will be ineffective for anyone who has a particular genetic anomaly, and that exactly one of the villagers has it. No one — including whoever it is who in fact has the genetic anomaly — knows who has it.

¹ Parfit [1991, 1997] distinguishes between *Telic* and *Deontic* versions of Egalitarianism. According to the former, it is in itself bad if some people are worse off than others. We should aim for equality because inequality detracts from the value of an outcome. According to the latter, we should aim for equality, but for other moral reasons — that is, for moral reasons that don't pertain to the goodness or badness of outcomes. I don't think this distinction will matter for our purposes, so long as for both sorts of views there are cases — like the one above — in which it would be wrong to administer Utilicycline. If you are an Egalitarian who thinks that, in the case above, it is *not* impermissible to administer Utilicycline because the considerations concerning equality are outweighed by those concerning the total amounts of wellbeing, feel free to adjust the example accordingly.

² For example, if you are a Consequentialist-Who-Regards-the-Value-of-People's-Lives-to-Be-Incommensurable, you might think that it is wrong to administer Utilicycline because doing so expresses the belief that your reasons in favor of administering Utilicycline outweigh your reasons against; and that belief is deeply and dangerously mistaken [e.g., Hare, 2016, p. 22]. (It's not obvious that administering Utilicycline would express such a belief in the case above — unless administering Egalitrex is, for some reason, acknowledged to be the default option (and, then, even still) — but the example can be amended to make this seem more plausible.) Also of interest: Deontological views that countenance constraints-against-killings (and constraints-against-rights-violations, more generally). In our example above, Arden's death would not be *causally necessary* for bringing about the greater good, but rather an unfortunate side-effect of doing so. Moreover, none of the villagers have a particular claim on receiving one or other of the treatments (e.g., it's *your* medicine, not *theirs*). So, you won't violate anyone's rights no matter what you do. But, again, the case can be easily amended to target these views as well.

They all regard it as equally likely that it might be any one of them. For future reference, let's call this case: **Opaque Two Treatments**. Is it morally impermissible to administer Utilicycline in this case too?

Here is a reason to think it is. If you know that Utilicycline will be ineffective for one of the villagers, then you know that, were you to administer Utilicycline, you would be sacrificing the life of someone in order to secure comparatively smaller benefits for the many. That someone might be Arden, or it might be Baldwin, ..., or it might be Zephyr. If it would be wrong to sacrifice Arden for the good of the others, and it would be wrong to sacrifice Baldwin for the good of the others, ..., and it would be wrong to sacrifice Zephyr for the good of the others, then — because you know that administering Utilicycline involves sacrificing one of their lives for the good of the others — it is morally impermissible to administer Utilicycline in this case. Call this the *Argument from Dominance*.³

On the other hand, here's a reason to think that not only is it *not* morally impermissible to administer Utilicycline in such a case, it's what you morally ought to do: doing so is in everyone's expected interest. Each of the villagers would be willing to endure a one-in-twenty-six chance of death in order to avoid paralysis in their left arm.⁴ (If this isn't so, we can adjust the case — by increasing the number of villagers (and thus reducing the risk of death) or by increasing the severity of the paralysis — to make it so.) You know that if you could consult each of them individually about what they would have you do, they would all agree: administer Utilicycline. And it's not implausible to think that if everyone who has a stake in the decision wants you to do something, that's what you morally ought to do. So, you morally ought to administer Utilicycline. Call this the *Argument from Individualistic Ex Ante Concern*.⁵

The conflict between these two arguments is the subject of this paper. I, like several others [most notably, Frick, 2013; Nissan-Rozen, 2017], think that the former argument — the one concerning Dominance — is mistaken. And, again like several others [most notably, Frick, 2015; Hare, 2016], I think that something like the Argument from Individualistic *Ex Ante* Concern is basically correct. However, I also think that, although the Argument from Dominance fails, it contains a kernel of truth. And this kernel can be cultivated into a powerful objection

³ This is a very quick and rough version of the argument defended by Fleurbaey and Voorhoeve [2013].

⁴ A one-in-twenty-six chance works out to be a little less than 3.85%, which, according to Ghaferi et al. [2009], is only slightly higher than the death rate associated with inpatient surgery in “very-low-mortality” hospitals (3.5%) and a little more than three percentage points lower than the rate at “very-high-mortality” hospitals (6.9%). It does not seem implausible to me to suppose that each of the villagers might be willing to undergo a surgical procedure at a somewhat-low-mortality hospital in order to avoid paralysis of the left arm.

⁵ This argument, or something reasonably close by, is appealed to by Frick [2015] and Hare [2016].

against Non-Utilitarian moral views that accept the Argument from Individualistic *Ex Ante* Concern. These views will, in some cases, enjoin you to *play ostrich*: to avoid morally relevant bits of information, even when doing so involves incurring costs to yourself or to the very others to whom you are meant to show concern. Or so I will argue.

Of course, one person's objection is another's "interesting consequence." But if you, like me, find it absurd that morality might require us to make others worse off so that we can maintain our willful ignorance, something must go: either the Argument from Individualistic *Ex Ante* Concern or the Non-Utilitarian verdict that we shouldn't sacrifice the wellbeing of the few for the good of the many. Which to give up I leave up to you.

Here's a roadmap. Our first stop is to take a closer look at the Argument from Individualistic *Ex Ante* Concern. It needs to be clarified in various ways. Then, in §3, I will present the Argument from Dominance and argue that it fails. In §4, I will argue that Non-Utilitarian views that accept *Ex Ante* Concern are objectionable.

2 A Defense of Ex Ante Concern

Suppose that you must make a decision, the effects of which will befall only one person. It's not implausible to think that, in such cases, you should aim to do what you think *that* person would decide to do were they in your position, knowing what you know and deciding on their own behalf.

DEFERENCE TO PRINCIPAL'S PERSPECTIVE

When making a decision that affects only one person (call them 'Person X'), if you know that your decision will only affect Person X and that Person X would choose to ϕ were they in your position, you morally ought to ϕ .

What do I mean by "in your position"? I mean something more than just "able to make this decision for themselves." For suppose that you, the decision-maker, know that ϕ ing will result in Arden's death, but that Arden does not know this. If Arden were making this decision for themselves, they (unfortunately) might opt to ϕ . If you know that Arden would ϕ in such a case, does that mean that you morally ought to ϕ , condemning Arden to an untimely death? No. You shouldn't defer to what you think a person might *actually* choose given the opportunity to do so for themselves if you know that you are better informed than they are and that, if they knew what you knew, they would choose differently. Instead, you should only defer to what a person might *hypothetically* choose to do were they to be as

equally informed about their situation as you are.⁶

The principle requires various restrictions. For example, if Arden has done something deserving of punishment and it's your job to punish those who deserve it, choosing to punishing Arden might be what you morally ought to do even if you know that, if left up to them, they would choose otherwise. So, the principle doesn't hold across the board. It should, at the very least, be restricted to cases in which desert-based considerations aren't at play. Such considerations aren't at play in any of the cases we will encounter, so let's ignore this complication.

