

COMPLAINTS AND

TOURNAMENT POPULATION ETHICS

Abelard Podgorski

National University of Singapore

Abstract: In this paper, I develop an approach to population ethics which explains what we are permitted to do in virtue of the possible *complaints* against our action. This task is made difficult by a serious problem that arises when we attempt to generalize the view from two-option to many-option cases. The solution makes two significant moves – first, accepting that complaints are essentially pairwise comparative, and second, reimagining decision-making as a *tournament* between options competing two at a time. The right view about when one option defeats another with respect to complaints, paired with the right view about how to pick the winners of the tournament given the relations of defeat, leads to a uniquely attractive view in the person-affecting tradition.

There is a picture of morality according to which there are no victimless wrongs. Whenever an action is wrong, it is because somebody *is wronged*, or, a bit more generally, somebody *has a complaint* against the action.¹ Call this view the *complaints model*. There are various reasons to like this picture – it is arguably intuitive on its face, and it pairs well with a roughly contractualist overarching view of morality as a matter of interpersonal duties we owe each other. But the motivation that is the closest to my own heart is that the complaints model promises to justify a distinctive set of judgments about cases where the number and identity of the people who exist depends on what we do.

The complaints model provides a very natural explanation of a bit of moral commonsense called the *procreation asymmetry*. According to the asymmetry, while it is wrong to bring someone into existence if they would have an absolutely terrible life, bringing someone into existence with a good life is merely morally neutral – we may decide to have a happy child, but we are not obligated to. The complaints model, in the tradition of *person-affecting* approaches to morality, contrasts most strikingly with views, like classical utilitarianism, that explain what is wrong by appeal to what is good *impersonally* – not what is good for this or that individual, but what makes the world better or worse – and these views have trouble accommodating the asymmetry while avoiding implausible implications. By

¹ To say someone has a complaint is meant to be weaker than to say they are wronged, in not entailing that the action is all things considered wrong.

contrast, the complaints model looks like it explains the procreation asymmetry quite straightforwardly, aided only by the plausible assumptions that someone who exists with a terrible life has a complaint against your deciding to create them, and that someone who never exists if you make a certain choice has no complaint against that choice.

The complaints model also inherits a number of challenges, however, from other person-affecting approaches. Some of these, including the infamous *non-identity problem*, manifest even in the simplest cases considered in population ethics. But there is a special class of problems that arise specifically in cases where an agent has three or more options. It turns out to be far from obvious how to generalize a natural treatment of two-option cases in a way that doesn't generate unacceptable consequences.

The main task of this paper is to develop the best version of the complaints model. The plan is this: In the first half of the paper, I will build, from the ground up, a very bare-bones complaints-centered view about decisions with two options which will capture our main motivations, and discuss some objections. Then, in the second half, we will look at a problem that arises when additional options are introduced which raises doubts about the possibility of generalizing the view we've developed. I'll go on to describe a solution, which makes two significant moves – first, accepting that complaints are essentially pairwise comparative, and second, reimagining decision-making as a *tournament* between options competing two at a time. Finally, I'll explain why the tournament approach is worth taking seriously even if one is inclined to reject other aspects of the view.

Part I: Building a View: Two-Option Cases

I'll be understanding complaints in a semi-theoretical way. Roughly, a complaint is the sort of thing that *ceteris paribus* would generate a legitimate objection to my action on behalf of someone. The correspondence between this and our ordinary talk about “complaints” might be imperfect. Two things are worth emphasizing to preclude some potential confusions. First, on my usage, someone could have a complaint against me even if they're unable to complain in the ordinary sense because, for instance, I've killed them. Second, complaints in my sense have a *prima facie* character – someone can have a complaint even if in the end they would not be justified in objecting to my action (for instance because of the greater complaints by others against the alternatives).

Any developed version of the complaints model will have at least two parts. It needs a *theory of complaints* – an account of when individuals have complaints and of the weight of these complaints, and it needs a *deontic principle* – an account of how the facts about various individuals’ complaints determine which options are morally permissible.²

I’ll begin by trying to fill out such a picture, limiting the discussion, for the moment, to decisions with just two options. For clarity and simplicity, and in order to best highlight the way the complaints model comes apart from impersonal, consequentialist alternatives, I will hew as closely as possible to classical utilitarianism while capturing the population-ethical implications that motivate thinking in terms of complaints.³ Ultimately, I think the best version of the complaints model will depart from utilitarianism in more ways than we will pursue here, and I will flag some of these possible departures as they arise.

In trying to keep close company with utilitarians, I will focus only on what we can call *welfare complaints* – complaints that derive entirely from facts about the existence and welfare of agents given each of our actions. I will also assume that we can represent welfare numerically, and that we can make sense of the boundary at which lives go from being good for their bearer to being bad for their bearer, the latter being represented as having negative welfare. A complete picture of morality can, and probably should, acknowledge non-welfarist complaints, and some of these may have implications for population ethics – I will mention a couple when we discuss the non-identity problem later. But I will try to show how a complaints-style view can get desirable results even while limiting itself to the resources shared by its impersonal welfarist competitors.

Classical utilitarianism is a maximizing view. Accordingly, we’ll start by making our version of the complaints model a minimizing view. Let’s begin with a very simple deontic principle:

Minimize Aggregate Complaints: An option X is permissible iff the total strength of complaints against X is no greater than the strength of complaints against any other option.

Minimize Aggregate Complaints is not the only candidate worth considering – some anti-aggregationists, for instance, might hold that we should minimize the strongest individual complaint

² I will be setting aside questions about uncertainty and risk, which raise additional complications beyond the scope of this paper. For our purposes the relevant kind of permissibility will be “objective” – that is, permissibility in light of all the facts.

³ In this way, our approach is similar to what Melinda Roberts (2003) does in developing her “person-based consequentialism”.

instead.⁴ I will later propose an amendment of my own. But assuming some deontic view like this will allow us to derive practical implications of different theories of complaints and help us choose between them.

A Theory of Complaints

Suppose, then, that an agent faces a choice between two options, X and Y, which would bring about outcomes W_X and W_Y , respectively. And suppose we know the facts about which agents exist in W_X and W_Y , and each of their welfares – all of the facts that, given our constraints, determine who has complaints and how strong these complaints are. Let A be some possible person who lives in either W_X or W_Y . When does A have a complaint?

