

Consequentialism and the

Agent's Point of View

Nathan Robert Howard --- Texas A&M University --- nathanrhoward@gmail.com

Forthcoming in *Ethics*; please cite published version

Abstract: I propose and defend a novel view called ‘*de se* consequentialism’, which is noteworthy for two reasons. First, it demonstrates — *contra* Doug Portmore, Mark Schroeder, Campbell Brown, and Michael Smith, among others — that a consequentialist theory employing agent-neutral value is logically consistent with agent-centered constraints. Second, *de se* consequentialism clarifies both the nature of agent-centered constraints and why philosophers have found them puzzling, thereby meriting attention from even dedicated non-consequentialists. Scrutiny reveals that moral theories in general, whether consequentialist or not, incorporate constraints by assessing states in a first-personal guise. Consequently, it is no coincidence that *de se* consequentialism mimics constraints: its distinctive feature is the very feature through which non-consequentialist theories enact them.

Word Count: 13 861, including all notes and references

Keywords: Consequentializing, Agent-Neutral Value, Agent-Relative Value, Subjective Consequentialism, the *De Se*

Consequentialism and the

Agent's Point of View

Forthcoming in *Ethics*; please cite published version

1. Introduction

Budding philosophers learn early that act consequentialism (hereafter, just *consequentialism*) has surprising implications. For example, its defining claim that, roughly, you may always do whatever produces the best outcome appears to permit killing an innocent person if doing so will save others. However, ‘New School’ consequentialists argue that surprising implications like permissions to kill the innocent follow only from optional interpretations of the basic consequentialist framework. According to them, consequentialism can mimic, for example, a so-called ‘agent-centered’ *constraint* that forbids a killing no matter its outcome, forbidding killing someone even to prevent another from killing that very same person and several more. If New Schoolers are right, very little about consequentialism’s reliance on *good* shapes theorizing since it is consistent with agent-centered constraints on acts. If that’s so, then we might understand all of morality through an evaluative lens.¹

But the New School view comes with a catch. Constraints appear logically incompatible with consequentialism unless we understand value in a particular way. New Schoolers assume that, “in order to [mimic constraints], [they] must deny that agents ought always to produce the most good and accept instead that each agent ought always to produce the most of what is good-relative-to-her”.² This

¹ See Dreier (1993), Louise (2003), and Peterson (2010) for remarks in this vein. For further discussion, see McNaughton & Rawling (1992), Brown (2011), Hurley (2013), and Schroeder (2017).

² Portmore (2007: 13). See also Portmore (2009: 330): “Of course, the rankings of outcomes will need to be agent-relative (and, perhaps, also time-relative) in order to consequentialize certain non-consequentialist theories.” Brown (2011) offers a similar thought: “Some views of right and wrong are incompatible with [agent-neutral] consequentialism, no matter what theory of the good one adopts.”

commitment is costly. Critics allege that consequentialism's compelling claim that you may always produce the *best* outcome cannot mean anything like 'best-outcome-relative-to-you' since the latter is agent-relative but the former is not — it is 'agent-neutral'. Critics insist that New School consequentialism is, therefore, not recognizably consequentialist.

As I argue below, the New Schoolers' assumption is false. Agent-neutral consequentialism and constraints are consistent. Standard-issue *subjective* consequentialism, which holds, roughly, that an agent must do what will have the best *expected* outcome, is consistent with constraints. Expected outcomes reconcile consequentialism and constraints because they are individuated by the agent's expectations or, more broadly, by her 'point of view'. Expected outcomes are thus already agent-relative, in a sense explained below. As a result, mimicking constraints does not require both agent-relative conceptions of the good and expected outcomes. Expected outcomes alone suffice.

This paper has two main aims. The first is to show that agent-neutral consequentialism is compatible with constraints. It turns out that the best way to achieve this aim is to hold that two outcomes that represent the same event can differ in value. In particular, consequentialists who wish to mimic a constraint on killing should hold that, roughly, presenting a killing in the *de se* or first-personal mode makes it worse.³ The second aim is to show that this position is coherent.

The proposed view merits attention from even dedicated non-consequentialists for two reasons. First, it illuminates how constraints are enacted in a moral theory. Second it shows how some normative concepts, like *good*, are more plastic than is often assumed. As a result, the common method of dividing ethical theories by their privileged normative concept — that is, by whether they understand rightness through rules, reasons, rationality, value, etc. — does not carve at the deepest joint in ethics. Rather, a more fundamental fault line lies between theories that take the 'point of view of the universe' and those that take the agent's point of view.