Here's another complication. Suppose that, in order to survive, Arden requires an injection but that they have such a strong fear of needles that, were they deciding for themselves, they would find holding the syringe so terribly unpleasant that they'd opt to instead say their goodbyes and throw in the towel.⁷ Although Arden would choose death over administering their own injection, they might be happy with *you* administering the injection for them. As stated, the principle says that you should let Arden die. That doesn't seem right. In light of cases like these, we might want to revise the principle to say something like, "if you know that Person X would *want you* to choose to ϕ were they as informed as you are, then ...," but, because nothing in this paper will turn on the distinction, let's ignore this complication as well.

What if Arden is a self-destructive masochist who you know is disposed, at least in some contexts, to act contrary to what is in their best interest? (Or, to choose a more quotidian example: What if Arden is weak-willed about certain matters, including ones relevant to the decision at hand?) Should you obey the principle and defer nonetheless? This is a notoriously thorny issue, pitting two different important values against each other: expressing respect for a person's autonomy and reasonable concern for a person's welfare. Consider the following alternative principle:

DEFERENCE TO EXPECTED BENEFICENCE

When making a decision that affects only one person (call them 'Person X'), if you know that your decision will only affect Person X and that ϕ ing is what's expectedly best for them, then you morally ought to ϕ .

These principles will issue conflicting recommendations if someone is disposed to act contrary to their interests even when those interests are the only ones at stake. How this conflict should be resolved isn't something to be settled here. We will only be concerned with cases in which these two principles coincide: that

⁶ What if you know that Arden is *better* informed than you? Suppose, for example, that *you* don't know whether Arden has the genetic anomaly, but you know that *Arden* knows. This is an interesting case. Let's ignore it for now. We will return to it later.

⁷ This example is inspired by the Anxious Patient in [Enoch \[2017\]](#).

is, cases in which if X would choose to ϕ were they in your position, this is (at least, in part) because they would recognize that ϕ ing is what's expectedly best for them.

A person's hypothetical choices might depart from what is expectedly best for them because they are weak-willed or self-destructive or prone to bouts of practical irrationality, etc. As mentioned, we'll set these cases aside. But it's worth mentioning an entirely different source of potential conflict: Person X might not be disposed to choose to ϕ , even though ϕ ing is what's expectedly best for them, because they are sensibly risk-averse with respect to their own wellbeing. Offhand at least, there doesn't appear to be anything irrational about being risk-averse in this way.⁸ It seems quite sensible to play it a bit safe when your wellbeing is on the line. And so, in cases like these, I think it would be wrong to choose the risky, expected-wellbeing-maximizing option when you know that Person X would, were they making the decision for themselves, choose a less risky option.⁹ Ideally, then, we'd revise the Deference to Expected Beneficence principle to account for differing attitudes toward risk. We could, following Hare [2016], talk about acting "from reasonable concern" for the person; or we could, following Frick [2015], talk about doing what is "justifiable *ex ante*" to the person. Nothing too important (at least for my purposes) turns on this issue, however, so let's ignore it.¹⁰

⁸ I say "offhand at least" because Broome [1991, 2015] argues that, because we don't have an independent basis for a quantitative conception of wellbeing (or, in Broome's terminology, "personal goodness") that is measurable on a cardinal scale aside from the role it plays in aggregating wellbeing under uncertainty, we should treat *utility* as a scale of personal goodness. Broome thinks: if that's right, so that the two (utility and personal goodness) stand in a linear relationship to one another, risk-aversion about personal goodness makes no sense. And so, while it might *off-hand* seem sensible to be risk-averse with respect to wellbeing or personal goodness or whatever-you-want-to-call-it, if Broome is right, it actually makes no sense. As far as I can tell, however, Broome's argument presupposes that what it is to be risk-aversion with respect to some quantity is for there to be diminishing marginal utility for that quantity. There are other, arguably more accurate, ways to think about risk-aversion [e.g., Buchak, 2013], and it's not obvious to me that the argument works if risk-aversion is understood in these ways instead.

⁹ Buchak [2017, p. 631-2] argues that, in fact, there is an interesting asymmetry concerning risk-taking when deciding on the behalf of others: when making a decision for another, if you don't their attitude toward risk, it is better to err on the side of caution by being more risk-averse, than more risk-loving, than they are. She also claims (plausibly enough) that, if you *do* know a person's attitude toward risk, you should adopt that same attitude when choosing on their behalf. In any case, it definitely seems wrong to knowingly take greater risks, when deciding for others, than they'd be willing to take for themselves.

¹⁰ Although perhaps not too important for my purposes, something important *does* turn on this issue, which is worth mentioning. If people are risk-averse with respect to their own wellbeing, and you aim to act "from reasonable concern for," or to do what is "justifiable *ex ante* to," each person, there is a class of Non-Utilitarian consequentialist views that avoid the objections considered in this paper: namely, Prioritarian views. It's hard to say why without getting ahead of ourselves, but this is the basic idea. Being risk-averse, roughly, involves putting extra weight on the possible

One final worry. Suppose that, as far as you and Arden are concerned, ϕ ing is the thing to do. It maximizes their expected interest, it's what they would choose in your shoes, their wellbeing is the only relevant consideration, etc. But ϕ ing isn't guaranteed to be in Arden's interest; it is what's *expectedly* best, not necessarily what's *actually* best. One might object to the principle on the grounds that, ultimately, what matters is Arden's interests, not their *expected* interest. And so, what you ought to do is whatever would actually — and not expectedly — be best for Arden. Of course we won't always know what would actually be best, and so we won't always know what we ought to do, but so what? It might be right there is a perfectly good sense of 'ought' — what is often referred to as the *objective* 'ought' — according to which you ought to do whatever it is that would actually be best. But there is also a perfectly good sense of 'ought' — the *subjective* 'ought' — that's sensitive to your and Arden's uncertainty about how the world is. We should construe the 'ought' in the consequent of the principles above subjectively.

So, when your decision affects only one person, you (subjectively) ought to do what you think that person would want you to do were they as well-informed as you are. Things become more complication when your decision affects more than one person. Suppose that your decision affects the wellbeing of both Arden and Baldwin and no one else. You might know that Arden, were they in your position, would have you ϕ ; and you might know that Baldwin, were they in your position, would have you ψ . You can't do both. So, in deciding what to do, you will need to find some way to navigate this conflict in interests.

But suppose instead that you know both Arden and Baldwin would, were they in your position, want you to choose ϕ . It's very tempting to think that in cases like these — cases in which you must make a decision affecting exactly n people, and in which you know of some particular option available to you that each of the n people would, were they able to make the decision on their own behalf, choose that option — you ought to choose ϕ . If everyone who would be affected by your decision agrees about what you should do, it's simple: you should do it.