We can begin with the easy case – one where A exists in both W_X and W_Y . The most natural view here is straightforward:

Common Existence Complaints: If X and Y are our only two options, and A exists in both W_X and W_Y , then A has a complaint against X iff A is worse off in W_X than W_Y , and the strength of this complaint is the difference between their welfare in W_X and W_Y .

Again, this is not the only view one might have – most notably, prioritarrians might hold that the strength of someone’s complaint is magnified by their absolute welfare, such that someone who would be worse off by a given amount has a stronger complaint if they are already poorly off. But the debate over fixed-population moral questions like this will largely be orthogonal to the issues that concern us in this paper.

It is easy to see that in cases where everyone who exists will exist no matter what we do, *Common Existence Complaints* paired with *Minimize Aggregate Complaints* gives identical recommendations as classical utilitarianism, which tells us to maximize total welfare. So the interesting work will be done by what we say about complaints when the identities of the people who exist in each option do not completely overlap.

⁴ See Horton (2017) and (2018) for a discussion of and arguments against several such anti-aggregative views.

Neutrality and the Asymmetry

This is a good time to get some of our main desiderata on the table. Suppose that everyone who exists will be equally well off no matter what we do, and that we have the option of creating a new person. Classical (totalist) utilitarianism will imply that we must create the person if, and only if, their welfare would be positive.

However, many people have the intuition that we are not obligated to create new people in this sort of case, even if their life would be quite good. It is especially intuitive that we are not obligated to make significant *cuts* to existing people's welfare merely to create new people, even when this might raise the total. The interests of these merely hypothetical people do not have the same kind of claim on us, the thought goes, as those of people who exist. Or as Jan Narveson puts it, we have reasons to make people happy, but not to make happy people. (1978) This is sometimes called the intuition of *neutrality*:

Neutrality. Ceteris paribus, bringing someone into existence with a good life is *morally neutral* – neither obligatory nor wrong.⁵

Bringing someone into existence with a life that is truly terrible, however – filled, for example, with nothing but pain – does *not* seem morally neutral. It is wrong, in the absence of sufficiently strong countervailing reasons. That is, we have reasons both to avoid making people sad, and to avoid making sad people. Folding these thoughts together, we get what is sometimes called the *procreation asymmetry*:

The Procreation Asymmetry. Ceteris paribus, bringing someone into existence with a good life is morally neutral, and bringing someone into existence with a bad life is morally wrong.

Not everybody shares these intuitions, but they are taken seriously even by many who ultimately reject them (Broome 2004, McMahan 2013), and a view which could capture them would have something significant going for it.

We have already noted that classical utilitarianism is inconsistent with neutrality, and by extension, the asymmetry. But it is quite difficult in general to give an explanation of the asymmetry in terms of what makes the world impersonally better or worse, while avoiding implications that are

⁵ People sometimes characterize the intuition of neutrality in other ways – for instance, in terms of the effect of additional people on the impersonal value of the world. I think the deontic formulation here is less controversial, and allows us to remain neutral about *why* bringing new people into existence is not obligatory.

much more counterintuitive than denying the asymmetry outright. Nearby alternatives like average utilitarianism all deny the neutrality of additional happy lives.

What kind of impersonal view could give us that result? One could claim that even the best lives are at best neutral for their bearer, so that welfare never goes above zero (for instance, by claiming that welfare is determined only by amount of suffering). But this threatens an extreme anti-natalism, and is surely more implausible than rejecting the asymmetry. Alternatively, one could maintain a normal view of welfare and claim that nevertheless only the negative lives matter for the value of the world (for instance, by claiming that a world is better just in case it has less suffering-in-overall-bad-lives). But then we run into the unwanted implication that the distribution of well-being in lives is morally irrelevant provided the lives are overall good, since two worlds including some good lives are each no better or worse than a world without those lives. One could deny that the value of worlds is explained by welfare at all, but then we lose our explanation for the wrongness of ordinary harm, not to mention of bringing into existence a person with a miserable life.

This is not an exhaustive inventory of the options available to an impersonal theory. But it is very difficult to explain the asymmetry in impersonal terms without embracing an implausible view about welfare or an implausible view about the relation between welfare and the value of the world. The best strategy for the impersonalist, I think, is simply to bite the bullet and find some way to cast doubt on the intuitions behind the asymmetry. And this opens the door for a different kind of view to step in and claim territory. Because the complaints model does not need to appeal to a ranking of worlds in terms of betterness, we do not face the same obstacles as impersonal views, and that gives us an opportunity to make good on what has proven a troublesome implication to capture.

Contingent People

With that in mind, it becomes clearer what our view should say in two-option cases where the existence of some people is contingent on our action. One of the nice things about the complaints view is that it is very easy to explain the intuition of neutrality and the asymmetry. All we need is the following two, independently plausible claims about complaints:

No Ghostly Complaints: If A does not exist in W_x , then A has no complaint against X .

Existential Harm Complaints: If A exists only in W_x , then A has a complaint against X iff A's welfare is negative in W_x , and the strength of their complaint is the magnitude of their negative welfare.

Because there are no ghostly complaints, there will be nothing to be said against failing to bring into existence a person with a good life, all else equal. Bringing someone into existence whose life is worse than nothing, on the other hand, produces an existential harm complaint. This gives us the asymmetry.⁶

At this point, we have a complete story about two-option decisions. *Common Existence Complaints* tells us what complaints a person has when they exist in both alternatives, and *No Ghostly Complaints* and *Existential Harm Complaints* tell us what complaints they have when they exist in only one. Finally, *Minimize Aggregate Complaints* tells us how to translate the facts about these complaints into the permissibility of the options. Our view behaves just like utilitarianism in cases where people's existence is not at stake, and unlike utilitarianism when their existence is at stake, in precisely the right way to respect neutrality and the asymmetry. We can call this view the *Simple Complaints View* (SCV).

SCV faces some worries. I'll start by considering some objections to the way it treats two-option cases. Then, we will look at what I think is a more troubling puzzle – the question of how to generalize it to cases with more options.

The Non-Identity Problem

By far the most familiar objection to views like the one we've developed is the *non-identity problem*. In short, our view permits us to create worse-off people instead of different, better-off people, provided they all have lives worth living. In the interests of precluding, or perhaps hastening, the reader's disappointment, I will not have much to say about this objection. But it is worth understanding because it represents an important choice point in population ethics.