³ Perhaps, as Louise (2003) and Smith (2009) observe, temporal considerations force a parallel treatment of *de nunc* outcomes. However, what I say for one can be extended to the other, so I bracket the issue.

2. The Standard Recipe

Consequentialists characteristically claim that, roughly, you may always do what produces the best outcome. No action is forbidden independently of its consequences. This claim conflicts with the idea that some actions, like killing, are intrinsically forbidden. I'll call a rule that forbids you from acting some way regardless of the consequences a *constraint*.⁴

To illustrate, suppose that a constraint on killing exists and that, as a result, you are never permitted to kill, even in cases like the following:⁵

Mafia: The Mafia are certain to kill five innocent strangers unless you kill a sixth, which is certain to stop the Mafia from killing the five.

The constraint implies that you may not kill even as the only means to prevent more killing. Consequentialism, by contrast, appears to permit you to kill the sixth if doing so produces a better outcome. Many find this permission heinous. As a result, there is considerable pressure on consequentialists to look for interpretations of their view that forbid agents from killing in cases like *Mafia*. New Schoolers allege that a 'standard recipe' offers just this. It has two ingredients: a 'broadened' conception of outcomes and an 'agent-relative' notion of value.

Consequentialists often hold that an act's outcome is the bounded sum of its downstream effects.⁶ New Schoolers, in contrast, conceive of outcomes more broadly. 'Broad outcomes' include not

⁴ I'm using 'constraint' to mean what Samuel Scheffler means by 'agent-centered restriction'. He writes that "an agent-centered restriction is, roughly, a restriction which it is at least sometimes impermissible to violate in circumstances where a violation would serve to minimize total overall violations of the very same restriction, and would have no other morally relevant consequences" (1985: 409).

⁵ I'm assuming that moral dilemmas where, for example, you should kill and you should not kill, are impossible. Suikkanen (2020) argues that moral dilemmas pose a special challenge to accommodating constraints in a consequentialist theory. These arguments are orthogonal to my aims here.

⁶ If outcomes extend across the future without limit, knowing whether an act is permissible implausibly requires knowing elusive facts about the distant future — see Lenman (2000) and Mason (2004) for more. I assume that the outcome of, for example, refraining to kill someone can include the killing of two by others.

only the act's downstream effects but also the act itself. For example, when I kill, the broad outcome involves not only the fact that someone has died but also the fact that I've killed them.⁷

Broadened outcomes do not, by themselves, reconcile consequentialism and constraints. Schroeder (2006) describes a case where each of Hans, Franz, and Jens must choose whether to kill in order to prevent the other two from killing. I'll represent the outcome where Hans kills as 'H', and similarly for F and J. We can derive a contradiction from this setup if we assume that killing an innocent person is always bad.⁸ This assumption implies that for any X and $Y \in \{H, J, F\}$ such that $X \neq Y$, $X > X \& Y$, where '>' expresses agent-neutral 'is better than' and ' $X \& Y$ ' names the outcome where X and Y kill.⁹ To mimic a constraint on killing, consequentialists must claim that, from Hans's perspective, $J \& F > H$. That's because, according to them, Hans ought not kill only if the outcome of his killing is worse than the outcome of both Jens and Franz killing. But the parallel fact from Jens's perspective, $H \& F > J$, proves a contradiction.¹⁰

Therefore, unless New Schoolers also employ 'agent-relative' value, where an outcome has a value only relative to an agent, their view is inconsistent under plausible assumptions. I'll represent their asymmetric, transitive, irreflexive agent-relative better-than relation using ' $>_X$ '. For example, rather than claiming that Hans's perspective implies that $J \& F > H$ and that Jens's perspective implies $H \& F > J$, New Schoolers instead claim that the agents' differing perspectives imply different but consistent rankings: $J \& F >_H H$ and $H \& F >_J J$.

A central challenge for New Schoolers is interpreting the concept expressed by ' $>_X$ '. It can be taken to mean *better-for- X* , such that ' $J \& F >_H H$ ' means that outcomes where Jens and Franz each kill someone are better *for Hans* than ones where Hans kills someone. The notion of something being better for someone than an alternative is simple and familiar. Eating two bowls of ice cream is better for

⁷ For suspicion about this idea, see Hurley (2017; 2019a; 2019b).

⁸ I'll drop the explicit 'innocent' qualification going forward and implicitly assume that all killings discussed below are killings of innocent people.

⁹ '&' does not express conjunction and is syncategorematic, saving me from inserting parentheses throughout.