THE PRINCIPLE OF EX ANTE CONCERN

When making a decision that affects n people (Person A, Person B, ...), if you know, for each of these people, that they would want you to ϕ were they in your position because ϕ ing is what's expectedly best for them, you

outcomes in which the decision would make you worse off. For example, if you are *maximally* risk-averse, you will (pessimistically) only attend to the worst-case scenario. If you aim to make a decision that would be acceptable to a number of risk-averse individuals, in giving greater weight to the worst case scenarios, you'll end up bringing about outcomes in which the worse-off aren't as worse-off as they might've been had you aimed to maximize what was in each individual's expected interest. Rawls [1971]'s Veil of Ignorance argument for the Difference Principle serves as a stark, and obviously quite famous, example.

morally ought to ϕ .¹¹

All of the same clarifications, restrictions, complications, and apologies mentioned in connection to the principles above apply here as well. I won't repeat them.

What can be said in favor of this principle? There are several things. Deontologists might emphasize its connection to *hypothetical consent*. If everyone affected by the decision would ϕ if they could, refusing to ϕ might seem like an affront to their autonomy.¹² And while some people think it's okay to paternalistically act against the wishes of others when doing so will significantly promote their wellbeing, in this case, you'd be acting against their expected interests as well.

Contractualists might support the principle by emphasizing the importance of acting in ways that are justifiable to each person *individually*. If everyone affected by the action wants you to ϕ , it's hard to see how ϕ ing wouldn't be justifiable to each person: those not affected by the decision have no reasonable grounds to object (because it doesn't affect them), and those who *are* affected by the decision surely cannot reasonably object to you doing what they each want you to do.¹³ Denying this invites a challenge that, to me, seems difficult to meet. If someone *did* have a reasonable objection to you ϕ ing, would they also have one if the decision affected them and them alone? If so, this is either because it's never permissible to subject someone to a risk (even when it's a risk that they would want you to take) or because all that matters morally are the actual consequences

¹¹ This principle is very closely related to the much discussed *Ex Ante Pareto Principle*. Typically, however, this principle is put in terms of preferences: e.g., "If everyone *prefers* ϕ to ψ , then ...". Often, among economists, preferences are (implicitly) asked to play two different roles. First, they are what explains why agents choose what they do. In fact, sometimes, preferences are understood to be representations of an agent's choice-dispositions. Second, they play a normative role in characterizing an agent's welfare. Roughly, if you prefer something, it makes you better off to get it. These two different roles are reflected in the two different Deference principles above, the one having to do with what someone would choose to do were they making the decision for themselves and the other having to do with what would be in their expected interest. To forestall confusion, I think it's better to make this distinction more explicit by avoiding talk of preference.

¹² Developing this argument is not easy, and the issues become quite baroque [e.g., Kamm, 2001]. Moreover, the notion of hypothetical consent is itself somewhat puzzling when you really think about it [see, e.g., Enoch, 2017], and so it probably doesn't provide the most solid foundation for the *Principle of Ex Ante Concern*.

¹³ This is the idea behind, what is called, *ex ante contractualism* [Frick, 2015]. And, although I said "surely" above, some contractualists absolutely *would* deny that if ϕ ing is in everyone's expected interest, no one has grounds to reasonably object to you doing so. The reason is that these contractualists — *ex post contractualists* — think that the right perspective from which to assess whether an action is justifiable to each person is the hypothetical situation in which all those affected by the decision are fully-informed about what the effects of the decision will be. For these contractualist, what everyone would agree to under conditions of uncertainty is largely irrelevant. But it's not obvious why that should be. As far as I can tell, *ex post contractualists* reject the *Principle of Ex Ante Concern* because they dislike the consequences of accepting it (see §3 on the Dominance Argument), which is fair enough I suppose, but it is somewhat dialectically unsatisfying.

of your actions. Both of these are implausible. On the other hand, if the person would *not* have a reasonable objection in the single-person version, how does the introduction of other people make what was justifiable to a person unjustifiable given that their presence in no way affects that person's interests? There might be a satisfying answer to this, but I can't see what it would be.

Finally, here is what I take to be the strongest argument for the *Principle of Ex Ante Concern*: the Agglomeration Argument.¹⁴ Recall the variant of the case that opened the paper: Opaque Two Treatment. In this case, administering Egalitrex ensures that everyone will survive but with paralysis in the left arm, and (because no one knows who has the genetic anomaly) administering Utilicycline gives each person a twenty-five-out-of-twenty-six chance of making a full recovery and a one-out-of-twenty-six chance of death. Suppose that, unlike in the original version of the case in which one or the other treatment can only be administered by mixing it into their water supply, we have twenty-six bottles of water — one per person — that we can mail to each of them individually, mixing in one or none of the treatments. What should you do?

You must make twenty-six decisions: which of the treatments to mix into Arden's bottle of water, which of the treatments to mix into Baldwin's bottle, ..., and which of the treatments to mix into Zephyr's bottle. Each decision affects only one person, you know that it does, and you know who that person is. Furthermore, for each of the villagers, you know that they would want to be administered Utilicycline (and that they each want this for a good reason: it's what's in their expected interest). According to the Deference principles, then, for each of the villagers, you ought to mix Utilicycline into their bottle of water and send it on its way. In other words: you ought to mix Utilicycline into Arden's bottle (\mathbb{U}_A) and send it on its way, you ought to mix Utilicycline into Baldwin's bottle (\mathbb{U}_B) and send it on its way, ..., and you ought to mix Utilicycline into Zephyr's bottle (\mathbb{U}_Z) and send it on its way. But if you ought to do this and you ought to do that and you ought to do the other thing, then you ought to do this, that, and the other thing. And so, in this case, you ought to give them all Utilicycline (\mathbb{U}) even though this will surely result in one of their deaths. This last step appeals to an agglomeration principle governing oughts:

UGHT AGGLOMERATION

If you ought to ϕ_A , and you ought to ϕ_B , ..., and you ought to ϕ_Z , then you ought to (ϕ_A and ϕ_B and ... and ϕ_Z).

Why accept this principle? Hare [2016] offers a compelling argument: if a moral

¹⁴ Hare [2016] presents the most thorough version of this argument. Horton [2017] makes use of a somewhat similar idea, not in the service of arguing for something like the *Principle of Ex Ante Concern* (which he rejects), but rather as a counterexample to *ex post contractualism*.

theory violates Agglomeration, it will sometimes be the case that there is literally nothing you could do to comply with it. Any moral theory that says, of some actions, that you ought to perform *each* of them but then goes on to say that you ought *not* perform them *all* is putting you in an impossible situation. Moral theories are supposed to help guide our actions, though, and that's not compatible with recommending that we do things that cannot be done.

This shows that, when we can mail them the treatments individually, we ought to give them all Utilicycline, but does this show that we should administer Utilicycline rather than Egalitrex when, as it is in Opaque Two Treatments, our only options are to either give them *all* no treatment, give them *all* the one treatment, or give them *all* the other treatment? It's hard to see why this should matter. If you'd be mailing each of them the same treatment anyway, why not save yourself the trouble (and the postage) and just dump it into their water supply?

That's the Argument from Individualistic *Ex Ante* Concern. In Opaque Two Treatments, you ought to administer the Utilicycline rather than the Egalitrex. Why? Because that's what each person who has a stake in the decision wants you to do.