In his famous discussion in *Reasons and Persons*, Derek Parfit describes the following sort of case. Imagine a prospective mother (call her Mary) has a choice between conceiving a child now, which will lead to the birth of a child (call her Mariette) with a medical condition, who will go on to have a

⁶ It is worth noting that by accepting *Existential Harm Complaints* we are not committed to the idea that negative lives are worse for their bearer than nonexistence, which some hold is incoherent. These complaints may be grounded, for instance, in the fact that these lives are *noncomparatively* bad (See McMahan 2013).

mediocre life, or conceiving a child later, which will lead to the birth of a different, healthy child (call him Harry) who will go on to have a very good life. We can represent the case as follows:

	Mariette	Harry
Conceive Now	5	∅
Conceive Later	∅	10

Many people, including Parfit, have the intuition that it would be wrong for Mary to conceive now. But SCV implies that it is permissible. After all, nobody has a complaint if Mary conceives now – Mariette is as well off as she could be, and ghosts like Harry don’t get complaints.

Much ink has been spilled over this problem – about whether it really is an unacceptable consequence, and about whether a view like ours could make sense of it. My own take is that we ought ultimately to bite the bullet on the non-identity problem, for the sorts of reasons discussed in David Boonin’s book-length treatment (2014), as well as by Heyd (2009) and Roberts (1998). One of the lessons of the last few decades of work in population ethics is that every conceivable view will have some counterintuitive consequences,⁷ and I think the non-identity problem is the least bad of the bunch.

Still, if one finds this unpalatable, there are other options which won’t force us to abandon thinking in terms of complaints. For example, we might try to look for possible non-welfarist complaints that Mariette might have. In the spirit of Harman (2004), this might be a complaint against being brought into existence in an absolutely harmed state. In the spirit of Shiffrin (2012), it might be a complaint against being forced to suffer the predicaments of her life without her consent.

⁷ See for instance, the impossibility results shown by Arrhenius (2000) and the discussion in Greaves (2017).

These other options would require us to complicate our view significantly, so I will not build those suggestions in. But in case the reader is more concerned about the non-identity problem than I am, it is worth remembering they are available.

Tyrannical Complaints and Answers

But there is another problem, raised by Joe Horton (2021) which is not well appreciated. Unlike the non-identity problem, which suggests that views like SCV are too permissive, the challenge here is that they are not permissive enough.

Suppose we face the following choice: We can bring Amy into existence with a happy life, or we can bring Amy into existence with a very slightly less happy life, and also bring Bobby into existence with a fantastic life.

	Amy	Bobby
No Bobby	5	∅
Yes Bobby	4.99	10

SCV tells us that we must choose No Bobby, since nobody has a complaint against it, and Amy has a (small) complaint against Yes Bobby. But intuitively, this is the wrong result. Assuming we've made our peace with the non-identity problem, it is plausible that No Bobby is permissible – everyone who exists is as happy as they could be, after all. But Yes Bobby *also* seems permissible. If the reader is not convinced, consider that Amy's tiny complaint would, if SCV is correct, morally prohibit the creation of any number of other future people, at any level of welfare one can imagine. Horton calls this the problem of *tyrannical complaints*.

As I see it, the problem this reveals for SCV is that SCV gives complaints *too much veto power* over the creation of new people. This manifests not just with common-existence complaints but also with existential harm complaints. For example:

	Dolores	Bobby
No Bobby	∅	∅
Yes Bobby	-.01	10

Like Amy, Dolores has a tiny complaint against Yes Bobby. And there is no complaint against No Bobby. But as before, I think Yes Bobby should be permissible. Again, consider the extreme version of this kind of case. Suppose I can create an entire planet of beings, all of whom will be gloriously happy except for one person who will have a life very slightly not worth living. It does not seem that I must avoid creating this world. In some way, the gratitude of all the new lives should drown out the sole complaint.

One might think that this is unavoidable, given our commitment to the idea that creating new people is morally neutral. But I think there is a big difference between the force of the thought that the hypothetical well-being of purely contingent people doesn't *force* us to create them over others' complaints, and the thought that their well-being doesn't *allow* us to create them over others' complaints. It is the *obligatoriness*, not the *permissibility*, of the creation of new people in such circumstances that is particularly objectionable in standard impersonal views.

So I propose the following thought: if we do something which fails to bring into existence a contingent person, they don't get to reach out from another world and complain. But if we do something which does bring them into existence, we can use their gratitude to *answer* the complaints of others. An *existential* benefit to someone can be put on a balance against harms to others, much like ordinary benefits. The only difference is that the deprivation of an existential benefit has no victim, and so generates no complaint itself.

One can imagine the resulting view like this: at the end of time, we are all put on trial for the harms we have done, and each person who has a complaint against the way we treated them will testify against us. We can defend ourselves against the accusations either by showing that someone else would have had a complaint if we had chosen otherwise, or by calling those we have benefited by existence to testify on our behalf. If we can defend ourselves in this way against the full weight of the complaints leveled against us, we are in the moral clear.

To be more precise, we will introduce into SCV the idea of an *answer* according the following principle:

Existential Benefit Answers: If A exists only in W_X , then A generates an answer to complaints against X iff their welfare in W_X is positive, and the strength of this answer is the magnitude of their welfare.

Then we modify our deontic principle *Minimize Aggregate Complaints* to the following:

Minimize Aggregate Unanswered Complaints: An option X is permissible iff the total unanswered strength of complaints against X is no greater than the unanswered strength of complaints against any other option. The unanswered strength of complaints against an option is given by the strength of complaints against that option minus the strength of answers to those complaints (to a minimum of zero).⁸

The conjunction of *Minimize Aggregate Unanswered Complaints*, *Common Existence Complaints*, *Existential Harm Complaints*, and *Existential Benefit Answers* I will call the *Unanswered Complaints View* (UCV).⁹

UCV has some potential weaknesses relative to SCV. For one, the new normative dimension it introduces makes the view more complex.¹⁰ Secondly, some hold that it is desirable for a theory to *absolutely prioritize* existing or noncontingent people when distributing benefits, and so will take some of the implications I've been criticizing as a feature rather than a bug.

In addition, there is a concern that as a response to the problem of tyrannical complaints, UCV has erred too far on the side of permissibility. Consider the following case, which I owe to Jacob Ross:

⁸ One could have a more complicated view about how existential benefits can answer complaints. For example, one might think that the magnitude of an existential benefit must be quite a bit larger than a harm in order to adequately answer the complaint that harm provides.