¹⁰ Given Hans's perspective, $J \& F > H$. But $H > H \& F$ by the assumption. So $J \& F > H \& F$ by transitivity. Likewise, given Jens's perspective, $H \& F > J$. But $J > J \& F$ by the assumption. So $H \& F > J \& F$ by transitivity. However, $>$ is anti-symmetric, resulting in contradiction.

me than eating three. However, a moral theory that requires each to promote what's good for them seems like a kind of teleological egoism, not consequentialism.¹¹ As Schroeder claims, "ethical egoists do *not* believe that it is always permissible for you to bring about the most good. They believe that it is always permissible for you to bring about the most of what is good *for you*. But since *good* and *good for* express different concepts, the idea that egoists find compelling turns out not to be at all the same idea that consequentialists find [c]ompelling."¹² Using *better-for-X* appears to forbid New Schoolers from drawing on intuitive support for consequentialism.

New Schoolers can instead claim that ' $>_X$ ' expresses a different concept, *better-relative-to-X*.¹³ However, critics argue that agents lack a pre-theoretic grasp of *better-relative-to-X*, so that talk of *better-relative-to-X* makes as little sense as talking about *orange-relative-to-X*.¹⁴ Or it can be interpreted in terms of what one should prefer, such that $A >_X B$ just when and because the reasons for X to desire that A are stronger than the reasons for X to desire that B .¹⁵ But it's unclear whether I always have most reason to prefer what's best.¹⁶

Perhaps independent reasons favour embracing a broadly agent-relative notion of value. I don't know. My thesis is simply that consequentialists don't need it to mimic constraints. So I'll momentarily bracket discussion of agent-relative value to focus on a novel, alternative approach.

¹¹ Portmore (2007) disputes this claim.

¹² Schroeder (2006: 354). Of course, not everyone finds consequentialism compelling — see Hurley (2017).

¹³ Discussions of this concept include Hurka (2001), Smith (2003), Portmore (2007), Suikkanen (2008), and Dreier (2011). For arguments that *some* axiology must mimic constraints, see Peterson (2010).

¹⁴ Schroeder (2008).

¹⁵ See Smith (2009) and Portmore (2011) for views defended along these lines.

¹⁶ [redacted]

3. *De Se* Consequentialism

3.1 Presentational Differences Can Be Moral Differences

Objective consequentialists characteristically evaluate, roughly, an action's future effects.¹⁷ Whether you should press the button depends only on what happens once you do. This view has counterintuitive implications. Even if all of your evidence indicates that pressing the button will result in an otherwise avoidable fiery catastrophe and even if there is simply *no* reason to suspect that your pressing the button has *any* moral upside, you should still press the button if doing so provides the best possible outcome. This implication results from evaluating only the actual effects of an agent's action, which obscures the potentially morally significant effects of how the action appears from the agent's point of view, such as whether it is risky or reckless.

As a result, other forms of consequentialism attend to an act's intended or foreseen consequences, rather than to its actual consequences; they take the agent's point of view. Very roughly, what matters for some *subjective* consequentialists are actions' *expected* effects not their actual effects. Crucially, this shifts from thinking of outcomes extensionally, as actual future states, to thinking of them from the agent's point of view, as objects of attitudes like expectation, belief, credence, anticipation, etc. As others, particularly Mahtani (2017) note, this grants subjective consequentialism unexpected nuance.

To illustrate:

Lois Lane. Lois Lane has just learned the bad news: aliens are about to destroy the world. Worse still, her cell phone has enough battery to make only one call. She can either call her friend Clark Kent to come save the day or she can call her friend Superman. She knows Clark Kent as the nerdy guy in the office. She knows Superman as someone who's faster than a speeding

¹⁷ I'm using 'subjective' and 'objective' consequentialism in broadly the manner of Railton (1984: 152). I'm not the first to exploit the distinctive features of probability and expectation to defend consequentialism. Elements of my approach are anticipated by Jackson (1991) and Oddie and Milne (1991).

bullet, more powerful than a locomotive, and able to leap tall buildings in a single bound. Unbeknownst to Lois, Clark Kent is numerically identical to Superman.

By adding some numerical values to this scenario representing Lois’s evaluations and expectations, the table below describes the features of Lois’s choice relevant to some subjective consequentialists:¹⁸

Expected Effects ▷ ----- Actions▽	Superman arrives to save the day.	Clark arrives to save the day.
Lois dials Superman’s number	(1) 10 utility at 99% likelihood = 9.9 value.	(3) 1 utility at 1% likelihood = .01 value.
Lois dials Clark Kent’s number	(2) 10 utility at 1% likelihood = .1 value.	(4) 1 utility at 99% likelihood = .99 value.