3 A Critique of the Dominance Argument

We just saw that there is a fairly compelling argument for administering Utilicycline. But there is also an argument for administering Egalitrex instead. I think that this argument fails (but for instructive reasons), so I will discuss it here before in the next section turning to what I think is the more troubling consequence for Non-Utilitarians who accept the *Principle of Ex Ante Concern*.

Here's a very rough characterization of the Dominance Argument as applied to Opaque Two Treatments. Let's, first, represent your moral quandary as a decision-problem. It is helpful to think of a decision-problem as being characterized by three different entities: there are your *options*, which are the things you can do (in this case, *Do Nothing*, *Administer Egalitrex*, and *Administer Utilicycline*); there are the *states*, which are ways the world might be that are outside of your control, and are the loci of your uncertainty; and, lastly, there are the *outcomes* that would result from choosing an option in a given state, and are the loci of value.

Recall that in Opaque Two Treatments, one of the twenty-six villagers, you know not which, has the genetic anomaly making Utilicycline ineffective for them. So, there are (at least) twenty-six relevant states: one for each villager who, for all we know, might have the genetic anomaly. Let S^A be the state in which Arden has the genetic anomaly, S^B be the state in which Baldwin has it, ..., and S^Z be the state in which Zephyr has it. For each of these states, there are three outcomes: the outcome that would result from opting to *No Nothing* in that state (i.e., all twenty-

six die), the outcome that would result from opting to *Administer Egalitrex* in that state (i.e., all survive with paralysis in their left arm), and the outcome that would result from opting to *Administer Utilicycline* in that state (i.e., all make a full recovery except for the one person, singled out in that state, who dies). Dominance Arguments, characteristically, draw conclusions about the value of your options on the basis of comparisons between the outcomes of those options residing in the same states. And, as we saw with the example that opened the paper (i.e., the *transparent* version of Two Treatments in which Arden has the genetic anomaly), it seems like we shouldn't administer Utilicycline rather than Egalitrex in that state (S^A). But it looks like the same reasoning (from the introduction) should hold for the comparisons made in each of the twenty-five other states as well. And so we shouldn't administer Utilicycline rather than Egalitrex in Opaque Two Treatment either.

That's a very rough statement of the idea, and it's worth trying to be a bit more precise, so let's take a closer look at Dominance Arguments in general, and how this one relates. I'll illustrate the basic idea with an example.

Opaque Prize Box. Before you is an opaque box. What's inside? You don't know. Hopefully it's something good because you've been offered the following deal:

Box You can take the opaque box home with you, including whatever is inside.

Box⁺ You can take the opaque box home with you, including whatever is inside plus \$100.

Taking the opaque box plus the \$100 (*Box*⁺) is clearly better than taking the opaque box without the additional money (*Box*), so you rationally ought to choose *Box*⁺ over *Box*. The former is *clearly* better than the latter because it dominates it. Consider all the very many different things the opaque box might contain (e.g., a gym sock, the keys to a 2011 Chevy Impala, a half-eaten turkey club on rye, your grandfather's ashes, an exhaustive list of your shortcomings, the hope diamond, ...). For each of the things that might be in the box, it plus \$100 is better than it alone. And if, for every way the world might be, the outcome of the one option is better than the outcome of the other, the former is said to *dominate* the latter. And, if one option dominates the others, you should take. That's (very roughly) the Dominance Principle.

Why believe in the Dominance Principle? I have an opinion about this, but not the space to properly defend it. Instead, I'll briefly explain how I think we should understand the principle without properly arguing for it.¹⁵ In contexts

¹⁵ I attempt to offer a proper defense of the view in Doody [2017].

like these, our ultimate concern is securing value. If we were to know how valuable our options *actually* are, we should take the one that is actually the most valuable. Uncertainty is ubiquitous, however, so we rarely do know how valuable our options actually are. But, if one option (e.g., Box^+) dominates another (e.g., Box), you are thereby in a position to know for certain that the former is actually more valuable than the latter.¹⁶ You won't necessarily be in a position to know how valuable either option actually is in some absolute sense, but you will be in a position to know how the actual values of the options *compare*. And, when it comes to choosing, comparisons are all that are needed.

Offhand, though, it's not obvious how this decision-theoretic principle could be relevant to what one should do in cases like Opaque Two Treatment. Although the subjective notion of value (e.g., utility) often at play in decision theory can easily be substituted for a more objective one (e.g., overall goodness), many Non-Utilitarian moral views either don't endorse a straightforward relationship between value and choice in any case. Of the views mentioned so far, Egalitarians are probably the best targets of the Dominance Argument because they do assign values to outcomes (some of that value is a reflection of the wellbeing of others, some of it is a reflection of how equally that wellbeing is distributed), and do endorse a fairly straightforward, consequentialistic relationship between value and choice.¹⁷ But other Non-Utilitarian views aren't like this. For example, on some views, it just makes no sense to talk about the value of an outcome. Outcomes, worlds, states-of-affairs, etc. can't be good or bad, or better or worse, *simpliciter*; rather, they can only be good or bad, or better or worse, *for* something or someone [e.g., Thomson, 2001]. There are other views on which it does make sense to rank outcomes in terms of their overall value, but these rankings don't always (or

¹⁶ At the risk of belaboring a point that is both nonessential to this paper and probably clear enough already, the basic idea here is that the Dominance Principle (understood correctly) encodes a sort of reasoning-by-cases. Dominance holds relative to a partition of states. In general, we can think of the cells of the partitions (or, the "states") as corresponding to hypotheses about how valuable your options might actually be. When one option dominates another (relative to the right kind of partition), for each of these hypotheses H , we're entitled to believe the conditional: "If H , then the first option is actually more valuable than the other." Because these hypotheses form a partition, they are mutually exhaustive. This allows us to reason-by-cases, discharging the antecedents of these conditionals, arriving at the conclusion that the first option is actually more valuable than the other. Add to this a bridge principle that says: "If you know for certain that the one option is actually more valuable than the others, you rationally ought to take it."

¹⁷ Considering each state in turn, the Egalitarian might think that the inequity in the distribution of wellbeing resulting from administering Utilicycline makes that outcome worse than the outcome that would result in that same state by administering Egalitrex instead. Because this is true for every state, the Egalitarian can conclude that, however the world turns out to be, the outcome that would result from administering Utilicycline will be worse than the one that would result from administering Egalitrex instead. And because the Egalitarian doesn't think one ought to bring about morally worse outcomes, the view recommends administering Egalitrex in Opaque Two Treatment (just as it would in the transparent version of the case).

even typically) play a significant role in what you ought to do. Other things — like, rights — matter too, and matter more [e.g., Kamm, 2001]. And, even among consequentialists, some hold that because the value of different people’s lives are incommensurable, the relevant outcomes are neither better, worse, nor equally as good as each other [e.g., Hare, 2013].