⁹ An equivalent formulation of the view (which loses the explanatory structure in terms of complaints) is the following: UCV: Let $S(X)$ be the set of people who will exist if you do X. Add up the welfare of the members of $S(X)$ in W_X , and in W_Y – this tells us how good X and Y, respectively, are for $S(X)$. It is permissible to do X iff:

a) X is at least as good as Y for $S(X)$, or

b) X is worse than Y for $S(X)$ by no greater degree than Y is worse than X for $S(Y)$.

¹⁰ One might mitigate this worry by conceptualizing answers as just another kind of complaint – a complaint by happy people against worlds in which they don't exist. The limited role of these complaints is then explained by the fact that this kind of complaint can only be raised if it has already been respected.

	Alvin	Simon	Theodore
Create Simon	5	10	∅
Create Theodore	0	∅	5

UCV entails that either option is permissible in this case. Alvin has a complaint against creating Theodore, but Theodore has an answer which matches the complaint. But this is liable to strike many as implausible. Creating Simon is impersonally better, of course. But more to the point, creating Simon might seem like it should be better *even* from a complaints standpoint. After all, creating Simon leads to fewer complaints (since Alvin is better-off), and even more available answers (since Simon would have a better life than Theodore). The existential benefit to Theodore, the thought goes, shouldn't be able to answer Alvin's complaint, when the possibility of creating the even better-off Simon is on the table.

This worry becomes especially sharp when paired with the observation that answers are quite easy to come by. It is plausible that due to the sensitivity of personal identity to the conditions of conception, almost all of our actions will, by the propagation of small changes, have significant effects on the identity of people in the distant future. If this is true, then many future people will owe their existence to our choice, resulting in a huge weight of available answers to complaints no matter what we do.¹¹

There is a relatively simple revision we could make to UCV to avoid this problem. We could say that only *net* answers – that is, those that remain after we subtract the weight of answers favoring the alternative – are available to weigh against complaints. Since there are no *net* answers left over on the side of creating Theodore, Alvin's complaint would stand. And when each of our options would lead to the existence of a similar number of people in the distant future, the answers provided by those lives would on balance cancel out.

¹¹ I owe this worry to Joe Horton.

For my part, I am not convinced that this move is necessary. Within the complaints model, I find it helpful to think in terms of interpersonal justification. I imagine Alvin confronting me after I have created Theodore, on the grounds that I have made him worse-off than he would otherwise have been. Given what we've said about answers, Theodore can step in on my behalf and point out that if I had done otherwise, he would not have existed, and his life is good. The key question, then, is whether Theodore's defense can be undermined by pointing to the merely possible existence of Simon, who, *if* he existed, would also have a good life. If this is legitimate, then the net answers revision makes sense. But I think it is not clear that Simon's hypothetical gratitude can be invoked in this way. It seems to me sufficient to justify my action by reference to the interests of the people who exist, given that I've done it.

While I favor UCV, and will be taking it as our answer in two-option cases, most of the rest of this paper works equally well as a development of SCV, UCV, or its revision. I think something in this family is the most plausible treatment of such cases – more plausible than impersonal views, and more plausible than other person-affecting views that have been proposed. Whichever view of complaints we like, however, we do not yet have a complete theory. And we are about to hit a snag.

Part II: Many Options

UCV and its siblings work only as views about two-option cases because they don't provide a way to determine complaints from people who exist in at least two alternatives, when there are more than just two. *Common Existence Complaints*, recall, tells us that when there are exactly two options and someone exists in both, we measure their complaint against an option by comparing their welfare given that option to their welfare given the alternative. But if we add a third option, there is no unique alternative for us to compare to. So under what conditions does someone have a common existence complaint?

Initially, one might think the answer is surely obvious. In many-option cases, the natural thought goes, the weight of someone's complaint against an option is the difference between their welfare in that option and their welfare in the alternative that is *best* for them. This, which we might call *Measuring From the Top*, is the most straightforward way to generalize the idea behind *Common Existence Complaints*. Larry Temkin's (2011) view works essentially this way.

However, the resulting view is vulnerable to a serious problem introduced by Jacob Ross (2015). First, let's imagine a simple two-option case:

	Adam	Eve
X	10	10
Y	1	∅

Here, every view will agree on the obvious answer – X should be chosen. For UCV the explanation is that Adam has an (unanswered) complaint against Y, and nobody has a complaint against X. So far so good. But now suppose we add a third option:

Improvable Lives and Backfiring Complaints

	Adam	Eve
X	10	10
Y	1	∅
Z	-2000	1000

In this case, which we can call *Improvable Life*, we've introduced an option in which Eve is better off and given her, on the proposed view, a complaint against X. Moreover, now the complaints against X vastly outweigh the complaints against Y (Eve has no complaints against Y since she does not exist there), even if we allow existential benefits to provide a partial answer. So if we generalize common existence complaints by measuring from the top, we would be obligated to choose Y. But this can't be right. After all, Eve would be glad if we choose X over Y, since this would give her a good life rather than nothing. Adam would be glad if we chose X over Y, since he's better off. And there is nobody else whose interest is at stake. How can X be wrong?

It seems ridiculous to impose costs on Adam to avoid creating someone with a perfectly good life merely because that person's life could have been better in an alternative which is unthinkable horrible in other respects. For this reason, Ross calls this the problem of "improvable life avoidance".

It likewise seems ridiculous that a new complaint from Eve should force the agent to switch from bringing about an outcome in which Eve has a good life to bringing about an outcome in which she doesn't exist at all. Eve's new complaint seems to *backfire* against her. For this reason, Horton calls this the problem of "backfiring complaints".

These are two faces of the same issue. It should not be possible to start with a set of choices which permit us to create someone with a happy life, add an option under which they are better off, and thereby generate a complaint *on their behalf* which makes it impermissible to create them at all.

But if someone's complaints against an option aren't determined by how far their welfare under that option falls short of what is best for them, how are they determined? There are very few attractive candidates, given that we are limiting ourselves to facts about the welfare of the agent. We

could, for example, try measuring from a benchmark in the middle rather than the top – comparing against the agent’s average welfare in the options in which they exist, rather than the option in which they do the best. But a little calculation makes it clear this won’t help with our example. And there is a sort of arbitrariness about any particular choice of benchmark, even if it gets the right result in a given case. So, it looks like we’ve hit a dead end.

A Diagnosis

Here is what I think has gone wrong for our view. We have been assuming (implicitly, by appealing to principles like *Minimize Aggregate Complaints*) that there is such a thing as the complaint an agent has against an option, full stop. However, I want to claim, complaints are *essentially pairwise comparative*. That is, someone has a complaint against an option only *relative* to some other particular option. In two-option cases, there is only one other option to compare against. So we don’t run into any problems. But in three-option cases, there is no such thing as *the* complaint you have against X, or Y, or Z. There is only the complaint you have against X relative to Y, the complaint you have against X relative to Z, the complaint you have against Y relative to Z, and so on.