From Lois’s point of view, while each friend is equally likely to answer the call, having Superman arrive to save the day is better than having Clark Kent arrive. What should Lois do? Since no matter who Lois calls, the same person is very likely to show up, *objective* consequentialism implies that Lois can call either. But that’s counterintuitive. Given that Lois thinks of Clark as a mere office worker, calling Clark would be a mistake when she could call someone who’s more powerful than a locomotive. *Subjective* consequentialism ratifies this intuition. Since by Lois’s lights Clark is a desk jockey and Superman is faster than a speeding bullet, it’s better to call Superman than to call Clark.¹⁹

Clark arriving to save the day is metaphysically identical to Superman arriving. But they are distinct outcomes for Lois because she regards them as distinct. These outcomes differ only in how a particular presents itself to Lois — as an office worker in one case and as able to leap tall buildings in a single bound in the other. Nevertheless, these merely *presentational differences* — differences in how a particular presents itself to an agent — clearly have moral significance for many subjective

¹⁸ I’m using ‘utility’ as a catch-all term for the features of the possible state-of-affairs represented by an outcome that contribute to the outcome’s value. An outcome’s value is (partly) a function of its expected utility.

¹⁹ Like Mahtani, I reject “naive referentialism about credence”. On the view that we favour, it’s possible to have distinct credal attitudes towards distinct epistemic possibilities that are intrinsically metaphysically identical. See also Chalmers (2011: 590) and Fitts (2013).

consequentialists since they explain why Lois should call Superman and not Clark Kent.

A second presentational difference is particularly significant. Just as Clark and Superman are identical, the referent of ‘I’, in my use, and the referent of ‘the actual author of this article’ (@AU) are identical. And just as (merely epistemic) possibilities involving Superman can misleadingly present themselves as involving only an office worker and not a superhuman alien, possibilities involving @AU can misleadingly present themselves as involving only her and not me. For example, I can entertain what thrilling turns my life might have taken had I not been the actual author of this article. Likewise, I can imagine someone else, apparently distinct from me, writing it.

Now, if the presentational differences between Lois’s Superman possibilities and her Clark Kent possibilities can ground evaluative differences, we cannot be surprised when the presentational difference between me myself and someone else (who happens to be me) also makes an evaluative difference. To illustrate, suppose that the local grocery store has selected you and several others to win a month’s worth of groceries. The prize merely compounds your luck: you are extremely well off, having benefited from a generous inheritance. You line up in one of the two queues formed to distribute the prizes. As you progress, only you and a woman in the adjacent queue remain. You are then informed that, owing to a clerical error, only one prize is left. You wonder whether you or the woman would benefit more from the prize. Because you are likely better off than her, you reason that she would probably benefit more from receiving the prize. So you judge it better that she receives it and that you should, therefore, decline the prize. But, twist of fate, you failed to recognize yourself in the mirror: you were her all along.

Prior to recognizing that you are the woman next to you, the kind of subjective consequentialism that we are considering implies that the outcome where you receive the prize has less value than the outcome where the woman receives it. These two outcomes are distinguished only presentationally. In one, you are “presented to [your]self in a special and primitive way, in which [you are] presented to no one else”, the so-called *de se*.²⁰ In the other, you are presented to yourself third-personally in the impersonal mode: you are presented as a stranger to yourself. Outcomes that

²⁰ Frege (1956: 298).

involve *de se* thought — “*de se* outcomes” — can therefore differ in value from metaphysically identical outcomes presented differently.

Summing up, the standard recipe represents an agent’s point of view in *her* distinctive evaluative standards or in the reasons that *she* has to prefer certain outcomes to others. In contrast, some forms of subjective consequentialism represent an agent’s point of view in how the potential effects of her actions present themselves to her, not in how she evaluates them. These presentational differences can ground evaluative differences between otherwise metaphysically identical outcomes. Presentational differences, as we will see, reconcile agent-neutral consequentialism and constraints.