I think this worry is misplaced, however. While it’s true that the Dominance Principle is stated in terms of notions like value, utility, and preference, the idea underlying the principle can be liberated from them. In a decision-theoretic context concerning practical reason, if you know that one option is actually more valuable than the others, then you also thereby know something about the *objective* ought of rationality: that’s the option that you — in some more objective sense — rationally ought to take.¹⁸ Rather than ‘actual value’, we could just as well say that if one option dominates the others, you are in a position to know that it maximizes *objective* value; and rather than ‘objective value’, we could just as well say that you are in a position to know specifically what you objectively ought to do. In other words, we can try to excise all mention of ‘value’ so as to make certain decision-theoretic principles, like the Dominance Principle, more friendly to non-consequentialists.¹⁹ Call it the *Dominance for Everyone Principle*. It has two parts. The second of the two parts, which we’ve just discussed, provides a bridge from the knowledge of what you objectively ought to do to what you subjectively ought to do.

OBJECTIVE-TO-SUBJECTIVE BRIDGE

If you know that you objectively ought to ϕ , then you subjectively ought to ϕ .

I think this principle is undeniable; when I try to imagine how it could be false, I start to lose my grip on what ‘objectively ought’ and ‘subjectively ought’ are supposed to mean.²⁰ The first of the two parts cashes out what it is for one option

¹⁸ What is this “more objective sense”? The ‘ought’ of practical reason is doubly subjective: first, the value-function represents whatever it is that *you* care about; second, it is sensitive to how you take the world to be, rather than fixed to how the world actually is. The objective ought of rationality is still subjective in this first way; the values are up to you. But it’s not subjective in this second way; it tracks what would *actually* best serve your ends.

¹⁹ Alternatively, we can still talk about the “value” of outcomes, on various non-consequentialist views, in a nonstandard, derivative sense. One strategy here would be to identify the value of an outcome, according to some value-unfriendly non-consequentialist theory, with the value that an extensionally-equivalent consequentialized moral theory assigns to that outcome. Whether every non-consequentialist moral theory can be faithfully consequentialized, in a way that would allow for this, is controversial [e.g., Brown, 2011]

²⁰ While this principle might be undeniable, be careful not to mistake it for some nearby principles that very much are deniable. For example, it does *not* follow that you subjectively *ought not* to ϕ from the fact that you know you objectively ought not to ϕ . The Miners Puzzle (as well as the Pill example from Jackson [1991]) serve as counterexamples to that principle.

REASONING-BY-CASES: OPAQUE TWO TREATMENT

- P1** \mathcal{S}^A or \mathcal{S}^B or ...or \mathcal{S}^Z
- P2** If \mathcal{S}^A , then you objectively ought to ϕ .
- P3** If \mathcal{S}^B , then you objectively ought to ϕ .
- ⋮
- P27** If \mathcal{S}^Z , then you objectively ought to ϕ .
-
- C** You objectively ought to ϕ .

to “dominate” another without appealing to value quantities in a way that might make non-consequentialist squirm.

REASONING-BY-CASES

Let $\mathbf{S} = \{S_1, S_2, \dots, S_n\}$ be a partition of the ways the world might be. If, for all $S \in \mathbf{S}$, you objectively ought to ϕ if S is the way the world actually is, then you know you objectively ought to ϕ .

Putting these two principles together delivers us a kind of dominance principle that doesn’t make reference to things like value, or preference, or utilities. Armed with the *Dominance for Everyone Principle*, which cannot be so easily dismissed by non-consequentialists, let’s return to the Opaque Two Treatments case to see how the Dominance Argument holds up.

At first glance, it appears to hold up very well: $\{\mathcal{S}^A, \mathcal{S}^B, \dots, \mathcal{S}^Z\}$ is a partition of the ways, for all you know, the world might be given that you don’t know which of the twenty-six villagers has the genetic anomaly; because, according to the Non-Utilitarian moral views under consideration, you shouldn’t knowingly allow Arden to die in order to confer comparatively meager benefits on to the others, if \mathcal{S}^A is the way the world actually is, then you objectively ought to administer Egalitrex; because you shouldn’t knowingly allow Baldwin to die for the same reasons, if \mathcal{S}^B is the way the world actually is, then you objectively ought to administer Egalitrex; ...and so on; therefore, by *Reasoning-by-Cases*, you know you objectively ought to administer Egalitrex; and so, by *Objective-to-Subjective Bridge*, you subjectively ought to administer Egalitrex in Opaque Two Treatments, which contradicts the conclusion of the previous section.

First glances can be misleading. The argument, as others have noted [e.g., Frick, 2013], is flawed. The argument illegitimately slides between, on the one

hand, the situation in which you allow, e.g., Arden to die in order to confer benefits on others but don't know at the time of the decision that it will be Arden who dies and, on the other, the situation in which you *knowingly* allow Arden to die to confer benefits on the others. Non-Utilitarian moral views who think it's wrong to administer Utilicycline in the transparent versions of Two Treatments, but who also (because of the *Principle of Ex Ante Concern*) think it's what you ought to do in the opaque version, can respond to the Dominance Argument by maintaining that what you, the decision-maker, know about the identity of the person who would die were you to administer Utilicycline is morally relevant, not only to what you *subjectively* ought to do, but to what you *objectively* ought to do. Your ignorance about who has the genetic anomaly doesn't merely affect how you should try to conform to what morality requires of you, it alters the normative landscape itself.²¹ And this is a consequence that many Non-Utilitarians are more than happy to accept.

The argument fails, then, because it trades on an equivocation. In order for it to be valid, *Reasoning-by-Cases* requires, for each of the twenty-six states, a conditionals of the form:

If S^x is the way the world actually is, then you objectively ought to administer Egalitrex.

But one of those conditionals is if not false, very close to the very thing the argument is seeking to establish. In particular, suppose that S^A is, unbeknownst to us, the actual way the world is. Then, because its antecedent is true, the conditional,

If S^A is the way the world actually is, then you objectively ought to administer Egalitrex.

is true (and the argument sound) only if its consequent — that you objectively ought to administer Egalitrex — is true. But that's what's at issue. And, as just

²¹ Why might this be? The answer will depend on the particular kind of Non-Utilitarian view. For example, *ex ante contractualists* [like, Frick, 2015], will emphasize that doing what is *ex ante* justifiable to each person crucially depends on what we are, are not, ignorant about. The less we know about what will happen to whom, the easier it will be for offer *ex ante* justifications. Others [like, Hare, 2016, 2013, 2010], will emphasize the ways in which knowing more about the identities of those affected by your actions influence what reasons you have for action. Ignorance about who is who can give rise to symmetries between the reasons you have, shielding you from having other reasons that might, were you to have them, render currently permissible actions impermissible. That's not to say that nothing puzzling remains for these views. In particular, as Hare [2016] points out, it's far from clear what it even *is* to know who someone is, in the sense that seems to be operative in these case. For example, we can coin a name, 'Gene', stipulating that it refers to the villager who has the genetic anomaly and so who will die if we administer the Utilicycline. There now is a sense in which we know who will die if we administer the Utilicycline: Gene will! Or what if I know the social security number of the person who will die if we administer the Utilicycline? For very interesting discussions of these sorts of questions, see Setiya [forthcoming]; Mahtani [2017].

mentioned, it's open to the Non-Utilitarian to say that, while it *would* be true that you objectively ought to administer Egalitrex if you *knew* that S^A is the way the world actually is, the fact that you *don't* know this affects what you objectively ought to do. On the other hand, all of the conditionals of the following sort are true:

If you knew that S^x is the way the world actually is, then you objectively ought to administer Egalitrex.