The natural way we tried to generalize our view to three options, *Measuring From the Top*, effectively said that the only relevant complaint for each person is the one relative to the option in which they do best. But this obscures important facts about the structure of their relative complaints which matter to what we ought to do. It matters, in *Improvable Life*, that Eve has a complaint against X *only relative to* (the horrible option) Z. The fact that this is not also a complaint against X relative to Y should in some way help keep X in the running, once we see how unacceptable Z is for other reasons.

The suggestion that complaints are pairwise comparative does not come completely out of thin air. It is natural to think that complaints have a close connection to *harms* (indeed, one might think that welfarist complaints of the sort we have been discussing correspond one-to-one with harms). And Alastair Norcross (2005) has forcefully argued that we should understand harms as pairwise comparative in exactly the way I’ve suggested for complaints. When we say an action or an event harms someone, Norcross suggests, context supplies a salient alternative relative to which the action or event is harmful.

A diagnosis is not a cure. So for the rest of the paper, I will explain how we can make good on the details of this idea. The strategy has three parts: First, we will reinterpret everything we said about complaints in two-option cases as claims about *relative complaints*. Second, we will reinterpret our deontic principle in two-option cases as a claim about *pairwise defeat*. Finally, we will come up with a new deontic principle which assigns permissibility as a function of the facts about pairwise defeat.

Relative Complaints and Pairwise Defeat

We are going to abandon all talk of unqualified complaints against an option in favor of talk about complaints against an option relative to another option. But the work that we did in the first part of the paper is not wasted. We thought at the time that we were building an incomplete account of complaints limited to two-option cases. But in fact, we were building a *complete* account of relative complaints.

All we have to do is take each of our complaint principles, remove any restrictions to two-option cases, and wherever “complaint against X” appears, replace it with “complaint against X relative to Y”. For example:

*Common Existence Complaints**: If X and Y are two options, and A exists in both W_X and W_Y , then A has a complaint against X relative to Y iff A is worse off in W_X than W_Y , and the strength of this complaint is the difference between their welfare in W_X and W_Y .

*No Ghostly Complaints**: If A does not exist in W_X , then A has no complaint against X relative to Y.

*Existential Harm Complaints**: If A exists in W_X but not W_Y , then A has a complaint against X relative to Y iff A’s welfare in W_X is negative, and the strength of their complaint is the magnitude of their negative welfare.


If we prefer UCV to SCV, as I do, then we do the same with our principle about answers:

*Existential Benefit Answers**: If A exists in W_X but not W_Y , then A generates an answer to complaints against X relative to Y iff their welfare in W_X is positive, and the strength of this answer is the magnitude of their welfare.

Now we have a story about relative complaints and relative answers, and we need to connect it to a deontic principle. We're going to do this by first introducing an intermediate, essentially pairwise notion of moral *defeat*. Our principle of defeat is going to be read almost straight off our earlier deontic principle in two-option cases:

*Minimize Aggregate Unanswered Complaints**: An option X defeats option Y iff the strength of unanswered complaints against X relative to Y is less than the strength of unanswered complaints against Y relative to X.

Since we have a complete story about complaints, we can use this to determine the defeat relations between each pair of options in any decision. And our new view will say, effectively, that X defeats Y iff X would have been obligatory according to UCV supposing X and Y were the only options. We can call this complete account of defeat *UCV-Defeat*. With such a view in hand, we can represent any decision we might face using a *defeat graph*. For example, here is the defeat graph for *Improvable Life*, where defeat is indicated by an arrow from the winner to the loser:




Tournaments

The defeat graph corresponds to the results of a tournament in which each option plays a match against each other option. The only thing that is left is for us to provide a way of picking winners of such tournaments, where the winners correspond to permissible options. In the example above, X is what is called a *Condorcet winner* – it defeats every other option. So, any reasonable principle will tell us X wins the tournament – exactly what is intuitive in this case. And any two-option decision is also very easy – an option will be permissible iff it is not defeated by the other option. But we won't always have a Condorcet winner, and our decisions may have an arbitrary number of options. How do we, in general, determine the *tournament solution* – the set of winning options?

There is an entire literature about various tournament solutions and their interesting properties, considered mostly in relation to problems of social choice (where a family of voting systems uses a similar tournament structure).¹² Going through all of the solutions discussed in that literature and assessing their consequences for our application is a large project beyond the scope of this paper.

Instead, I'm going to begin approaching this question in a very flatfooted way. There are a total of seven possible defeat graphs for decisions with three options. For some, what our tournament solution should say is obvious. For the others, we can look at a case with the relevant structure and see what seems correct in that case. Then we can pick a tournament solution which delivers the results we want.


First, some easy ones:



Interestingly, we can set this one aside completely if we accept *UCV-Defeat*, since cycles of defeat are impossible on that view. This is because one option can defeat another, on that view, only if it generates more total welfare.¹³ And there mathematically cannot be a cycle where each member generates more total welfare than the option before it. A principle of defeat based on SCV would allow such cycles. If we have a view which allows cycles, symmetry forces us to say either that every option is in the solution (i.e. is permissible) or that none of them are. If we are averse to moral dilemmas, the former is preferable, and I think this is what someone attracted to SCV should say. If we like *UCV-Defeat*, we can sidestep that choice.


¹² See Laslier (1997) for a discussion of tournament solutions and proofs of some of their properties.

¹³ If W_X has at least as much total welfare as W_Y , then every complaint against X relative to Y must be matched by complaints against Y relative to X, or answers in the form of happy people in W_X , in order to make up for the welfare shortfall associated with the complaint.




The second possibility is also easily dispensed with. Symmetry again demands that we say each option has the same status. But here there is no temptation whatsoever to think that no option is permissible – this is, after all, the defeat graph for a decision where no option has any effect on anyone’s welfare. So, obviously, every option should be in the solution.


The next two we’ll treat together:



In these, X is a Condorcet winner. We have already seen the second – it is the graph for *Improvable Life*, where intuitively X should be chosen. A case where X is strictly better than any other option for every possible person will also take one of these forms. Clearly, then, X should be alone in our solution here.



In this case, X and Y are symmetrical, and Z is defeated by everything. So, the only reasonable thing to say is that X and Y should be in our solution.