3.2 Private Outcomes

A view that I will now develop, called *de se* consequentialism, employs *de se* outcomes to reconcile constraints and agent-neutral consequentialism, demonstrating that they are consistent, *contra* Schroeder, Smith, Portmore, Brown, etc. It requires two assumptions. First, it requires assuming that some outcomes involve *de se* thought, following discussion above. For example, when I drink in order to quench my thirst, the outcome that I intend by drinking — namely, that *I myself* quench my thirst — is necessarily *de se* or essentially indexical. Likewise, when deciding whether you will get the prize or whether the woman in the other line (who happens to be you) will get it, the first outcome is necessarily presented in the *de se* mode since that’s what distinguishes it from the second.²¹

The second assumption is also defended in the previous subsection: that *de se* thoughts are private. In *de se* thought, you are “presented to [your]self in a special and primitive way, in which [you

²¹ *De se* consequentialists therefore follow many philosophers in assuming that outcomes figure in the psychological explanation of action and that, as Bermudez (2017: 690) puts it, “psychological explanations are essentially indexical”. The phrase ‘essential indexical’ of course traces back to Perry (1979), but defenders of the thesis also include Castañeda (1971), Lewis (1979), Burge (2000), and Bermudez (2016). Cappelen and Dever (2013) criticize the thesis. Babb (2016) and Bermudez (2017) respond.

are] presented to no one else”. This claim is also widely assumed.²²

Combining these assumptions, if some outcomes are *de se* thoughts and if all *de se* thoughts are private, then some outcomes are private. For example, when Franz is deciding whether to kill or whether Hans and Jens will kill, the outcome entertained by Franz where he himself kills is private. I’ll use ‘ I_F ’ to represent that outcome, distinguishing it from other presentationally-distinct outcomes where Franz kills, such as F , the publicly-accessible third-personal counterpart of I_F . Since an agent’s outcomes, as we’ve seen, are partly a function of her thoughts and since only Franz is in a position to think his *de se* thoughts, necessarily, I_F is an outcome only for Franz. This makes I_F private.

I_F is essentially *de se* or first-personal; it is an outcome that only Franz can entertain. I_F contrasts with F which, as I’ll use it for the rest of the discussion, is an essentially *third-personal* outcome. Just as entertaining I_F requires thinking of Franz in a “special and primitive way” that only Franz can, entertaining F requires entertaining an outcome not only where Franz kills but also where Franz is represented as distinct from oneself. As a result, Franz cannot have F as an outcome unless he is suffering from some kind of *de se* confusion. The moment that confusion lifts, he loses F as an outcome since, necessarily, that outcome represents Franz as a distinct agent or, as I put it earlier, “as a stranger to himself”.²³

²² Bermudez (2016) denies that *de se* thought is private. According to him, “an account of the sense of “I” must allow tokens of “I” to have the same sense as tokens of other personal pronouns such as “you” in appropriate contexts” (79). However, Bermudez’s package of views offers no succor to opponents of *de se* consequentialism. Bermudez (2020) defends a notion of ultraintensionality, such that even if tokens of “I” can share a hyperintension with tokens of “you”, they nevertheless differ in ways that allow those tokens to be evaluated differently, which is all that *de se* consequentialism requires. This shows that *de se* consequentialism relies on weaker assumptions than those that I adopt above. I adopt them only to simplify exposition.

²³ One might worry that known identities pose a problem for *de se* consequentialism. When Lois knows that Clark is Superman, then she cannot rationally assign different values to her Clark and Superman outcomes. Likewise, when someone knows that F and I_F represent the same state, they must assign them the same value. But, as we’ll see, mimicking constraints requires assigning those outcomes different values.

This worry is less forceful than it appears. Since only I am in the position to think my *de se* thoughts, only I am in a position to assign values to $I_{@AU}$. When I know that @AU and $I_{@AU}$ represent the same state, @AU is *not* one of my outcomes, since it is essentially a third-personal representation of myself. Just as Franz can have F as an

Outcomes' privacy is weaker than it may seem in at least two ways. First, though your outcomes are private, they nevertheless permit publicly accessible explanations of your actions. This follows from a more general observation about first-personal thought. Even if it is essentially private, I can still grasp the *de se* thought expressed by your utterance of "I am thirsty". Perhaps I grasp it by implicitly grasping the rule that takes your utterance and its context to a Kaplanian content; perhaps I grasp some other explanation. But whatever that explanation is, *de se* thoughts are not totally inscrutable simply because they are private. By the same token, although *de se* outcomes are private, they contribute to explanations of intentional action. For example, when we're trying to explain why Mark is drinking and I claim that it's because he wants to quench his thirst, then my claim does not fully capture the outcome that Mark desires. But it's close enough that we can fill in what's missing using extra-linguistic means.

Second, that outcomes are private does not imply the absurd proposition that an action affects only its agent or anything of the sort. That confusion rests on misunderstanding the meaning of 'outcome' in this context. 'Outcome', as I'm using it, is a consequentialist term of art, denoting the bearer of value on which an act's rightness allegedly depends. As such, even if an action's intended outcome is private, its effects need not be.