But then the argument is no longer valid because *those* conditionals don't satisfy the antecedent of *Reasoning-by-Cases*.

Proponents of the Dominance Argument can respond, however, by amending the antecedent of *Reasoning-by-Cases* in order to rescue the argument from invalidity thusly:

REASONING-BY-KNOWN-CASES

Let $\mathbf{S} = \{S_1, S_2, \dots, S_n\}$ be a partition of the ways the world might be. If, for all $S \in \mathbf{S}$, you objectively ought to ϕ if you *knew that* S is the way the world actually is, then you know you objectively ought to ϕ .

But, while *Reasoning-by-Cases* encodes what is typically a valid form of deductive reasoning, *Reasoning-by-Known-Cases* does not.²² Here is a counterexample. I'm a curious person and I am very interested in knowing whether or not p is the case. You know the truth about p but — always the shrewd entrepreneur — you offer to tell me the truth for a \$1. That's a pretty good deal, so I reach for my wallet. But, first, I decide to do some reasoning to make sure that this transaction is one with which I ought to engage:

Either p or $\neg p$. If I knew that p , then I ought *not* pay you the \$1 because it would be a total waste of money; I already know that p is true, so what a I paying you for? On the other hand, if I knew that $\neg p$, then I ought *not* pay you the \$1 because, again, I would already have the very thing I was paying for. Therefore, I ought not pay you the \$1 to learn the truth about p .

This (obviously bad) argument generalizes, too, of course: I should never incur any cost to learn something. The lesson from this example, and the invalidity of

²² I say 'typically' because one lesson some have learned from the Miners Puzzle is that Reasoning-by-Cases *isn't* valid when it includes indicative conditionals with information-sensitive consequents. Come to think of it: if the *objective* 'ought' is sensitive to what you, the decision-maker, do and do not know about the identities of the people affected by your choices, we should be careful about appealing to Reasoning-by-Cases as a valid form of deductive reasoning, and therefore suspicious of *Reasoning-by-Cases* as well.

REASONING-BY-KNOWN-CASES: OPAQUE TWO TREATMENT

- P1** \mathcal{S}^A or \mathcal{S}^B or ...or \mathcal{S}^Z
- P2** If KNOW (\mathcal{S}^A), then you objectively ought to ϕ .
- P3** If KNOW (\mathcal{S}^B), then you objectively ought to ϕ .
- ⋮
- P27** If KNOW (\mathcal{S}^Z), then you objectively ought to ϕ .

C You objectively ought to ϕ .

Reasoning-by-Known-Cases more generally, is that what you would do were you better or fully informed is not always a good guide to what you should do when you aren't as well-informed.

This suggest one more move for the proponent of the Dominance Argument. Instead of asking yourself what you would do were *you* fully informed, maybe we should defer to a hypothetical, fully informed advisor who has our best interest at heart instead?

DEFER TO YOUR FULLY INFORMED ADVISOR

If you know that a fully informed advisor who has your best interests at heart would advise you to ϕ , then you subjectively ought to ϕ .

This nicely handles the silly example about (not) paying to learn the truth about p . Although you wouldn't pay to learn the truth if you were fully informed (because you would already know it), your fully informed advisor, who in virtue of being fully informed knows the truth about p , knows that you *don't* know the truth about p but would very much like to. So, although you know that if *you* were fully informed you wouldn't pay to learn about p , you don't know that your fully informed advisor would advise you against paying to learn about it. So that counterexample is one no longer.

Can *Defer to Your Fully Informed Advisor* resuscitate the Dominance Argument? It's not clear. Your fully informed advisor, in virtue of being fully informed would know that, e.g., \mathcal{S}^A is true and so that it is Arden whose life is on the line. She would also know that you don't know that. Furthermore, because she has your best interests at heart, and doing what you morally ought to do is something you very much care about, your fully informed advisor will aim to advise you to do as you morally ought to do in this case. If your advisor was the one deciding which of the treatments to dump in the water supply, she would have to choose

Egalitrex. She knows that Arden's life is on the line, and it is wrong to knowingly sacrifice Arden. But she knows that it is *you*, and not her, making the decision. And she knows that you, unlike her, *don't know* who would and wouldn't survive the Utilicycline treatment; and so your decision to administer Utilicycline can be seen as an expression of concern for the expected interest of all. And part of what is at issue is whether what you know and don't know about your situation is relevant to what you objectively ought to do. To presume that you are in a position to know that your fully informed advisor would advise you against administering the Utilicycline, when part of what is at issue is whether being as informed as your advisor is changes what one objectively ought to do, seems to beg the question.

So the Dominance Argument against doing what is expectedly best for all either fails or begs the question (which is its own kind of failure). But, nevertheless, I can't shake the feeling that there is *something* objectionable about doing something that you know would be deeply morally wrong were you to know a little bit more about what you were doing. In the final section of the paper, I will attempt to turn this feeling into an argument.

4 Playing Ostrich

Consider the following example:

Opaque Two-Course Treatment. As before, twenty-six villagers in a remote area (Arden, Baldwin, Clay, ..., and Zephyr) have contracted a terrible, life-threatening disease. If nothing is done, they will all die. There are two possible treatments that could be administered. The first is trusty, ol' *Egalitrex*. The second treatment, however, is different. It is a two-course treatment (*Utilicycline Diachronazole*). If you opt for this one, you'll need to administer the second exactly a week after the first. This treatment is a Utilicycline follow-on, and works similarly: when it is effective, it ensures a full recovery; it will be effective for everyone but one (owing to an undiagnosed genetic anomaly), but no one knows who.²³ If Utilicycline Diachronazole is administered, those of whom for which it will be effective will start showing signs of improvement almost immediately; the poor soul for whom it won't, won't. Because the two-course treatment requires both courses, it's important to follow through on administering the second course; otherwise, everyone will get very sick. You can mix the first of the two-courses with *Egalitrex* instead of the second course, and that's better than foregoing the second course en-

²³ It's Arden.

tirely, but the treatments have deleterious effects if taken together, so doing so is sub-optimal. Like before, the only way to administer the treatment is by mixing it into their water supply.

Options	Results
<i>Do Nothing</i>	Arden, Baldwin, ..., and Zephyr all die (o).
<i>Administer Egalitrex</i> (E)	Arden, Baldwin, ..., and Zephyr all survive with paralysis in their left arms (7).
<i>Administer Utilicycline Diachronazole</i> (UD)	Baldwin, Clay, ..., and Zephyr make a full recovery (10); Arden dies (o).

If you're a Non-Utilitarian who accepts the *Principle of Ex Ante Concern*, what treatment ought you to administer? First thought: just as in Opaque Two Treatments, you should administer the drug that is in each person's expected interest. In this case, that drug is Utilicycline Diachronazole. Each person regards their chance of being the one for whom the treatment will be ineffective to be low enough, and the costs of being paralyzed in the left arm high enough, that they'd want you to dump Utilicycline Diachronazole into their water supply. So far, this example is exactly like the one we encountered earlier.