This is the only graph with exactly one defeat arrow. It is less immediately obvious what to say. But here is a case with the relevant structure:

	A	B
X	10	∅
Y	5	∅
Z	∅	10

I think it is clear that X and Z should be permissible, and Y, being strictly worse than X, should not. So this is what we'll want from our eventual solution.

Finally:



This one is the trickiest. Just visually, one might be tempted to say that X should be the unique winner. But consider the following case, which has the above structure:

	A	B
X	20	∅
Y	5	10
Z	∅	5

X and Z would both be permissible if they were the only two options (we chose to bite the bullet, recall, on the non-identity problem). If we say, then, that X is the only permissible option once we add Y, then the permissibility of creating B has been undermined by B having an improvable life, and the backfiring complaint this generates. This is exactly the sort of thing we wanted to avoid. So if we really want to stick to our motivations, we should not let X be the only permissible option.


Letting all three options be permissible, or just X and Z, runs into a different problem. That solution would violate the following plausible principle:

No Inferior Options: If W_Y is better than W_X for everyone who exists in W_X , and includes no new unhappy people, then X is impermissible.¹⁴

¹⁴ *No Inferior Options* also poses problems for a defeat principle based on SCV, because given such a principle it is possible for an inferior option to be in a defeat cycle, and it is natural to treat all options in a defeat cycle as permissible.

W_Y is better for B, and better for A, who exists and is happy there. If W_Y is preferred to W_Z by everyone who exists under either choice, then Z is in an important sense strictly *inferior* and should not be permissible. In light of all this, I think the best thing to say about this case is that X and Y, but not Z, are permissible.

To summarize, here are the three-option defeat graphs again, with the options I've argued should be in the solution filled in:



Now, what we want is a general tournament solution which will have these results in the three-option special case. There are countless possibilities with different implications for cases with four or more options, but here is what seems to me the simplest. Let us say that X *covers* Y when X defeats Y and anything that Y defeats. Then we say:

Uncovered: An option is permissible iff there is no option that covers it.

One can easily verify that *Uncovered*, which is one version of a commonly discussed tournament solution in the social choice literature,¹⁵ gets the results above for three-option cases.¹⁶ But better than that, with a little work we can prove that, paired with *UCV-Defeat*, it entails *No Inferior Options*, and even more interestingly, that it is guaranteed to avoid the problem of improvable life avoidance and backfiring complaints no matter how many options we add.

No Inferior Options

Let's begin with *No Inferior Options*. Suppose X is inferior to Y – that is, W_Y is better than W_X for everyone who exists in W_X , and has no new unhappy people. So, the people in W_X have unanswered complaints against X relative to Y and there are no unanswered complaints against Y relative to X. So, Y defeats X.

Suppose further that X defeats some other option Z. This means that the balance of unanswered complaints favors X over Z. We can now show that the balance of unanswered complaints favors Y over Z by at least as much. Consider any possible individual A. There are three ways the facts about A could favor X over Z. Taking each in turn, it follows straightforwardly from X's strict inferiority to Y that:

- 1) If A provides a common existence complaint against Z relative to X, then A provides a common existence complaint of greater strength against Z relative to Y.
- 2) If A provides an existential harm complaint against Z relative to X, then either A provides an existential harm complaint of equal strength against Z relative to Y (if A doesn't exist in W_Y) or A provides an existential harm complaint of greater strength against Z relative to Y (if A does exist in W_Y , where by supposition they are happy)
- 3) If A provides an answer to complaints against X relative to Z, then A provides an answer of greater strength to complaints against Y relative to Z.

¹⁵ I say one version, because the tournament solution social choice theorists call the *uncovered set* is typically defined only for what are called “strong” tournaments – tournaments where there are no ties – and there are several alternative definitions which are equivalent under that constraint. What we have here corresponds to one of those definitions applied to tournaments allowing ties. For discussion of various definitions of the uncovered set, see Penn (2006) and Duggan (2013).

¹⁶ One other tournament solution which is often discussed in the social choice literature and which gets the results we desire for three-option cases is the *Banks Set* – the set of all winners of maximal transitive sub-tournaments. The Banks Set is a subset of the Uncovered Set, but they come apart in tournaments with more than three players. Because I find it difficult to have a clear intuition about the complicated cases in which they differ, I have nothing special to say in favor of one rather than the other. I go with the Uncovered solution here because it is simpler to understand.

There are also three corresponding ways facts about A could favor Z over Y:

- 4) If A provides a common existence complaint against Y relative to Z, then either A provides a common existence complaint of greater strength against X relative to Z (If A exists in W_x), or A provides an answer to complaints against Z relative to X of greater strength (if A doesn't exist in W_x).
- 5) If A provides an existential harm complaint against Y relative to Z, then A provides an existential harm complaint of greater strength against X relative to Z.
- 6) If A provides an answer to complaints against Z relative to Y, then A provides an answer of equal strength to complaints against Z relative to X.

So any possible person who weighs in X's favor against Z on the scale of complaints and answers weighs at least as much in Y's favor against Z, and any possible person who weighs in Z's favor against Y weighs at least as much in Z's favor against X. Since, by supposition, the balance favored X over Z, it must also favor Y over Z, and so Y defeats Z as well.

Since, Y defeats X and defeats anything that X defeats, Y covers X. By our tournament solution, X is impermissible, and *No Inferior Options* is validated.¹⁷

Backfiring Complaints and Improvable Lives Revisited

The problem with backfiring complaints and improvable life avoidance, recall, was this: we could start with a case where it was permissible to create someone with a happy life, add an option under which they would be better off, and in virtue of their new complaint make it impermissible to create them at all. We can show that *Uncovered* rules this out.

Suppose X is permissible. By our view, there are no options that cover X. Now suppose we add a new option Y. There are two possibilities. If Y does *not* cover X, then X is still not covered – the addition of any new option cannot make existing options which did not cover X cover it. So X remains permissible.

¹⁷ This result depends both on our choice of tournament solution and on our view about defeat. Notably, if one is moved by the worries in Part I to accept a revision of *UCV-Defeat* which looks only at net answers, that will work as well, since the resulting view is strictly less permissive than ours. However, pairing *Uncovered* with *SCV-Defeat* will not (See footnote 14).

If Y *does* cover X, then Y must be permissible. This is because otherwise Y itself would have to be covered. But anything that covers Y *also* covers X, since the covering relation is transitive. And this third option would already have covered X before the introduction of Y, and X could not have been permissible. This proves an important property of *Uncovered*:

Losers Can't Dislodge Winners: If X is permissible, then the addition of Y can make X impermissible only if Y is itself permissible.