Private outcomes reconcile agent-neutral consequentialism and agent-centered constraints because their imagined conflict relies on certain outcomes being public. For example, on the standard view, the total set of outcomes in Schroeder's scenario is {J, F, H, J&H, J&F, H&F}. Jens's perspective

outcome only if he is suffering *de se* identity confusion, I can have @AU as an outcome only if I'm similarly confused and I can rationally assign @AU and $I_{@AU}$ different values when confused.

Some might insist that I can have @AU as an outcome even if I'm not confused because I can imagine myself from a third-personal perspective while identifying with the imagined person. First, such thoughts have long puzzled philosophers. It seems plausible that when we imagine ourselves third-personally, we are engaging in the mere pretense that we are distinct from the imagined person. Or perhaps our identity is distributed across two distinct perspectives. In any case, neither description seems to fit outcomes like @AU, which essentially present their agents as distinct from oneself. Second, even if I can imagine myself third-personally while maintaining the integrity of my identity, *imagining* @AU and *having* @AU as an outcome are two different conditions. I'm concerned only with the latter. Observations about the former are orthogonal.

implies that $H\&F > J$ and Hans's perspective implies that $J\&F > H$. Because we're assuming that the outcome where Jens kills (J) in Jens's choice set (that is, $\{J, H\&F\}$) overlaps with one that appears in Hans's choice set (that is, $\{H, J\&F\}$), two instances of the principle that more killing is always worse — $J > J\&F$ and $H > H\&F$ — combine to yield a contradiction.²⁴

We can derive a similar result more informally. Jens absolutely must not kill only if $H\&F > J$. A constraint on killing therefore implies that J is extremely bad. But if J is extremely bad, then, plausibly, others should sometimes kill in order to avoid J, which straightforwardly conflicts with a constraint on killing. Because J is part of at least two different agents' outcomes, assigning J the value necessarily for securing a deontic verdict for one agent may ensure the wrong deontic verdict for the other agent. The contradiction falsely thought to demonstrate the inconsistency of agent-neutral consequentialism and constraints thus depends on public outcomes.

De se consequentialism's private outcomes, however, avoid this implication. Privacy limits their deontic import to only the single individual whose *de se* thoughts constitute the outcome. For example, because I_F is only Franz's outcome, only his actions are affected by I_F 's value (or disvalue). Because the deontic significance of *de se* outcomes is limited to a single agent, we can claim that I_F is very bad indeed without implying anything at all about what Jens or Hans should do (and similarly for I_H and I_J).

Accordingly, private outcomes let consequentialists mimic constraints. *De se* consequentialism implies that Hans's choice set is $\{I_H, J\&F\}$, Franz's choice set is $\{I_F, H\&J\}$, and Jens's choice set is $\{I_J, F\&H\}$, yielding a different total set: $\{I_H, I_J, I_F, H\&J, H\&F, J\&F\}$. Critically, since *de se* outcomes are private, no *de se* outcome where an agent herself kills is part of a different agent's choice set. As a result, the total set of outcomes $\{I_H, I_J, I_F, H\&J, H\&F, J\&F\}$ can be ordered to mimic a non-consequentialist prohibition on killing just so long as the *de se* outcomes are the worst possible outcomes. This assumption yields the following ranking in terms of the agent-neutral better-than relation:

²⁴ Recall that, given Hans's perspective, $J\&F > H$. But $H > H\&F$ by the principle. So $J\&F > H\&F$ by transitivity. Likewise, given Jens's perspective, $H\&F > J$. But $J > J\&F$ by the principle. So $H\&F > J\&F$ by transitivity. However, $>$ is anti-symmetric, resulting in contradiction.

$H\&J=H\&F=J\&F > I_H=I_J=I_F$.²⁵ According to this ranking, Hans's choice set $\{I_H, J\&F\}$ implies that Hans should not kill since $J\&F > I_H$ and similarly for the rest.

To be clear, I have not yet *justified* ordering outcomes in this way. I've merely shown that *de se* consequentialism lets us do so, thereby squaring constraints and agent-neutral value, which is the paper's main aim. Whether we can justify this ranking on independent grounds and whether *de se* consequentialism is an independently attractive view are *further* questions, which I begin to address below. But I want to be explicit that articulating *de se* consequentialism as I've just done suffices to rebut Portmore, Schroeder, Brown, Smith, Nair, *et al.* since it establishes that agent-neutral consequentialism can be reconciled with constraints.