However, here is the catch. Utilicycline Diachronazole is a *two-course* treatment. And after administering the first course, but before administering the second, it will become clear for whom the treatment is not working, it will become clear whose life is being sacrificed for the others. So, if you administer the first course of the treatment, the opaqueness will be lifted, transforming the case into a transparent one. Because it's your job to monitor the health of the villagers, you know now that, if you administer Utilicycline Diachronazole, you will come to know for whom the treatment is ineffective (sadly, it will turn out to be Arden, but you don't know that yet) prior to administering the second course of the two-course treatment. As we've seen, for the Non-Utilitarian, learning this information really matters; it changes the facts about what you objectively ought and ought not do.

What would happen if you tried to do what was expectedly best for each person at this time in Opaque Two-Course Treatment? Suppose you dump the first of the two courses of the Utilicycline Diachronazole into their water supply. And, over the course of the week, you learn that the treatment is not working for Arden. What should you do? Let's game it out.

Two-Course Treatment, Part II. You've given everyone the first course of the two-course treatment. Everyone — with the notable exception of Arden — seems to be doing much better. But you know that if you do not administer the second dose of treatment at the right time, it

will be a disaster: everyone will die. On the other hand, if you administer the second course of Utilicycline Diachronazole, you know that Arden is sure to die, while everyone else makes a full recovery. But, even though mixing the two treatments will have deleterious effects on everyone, no one will die; instead, they will all be paralyzed in *both* arms. These, then, are your only options:

Options	Results
<i>Do Nothing</i>	Arden, Baldwin, ..., and Zephyr all die (o).
<i>Administer Egalitrex</i> (\mathbb{E})	Arden, Baldwin, ..., and Zephyr all survive with paralysis in both arms (6).
<i>Administer 2nd course of Utilicycline Diachronazole</i> (\mathbb{UD}_2)	Baldwin, Clay, ..., and Zephyr make a full recovery (10); Arden dies (o).

What should you do? The decision should feel eerily familiar; it is merely a slightly worse version of the example that opened the paper. Because you now know who it is — it's Arden — who will die if you continue on with the second course of the two-course treatment, administering it is no longer what's expectedly best for each person (it's certainly not expectedly best for *Arden*, for example) and, if you're a Non-Utilitarian of the stripe we've been discussing, it is not morally permissible to administer the second course of Utilicycline Diachronazole because you should instead swap it out for Egalitrex — resulting in upper body paralysis for all of the villagers. That's obviously not *great*, but what else can you do in such a situation?

One thing you could do is to try to avoid being in that situation in the first place. Notice that (nearly) everything that will happen in Two-Course Treatment, Part II is completely predictable *ex ante*. You, of course, can't accurately predict for whom the Utilicycline Diachronazole will be ineffective, but you know that there will be *someone* for whom it doesn't work; and, you know that if you attempt to administer both courses of Utilicycline Diachronazole, you will come to know who that someone is; and, you know that, once you come to know who that someone is, it will not be morally permissible to continue on with the second course of treatment; and, you know that if you morally cannot finish administering the Utilicycline Diachronazole, the only acceptable option that remains is to save each of their lives by administering a dose of Egalitrex which will paralyze them all in both of their arms; and, you know that each of them would rather be paralyzed in just their left arm than in both, and that, before administering the first course of the two-course treatment, you had the opportunity to do something — dump the Egalitrex into their water supply — that would have saved

each of their lives at the lower cost of only left-arm paralysis. If you look ahead to what would happen were you to try to administer the two-course treatment, you should decide against dumping the first of the two-courses because you should recognize that doing so will result in something that is foreseeably worse — being paralyzed in both arms, rather than merely one — for each person. And so, in Opaque Two-Course Treatment, it looks as if you morally ought to administer the Egalitrex even though, at least at that specific time, it is in each person's expected interest *ex ante* for you to, somehow, successfully administer both courses of the Utilicycline Diachronazole.²⁴

So far so good, but now here is the objection. Consider the following adjustment to Opaque Two-Course Treatment:

Two-Course Treatment + Toxin. Everything is the same as before. You've reasoned through your options, and you've decided that what you morally ought to do is administer the Egalitrex. But before you do, something occurs to you: you could, when mixing the treatment into their water supply, also add a mild toxin. The toxin leaves no lasting (physical) damage to anyone who ingests it, but it will make you profoundly ill for several weeks after coming into contact with it. The pain is debilitating; the other symptoms, too stomach-churning to mention. But it is only temporary. If you slip some of the toxin into the water supply with the first course of the two-course Utilicycline Diachronazole, the toxin will make each of the villagers so repulsively ill that it will be impossible to tell who is, and who is not, benefiting from the first course of the treatment. These, then, are your only options:

²⁴ This phenomenon — policies that, while agreed to by each person *ex ante*, are (arguably) immoral to implement once the natural veil of ignorance has partially been lifted — is explored in some depth by Kamm [2001]. Frick [p. 205-12 2015] addresses this issue as well, labeling himself a “stage-wise *ex ante* contractualist” because of his endorsement of the “Decomposition Test”, which roughly says that, in order for a procedure whose implementation involves multiple, temporally-extended stages of voluntary actions to be morally permissible, the actions at every stage must be justifiable to each person at the time they are to be performed. In essence, the Decomposition Test disallows forming, and then following through on, resolutions merely on the basis of that being successfully resolute is justifiable *ex ante*. Clearly, the Decomposition Test would say that it is not morally permissible to somehow successfully implement the two-course treatment of Utilicycline Diachronazole.

Options	Results
<i>Do Nothing</i>	Arden, Baldwin, ..., and Zephyr all die (0).
<i>Administer Egalitrex</i> (\mathbb{E})	Arden, Baldwin, ..., and Zephyr all survive with paralysis in their left arms (7).
<i>Administer Utilicycline Diachronazole</i> (\mathbb{UD})	Baldwin, Clay, ..., and Zephyr make a full recovery (10); Arden dies (0).
<i>Administer Utilicycline Diachronazole + Toxin</i> (\mathbb{UD}_{+toxin})	Baldwin, Clay, ..., and Zephyr make a full recovery (9); Arden dies (-1).

What should you do? We established earlier that, if you are a good Non-Utilitarian, you won't try to administer the two-course Utilicycline Diachronazole because you can predict that, by the time it comes to administer the second course, it will not be permissible for you to do so because you will, at that point, know whose interests are not being served by carrying on with the two-course treatment; administering the second course of the treatment cannot be justified to each person or seen as an expression of reasonable concern for all. And, because your other options (after having administered the first of the two courses) aren't great, you shouldn't try to administer the Utilicycline Diachronazole at all, but instead opt for the Egalitrex.