This means, as a special case, that if creating A is permissible, then there is no way to add options in which A exists such that creating A is impermissible. At most, A's complaint relative to the option in which they are better off will make *that* option permissible instead. So given our story about defeat and our tournament solution, there *cannot* be cases of backfiring complaints, with any number of options, and as a corollary, the fact that a life is improvable never obliges us not to bring it about.

Ross's View and the Hybrid Strategy

It will be useful here to contrast our view with one proposed by Jacob Ross (2015). Ross endorses a hybrid view – he thinks morality ought to acknowledge a dimension corresponding to impersonal value (for instance, total utility) as well as one corresponding to person-affecting value (or complaints in our terminology). And although he does not put it in these terms, the person-affecting dimension has precisely a defeat/tournament structure.

Ultimately, I think Ross's view is less attractive than the one we develop here. I want to begin by looking at the person-affecting component of his view alone, which we can see as a competitor to our account of complaints and their resolution. Then we will consider the virtues and vices of going hybrid with an impersonal view. Ross proposes the following principle of defeat:¹⁸

Differential Utility Principle of Defeat (DUP-Defeat): There is one important evaluative dimension such that, on this dimension, a first outcome defeats a second if and only if, among those who exist in both outcomes, the total utility is greater in the first outcome than in the second.

¹⁸ Strictly speaking, Ross describes but does not go so far as to commit himself to this principle. Rather, it is the most natural generalization of a more narrow principle that he does endorse.

Although it is not framed in terms of complaints, *DUP-Defeat* essentially amounts to the claim that an outcome defeats another if it minimizes *common existence* complaints as we have characterized them. He pairs this view of defeat with the following tournament solution:

Insubordinate: Avoid options that are *subordinate* to other options, where X is subordinate to Y iff Y defeats X directly or indirectly (by a chain of successive defeat), but not vice versa.


Ross shows that the resulting view gets the desirable result in *Improvable Life*. But although I think the structure of the view reflects an important and correct insight about the strategy theorists in the person-affecting tradition should take, there are a number of important things it does not explain which our view does.

First, common-existence complaints are not sufficient to explain the procreation asymmetry. Indeed, they are not even sufficient to explain the most uncontroversial part of the asymmetry – that it is impermissible, *ceteris paribus*, to create lives that are worse than nonexistence. In this way Ross’s person-affecting component explains less even than *SCV*. Because it doesn’t include anything like answers, moreover, it does not avoid the problem of tyrannical complaints.

A second issue is that while it gets the right result in some specific cases like *Improvable Life*, it does not solve the underlying problem, and there are related cases we have already seen where it has implications that are objectionable for similar reasons. Recall the following case:

	A	B
X	20	∅
Y	5	10
Z	∅	5

Ross's view would give this case the same defeat graph as *SCV-Defeat* and *UCV-Defeat*:



But while *Uncovered* tells us X and Y are both permissible, *Insubordinate* tells us that picking X is obligatory. As we argued earlier, this result would make the case another instance of a backfiring complaint, where the fact that B's life is improvable forces us not to create them. On *Insubordinate*, losers *can* dislodge winners. Avoiding the counterintuitive implication in *Improvable Life* is one of Ross's explicit motivations in choosing his defeat principle and tournament solution, so it is particularly worrying that a version of the problem still remains.

If the person-affecting component exhausted the view, these would all be very serious disadvantages. But Ross does not intend the view to exhaust morality - he holds that there are also reasons given by impersonal value, and these reasons might help plug the explanatory holes. But I think this strategy is not very appealing for several reasons.

First, it can at best fill only *some* of the gaps left by the person-affecting component of the view. It can, for example, explain why creating miserable lives is wrong – doing so lowers the value of the world by lowering total utility. Even here, though, it is not obvious that it gives the *right* explanation. It seems plausible that when something is wrong because of impersonal reasons alone, it is wrong without *wronging* anyone – the idea of a *victim* is connected to whatever sphere of morality is concerned with obligations to individuals, where person-affecting reasons reside. But when we create someone with a life that is worse than nonexistence, not only have we done something wrong, but we have *wronged that person* – we have done something against which they have a legitimate grievance on their own behalf. In other words, it is precisely the sort of thing that should be on the complaints side of a hybrid view.

It might look like the hybrid view helps with the problem of tyrannical complaints as well. Total utility will be higher if we create happy people, and this might outweigh complaints that would otherwise be tyrannical, robbing them of their veto power. But even then it is not clear that the appeal to impersonal value gets the intuitive result, which is that creating happy people over someone's small complaint is *permissible but not obligatory*. That will depend on how the person-affecting and impersonal dimensions interact, and it is difficult to see what plausible view of this will generate this result.

This is connected to what seems to me the biggest lacuna left by the hybrid strategy, which is that it cannot explain the thing that most motivated us – the procreation asymmetry. We already saw that impersonal views are in a poor position with respect to the asymmetry, and attaching an impersonal view to a view which only acknowledges common existence complaints does not help at all.

So adding an impersonal component is at best a mixed bag when it comes to closing the explanatory gap between Ross's account of a person-affecting dimension of evaluation and ours. On top of this, I think there are additional grounds not to like a hybrid view of welfarist reasons. One drawback comes from very general considerations of theoretical unity. It is not just that the hybrid view is pluralist. It is that impersonal reasons and person-affecting reasons seem to represent quite different overarching views of what morality is all about – whether it is about our relation to value in the world, or a system of interpersonal obligations. In this way it is less unified than, for instance, a wholly impersonal view which recognizes more than one way a world might be better than another, or a wholly interpersonal view which recognizes more than one irreducible kind of obligation.

There are also concerns about how the different moral dimensions interact to determine what we ought to do, all-things-considered. If the person-affecting welfarist reasons favor one option and the impersonal welfarist reasons favor another, what determines their relative weight (if indeed, weighing is even the way to understand their interaction)? Again, in part this is a general problem for pluralist views, but I think it is especially pressing here because all of the relevant reasons are welfarist. It is a special disappointment to find that even once we've assigned numbers to every individual's welfare in every option, and we fully understand how welfare matters to the impersonal and person-affecting components of the view individually, it remains a mystery what we ought to do.