3.3 *De Se* Value

3.3.1 For all agents X , $X > I_X$

De se consequentialism implies what I'll call '*de se* value', which is ordinary, agent-neutral value (or disvalue), grounded in what's distinctive about first-personal thoughts about killing (or whichever other act is forbidden by a constraint). We can demonstrate this implication by adding the perspective of a bystander to the case of Hans, Franz, and Jens. Suppose that Klaus is watching the three bicker about who's going to get their hands dirty. We're assuming that every killing is bad. Accordingly, from Klaus's perspective, Jens killing one is better than Hans and Franz killing two, and *mutatis mutandis*

²⁵ To avoid proliferating symbols, I'm not using '=' to express identity but rather 'is equal in value to', the binary relation between outcomes that holds when their value is identical. I should note that the ordering above violates a stronger principle than the principle that killing is always bad, according to which two killings are always worse than one: for any X, Y , and $Z \in \{H, J, F, I_H, I_J, I_F\}$ such that $Y \neq Z, X > Y\&Z$. But this principle is straightforwardly inconsistent with a constraint on killing, since it entails that it is always better to kill one in order to save two.

We might reasonably deny this stronger principle on independent grounds even if we are drawn to the weaker one. The weaker principle does not make comparisons between the disvalue of various killings. It holds only that killing is always bad. The stronger principle, by contrast, holds that two killings are always worse than one. If we think that killings are not comparable in disvalue or that killings can differ sufficiently in value, we should reject the stronger principle regardless of our sympathy for consequentializing.

for the rest. Klaus's perspective suggests, quite intuitively, that $H=J=F > H\&J=H\&F=J\&F$. However, if we assume that $H\&J=H\&F=J\&F > I_H=I_J=I_F$ to mimic constraint on killing, transitivity implies that $H=J=F > I_H=I_J=I_F$, from which we can extrapolate that for all agents X , $X > I_X$ — that is, *de se* killings are worse than their third-personal counterparts. I'll focus on *de se* value for the rest of the paper.²⁶

De se value might seem puzzling: that an outcome's presentation can make it worse is odd. But this oddness offers no objection to *de se* consequentialism. That's because I'm not relying on pre-theoretically plausible claims about *de se* value to justify an absolute constraint on killing in the way that we might justify a particular policy by the number of people it helps. Rather, the order of explanation is the opposite: I'm showing what a constraint on killing entails about the structure of value in an agent-neutral consequentialist framework. Consequently, *de se* consequentialists might agree that *de se* value is odd but insist that the evidence for agent-neutral consequentialism is so great and that the intuitive force of a constraint on killing is so strong that *de se* value is worth embracing as a genuine moral discovery, regardless of how puzzling it may seem.

As a result, I don't need to show that *de se* value is plausible independently from our evidence for consequentialism and constraints. I need only to show that it is coherent. However, one could reasonably suspect otherwise, particularly given an initially compelling argument that I call *the Supervenience Argument*. According to it:

1. X and I_X represent the same metaphysical possibility.
2. The normative supervenes on the non-normative.
3. If X and I_X represent the same metaphysical possibility, they are non-normatively alike.
4. (MP1-3, 2) Therefore, X and I_X are normatively alike, so *de se* value violates supervenience.

²⁶ Muñoz (2020) argues that intrinsic differences between actions cannot ground moral differences between outcomes for consequentialists. One might worry that this threatens *de se* value. However, since, for example, J and I_J are caused by *token identical actions*, moral differences between those outcomes are not the result of intrinsic differences between actions. As a result, *de se* value is a 'narrow', not 'broad', feature to use Muñoz's terminology. Appealing to that feature therefore complies with the stipulation that only narrow differences between outcomes are moral differences.

I've defended (1) and I assume (2), so I'll rebut the Supervenience Argument by showing how (3) mischaracterizes outcomes, at least as the *de se* consequentialist understands them.

3.3.2 Representation and Supervenience

It's natural to assume that evaluative differences between outcomes wholly supervene on differences between the metaphysical possibilities that they represent. But that's false: Lois Lane's Superman outcomes differ from her Clark Kent outcomes while being metaphysically identical to them. Hence, it's a mistake to treat outcomes simply as proxies for metaphysical possibilities. Treating them as proxies collapses subjective consequentialism into objective consequentialism.

An analogy helps to diagnose the mistake behind the Supervenience Argument. It is natural to think of, for example, @AU and $I_{@AU}$ as two different vistas on the same landscape. One vista might provide a better view than the other, but that doesn't make the landscape any more beautiful. Likewise, it's tempting to think that though the moral repugnance of a killing may be more vivid from a first-personal standpoint, that doesn't make the killing itself any more repugnant.