But what about administering the first course of the Utilicycline Diachronazole with the toxin? The toxin will make the twenty-six villagers so sick that you, and they, will be unable to tell for whom the Utilicycline Diachronazole is working. And if no one knows, when it is time to administer the second course of treatment, exactly who will eventually make a full recovery and who sadly will not, carrying on with the two-course treatment is still what's in everyone's expected interest and it is still justifiable to each of them. At the beginning, the course of action that each of them wants you to perform is to succeed in implementing both courses of the two-course Utilicycline Diachronazole treatment. But, for moral reasons, you are unable to do that. Your Non-Utilitarianism prevents you, when the time comes, from administering the second course of the treatment. Because *that* option is morally unacceptable, but administering both courses of the Utilicycline Diachronazole *plus the toxin* isn't morally unacceptable, and because using the toxin to help successfully implement both courses of the Utilicycline Diachronazole treatment is what's expectedly best (out of the options that are morally acceptable) for each person, it appears as though what you morally ought to do is administer the Utilicycline Diachronazole along side a healthy dose of the vile, stomach-churning toxin. If you accept the *Principle of Ex Ante Concern*, it's hard to see how this could be denied.

	S^A			S^B			...	S^Z			Expected Values		
	A	...	Z	A	...	Z		A	...	Z	A	...	Z
E	7	...	7	7	...	7	...	7	...	7	7	...	7
UD _{bind}	0	...	10	10	...	10	...	10	...	0	9.6	...	9.6
UD _{myopic}	6	...	6	6	...	6	...	6	...	6	6	...	6
UD _{+toxin}	-1	...	9	9	...	9	...	9	...	-1	8.6	...	8.6

The outcome that would result from administering the Utilicycline Diachronazole plus the toxin has a noteworthy property: it is guaranteed to make everyone worse off than they could've been. Consider the twenty-five people for whom the Utilicycline Diachronazole would be effective. They are guaranteed to be worse off than they would've been had you administered the Utilicycline Diachronazole sans toxin. Now consider Arden, the one person for whom the Utilicycline Diachronazole will not be effective. Arden, too, is guaranteed to be worse off than they would've been had you somehow been able to successfully implement both stages of the two-course Utilicycline Diachronazole treatment. Arden, sadly, will die either way. But if you succeed in administering the treatment without the toxin, Arden's final days would be significantly better. Adding the toxin into their water supply makes everyone worse off. And yet that is what you *morally* should do? That strikes me as absurd. Morality should not lead you to do things that are guaranteed to make everyone worse off. Something has gone wrong, and there are two potential culprits: either the *Principle of Ex Ante Concern* or the Non-Utilitarian refusal to countenance certain kinds of interpersonal trade-offs. Together, the two instruct you to make everyone worse off so that you can preserve your ignorance about who, specifically, is being sacrificed for the greater good. In contrast, imagine what the Utilitarian would do in this situation. She would do what everyone would *ex ante* want her to do — in this case, she would administer the first course of the Utilicycline Diachronazole, and then she would stay the course by administering the second course as well. In doing so, she must allow something to happen to someone that, at the time of it happening, is no longer justifiable to that person. That's not great. But let's ask how this very same person would fair in the hands of the Non-Utilitarian. Although there will be no time at which something is allowed to happen to her that is no longer justifiable to her, I imagine that this is exactly the sort of thing it is difficult to enjoy when spending — what is unbeknownst to her (and you, and me) — her final days violently and hopelessly ill.

The ostrich buries its head in the sand in order to maintain its willful igno-

rance.²⁵ The Non-Utilitarian who accepts the *Principle of Ex Ante Concern*, at least in some cases, buries the heads of others in weeks of debilitating illness in order to manipulate what they might learn about the consequences that their actions have for the specific people they affect. The point of manipulating this information isn't to make anyone better off — in these cases, it very much doesn't; it literally makes everyone worse off — or to bring comfort and joy to those who might need it, or to make the world a better place, but rather to maintain the requisite amount of epistemic distance from those affected by your actions necessary to construe the choices you make as unobjectionable them. Let's leave avoiding morally relevant information to the birds. We can do better.

References

- John Broome. *Weighing Goods: Equality, Uncertainty, and Time*. Blackwell Publishers Ltd., 1991.
- John Broome. General and personal good: Harsanyi's contribution to the theory of value. In Iwao Hirose and Jonas Olson, editors, *The Oxford Handbook of Value Theory*. Oxford University Press, 2015.
- Campbell Brown. Consequentialize this. *Ethics*, 121(4):749–771, 2011.
- Lara Buchak. *Risk and Rationality*. Oxford University Press, 2013.
- Lara Buchak. Taking risks behind the veil of ignorance. *Ethics*, 127(3):610–644, 2017.
- Ryan Doody. Actual value decision theory. unpublished manuscript, October 2017.
- David Enoch. Hypothetical consent and the value(s) of autonomy. *Ethics*, 128(1): 6–36, 2017.
- Marc Fleurbaey and Alex Voorhoeve. Decide as you would with full information! An argument against *ex ante* pareto. In Nir Eyal, Samia A. Hurst, Ole F. Norheim, and Dan Wikler, editors, *Inequalities in Health: Concepts, Measures, and Ethics*, pages 113–129. Oxford University Press, Oxford, 2013.
- Johann Frick. Uncertainty and justifiability to each person: Response to Fleurbaey and Voorhoeve. In Nir Eyal, Samia A. Hurst, Ole F. Norheim, and Dan Wikler, editors, *Inequalities in Health: Concepts, Measures, and Ethics*. Oxford University Press, 2013.

²⁵ They don't. It's a myth.

- Johann Frick. Contractualism and social risk. *Philosophy and Public Affairs*, 43 (3):175–223, 2015.
- Amir A. Ghaferi, John D. Birkmeyer, and Justin B. Dimick. Variation in hospital mortality associated with inpatient surgery. *New England Journal of Medicine*, 361:1368–1375, 2009.
- Caspar Hare. Take the sugar. *Analysis*, 70(2):237–247, 2010.
- Caspar Hare. *The Limits of Kindness*. Oxford University Press, 2013.
- Caspar Hare. Should we wish well to all? *The Philosophical Review*, 125(4):451–472, 2016.
- Joe Horton. Aggregation, complaints, and risk. *Philosophy and Public Affairs*, 45 (1):54–81, 2017.
- Frank Jackson. Decision-theoretic consequentialism and the nearest and dearest objection. *Ethics*, 101:461–82, 1991.
- Frances Kamm. *Morality, Mortality Volume II*, chapter Agreements. Oxford University Press, 2001.
- Anna Mahtani. The ex ante pareto principle. *The Journal of Philosophy*, 114(6): 303–323, 2017.
- Ittay Nissan-Rozen. How to be an ex-post egalitarian and an ex-ante paretian. *Analysis*, 77(3):550–558, July 2017.
- Derek Parfit. Equality or priority? Department of Philosophy: University of Kansas, 1991.
- Derek Parfit. Equality and priority. *Ratio*, 10:202–221, 1997.
- John Rawls. *A Theory of Justice*. Harvard University Press, 1971.
- Thomas M. Scanlon. *What We Owe to Each Other*. Harvard University Press, 1998.
- Kieran Setiya. Ignorance, beneficence, and rights. *Journal of Moral Philosophy*, forthcoming.
- Judith Jarvis Thomson. *Goodness and Advice*. Princeton, NJ, 2001.