Moreover, it is not at all straightforward to fill in a story about how to weigh these dimensions of value. Even a simple and natural claim like “the weight of welfarist reasons favoring an option are

the sum of the weight of impersonal reasons and the weight of person-affecting reasons” looks like it won’t be right. Suppose I face a choice between inflicting a horrible life on an existing person, or creating a new person and inflicting that life upon them, such that the welfare table is as follows:

	Aldo	Nova
Torture Existing	-100	∅
Torture New	0	-100

On a plausible impersonal view like classical utilitarianism, the impersonal reasons against each option are tied. But on Ross’s view, the first option, and only that option, has strong person-affecting reasons against it as well. So, if the simple and natural claim were true, the first option would be significantly more objectionable than the second. But intuitively it is not. It is easy to vindicate this intuition on a pure complaints view like ours, or on a purely impersonal view, but it is not easy on the hybrid view.

It would not be fair to claim that there are no advantages of the hybrid approach over our pure complaints picture. Notably, a hybrid view can generate the intuitive judgments about the non-identity problem, which we have not even tried to do. In part, whether one likes a hybrid view will depend on how much one values theoretical virtues like simplicity. But we’ve seen that a pure complaints model can do much of the work Ross requires of the impersonal component of his view, and even if we ultimately accept a hybrid view, there are reasons to prefer our account of the person-affecting side of the picture.

Conclusion

I’ve argued that *UCV-Defeat*, together with *Uncovered*, captures the procreation asymmetry, avoids the problem of tyrannical complaints, and escapes the worries introduced by additional options in a way that we can prove will not recur. I think the result, perhaps with some further tinkering, is not only the most plausible version of the complaints model, but the most plausible view in the tradition of person-affecting approaches to population ethics. Even if one doesn’t like various aspects of the view, however, I think the *strategy* that we used to get ourselves there is a very promising tool to keep in our pockets.

The tournament approach is very flexible – if you don’t like my theory of complaints, you can plug in any view about complaints you want, and keep the rest of the structure and much of its attractiveness. Or, if you don’t like complaints at all, you can come up with a principle of defeat between options which doesn’t appeal to complaints, and then attach whichever tournament solution has the properties that are most convenient. I’ve tried to show how some of the properties of tournament solutions can place significant constraints on the kinds of counterexamples that can be developed against a view, and we did not even scratch the surface of the work that has been done by mathematicians and social choice theorists studying tournaments.

It is particularly useful when we find ourselves in a position where we are confident in how we want to treat two-option cases, but find it difficult, for whatever reason, to generalize. In our case, I suggested a particular principled explanation for why two-option cases were so much more tractable – namely, that complaints themselves were essentially pairwise comparative. But the trick of taking a view about two-option cases, translating it into a view about pairwise defeat, and then combining it with a tournament solution, is one you could pull off in any normative domain, and in response to any roadblock making generalization troublesome.¹⁹

So there are two things I hope to have commended to the reader: a plausible substantive view about population ethics, and a philosophical maneuver with potential for application across normative philosophy.

¹⁹ I’ve argued elsewhere that a similar strategy can make good on an attractive idea about decision theory which otherwise resists generalization to many-option cases (Podgorski 2020). And Tarsney (2019) uses the tournament idea as a decision rule under conditions of normative uncertainty.

References

- Arrhenius, G. (2000). An Impossibility Theorem for Welfarist Axiologies. *Economics and Philosophy* 16, 247-266. doi: doi.org/10.1017/S0266267100000249
- Boonin, D. (2014). *The Non-Identity Problem and the Ethics of Future People*. Oxford: Oxford University Press.
- Broome, J. (2004). *Weighing Lives*. Oxford: Oxford University Press.
- Duggan, J. (2013). Uncovered Sets. *Social Choice and Welfare* 41, 489-535.
doi: doi.org/10.1007/s00355-012-0696-9
- Greaves, H. (2017). Population Axiology. *Philosophy Compass* 12. doi: doi.org/10.1111/phc3.12442
- Harman, E. (2004). Can We Harm and Benefit in Creating? *Philosophical Perspectives* 18, 89-113.
doi: 10.1111/j.1520-8583.2004.00022.x
- Heyd, D. (2009). The Intractability of the Non-Identity Problem. In Roberts, M. and Wasserman D. (Eds) *Harming Future Persons*, pp. 3-25. Dordrecht: Springer.
- Horton, J. (2017). Aggregation, Complaints, and Risk. *Philosophy and Public Affairs* 45, 54-81.
doi: doi.org/10.1111/papa.12084
- Horton, J. (2018). Always Aggregate. *Philosophy and Public Affairs* 46, 160-164.
doi: doi.org/10.1111/papa.12116
- Horton, J. (2021). New and Improvable Lives. *Journal of Philosophy* 118, 486-503.
doi: doi.org/10.5840/jphil2021118934
- Laslier, J. (1997). *Tournament Solutions and Majority Voting*. Berlin: Springer-Verlag Berlin Heidelberg.
- McMahan, J. (2013). Causing People to Exist and Saving People's Lives. *The Journal of Ethics* 17, 5-35. doi: 10.1007/s10892-012-9139-1
- Narveson, J. (1978). Moral Problems of Population. *The Monist* 57, 62-86. doi: 10.5840/monist197357134

Norcross, A. (2005). Harming in Context. *Philosophical Studies* 123, 149–173. doi: 10.1007/s11098-004-5220-3

Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.

Penn, E. M. (2006). Alternate Definitions of the Uncovered Set and their Implications. *Social Choice and Welfare* 27, 83-87. doi: 10.1007/S00355-006-0114-2

Podgorski, A. (2020). Tournament Decision Theory. *Nous*. Advanced online publication. doi: doi.org/10.1111/nous.12353

Roberts, M. (2003). Can it ever be better never to have existed at all? Person-based consequentialism and a new repugnant conclusion. *Journal of Applied Philosophy* 20, 159–185. doi: doi.org/10.1111/1468-5930.00244

Roberts, M. (1998). *Child versus Childmaker*. Lanham: Rowman & Littlefield.

Ross, J. (2015). Rethinking the Person-Affecting Principle. *Journal of Moral Philosophy* 12, 428-461. doi: 10.1163/17455243-01204004

Shiffrin, S. (2012). Harm and its Moral Significance. *Legal Theory* 18, 357-398. doi: doi.org/10.1017/S1352325212000080

Tarsney, C. (2019). Normative Uncertainty and Social Choice. *Mind* 128, 1285-1308. doi: doi.org/10.1093/mind/fzy051

Temkin, L. (2011). *Rethinking the Good*. Oxford: Oxford University Press.