But *de se* consequentialists don't conceive of outcomes in this way. Outcomes — an act's foreseen effects — themselves are the bearers of deontically-relevant value, not the states of affairs that they represent. Normative supervenience therefore implies that evaluative differences between outcomes necessitate non-evaluative differences between outcomes, not in the *states* that those outcomes represent. Thus, because outcomes are evaluatively significant in their own right, a more appropriate analogy is thinking of @AU and $I_{@AU}$ as two *paintings* of the same landscape that differ in value because one is beautiful and one is ugly. Evaluative differences between paintings do not supervene on what those paintings represent. It is clearly a mistake to think that two paintings differ in their degree of beauty only if they differ in what they depict. There are beautiful paintings of ugly landscapes and ugly paintings of beautiful people. Although paintings are representations, their value is at least partly independent from what they represent. If expected outcomes are like paintings in this respect, it's also a mistake to think that outcomes' evaluative properties supervene on what they represent.

Thus, premise (3) of the Supervenience Argument is false. Normative supervenience is consistent with *de se* value. For example, the state of affairs where Clark Kent arrives on the scene is identical to the one where Superman arrives. Nevertheless, the outcome corresponding to the second state of affairs exemplifies a property that the first lacks: the second is an outcome where it's likely that the world is saved, relative to Lois's information. This epistemic property — being such that it's likely that the world is saved — is *not* a property of the state of affairs that both outcomes represent since the outcomes differ relative to this epistemic property while representing the same state of affairs. Rather, it's a descriptive property of the outcome in its own right, *qua* representation, that makes a considerable difference to its value. The property of being *de se* is similar: a descriptive property of an outcome that affects its value.

The analogy between paintings and *de se* outcomes has limits. There are ugly paintings of beautiful things and beautiful paintings of ugly things. The analogy thus suggests an unconstrained relationship between an outcome and the corresponding event or state-of-affairs. But an analogous decoupling of the value of an outcome from the value of the corresponding event is unappealing; surely the value of the outcome is at least partly a function of the value of the event.

A second analogy is helpful here. Arpaly and Schroeder (2013) argues that co-referring representations can differ evaluatively because a particular kind of goodness, *good will*, depends on presentational differences like those that make *de se* consequentialism tick.²⁷ According to Arpaly and Schroeder, “complete good will is an intrinsic desire for the right or good, *correctly conceptualized*.”²⁸ This account straightforwardly implies that two representational states (in this case, desires) can vary in value (particularly, in their quality of will) while representing the same state. For example, the desires of an alien who finds the firing of c-fibers very unpleasant, like fingernails scraping over a chalkboard, will involve the same state of affairs as the desires of a pure-hearted altruist who abhors seeing people in pain. However, only the altruist's actions exemplify good will: when the alien is moved to alleviate someone's suffering, the alien's deployment of the concept *c-fiber firing* does not ‘correctly

²⁷ A referee for *Ethics* astutely observes that *de se* value is unusual because it seems to be a kind of *non-instrumental* but *derivative* value. The value of a good will is similar in this respect.

²⁸ Arpaly and Schroeder (2013: 163), emphasis mine.

conceptualize' the suffering. The quality of the alien's will is thus worse than the altruist's, though both are moved by desires that differ only in how they present a given state. However, the alien's actions do not merit praise because of this difference.

De se consequentialism is similar: simply swap 'good will' with '*de se* (dis)value' and swap 'merits praise' with 'is right'. Arpaly and Schroeder hold that some acts differ in their praiseworthiness based only on differences in good will grounded in presentational differences between desires. *De se* consequentialism holds the analogous position that some acts differ in their rightness based only on differences in *de se* value grounded in presentational differences between outcomes. The value of an outcome is, therefore, still closely related to what it represents: only outcomes associated with constraints can exhibit *de se* value, which can only intensify the badness of a killing. The link between an outcome's value and of the corresponding state is far more closely intertwined than the beauty of a painting and of what it represents.

This analogy between good will and *de se* value is no coincidence. Julia Driver finds a structural similarity between subjective consequentialism and Kantian ethics. She writes,

Subjective forms of consequentialism are also internalist in the same way the Kantian system is, since on one version of subjective consequentialism, for example, right action is defined in terms of expected consequences, and on this version of the subjective view a virtue would be a disposition to act so as to maximize expected good or to try to maximize the good. Across this parameter, then, subjective consequentialism and Kantian ethics are the same sort of ethical theory, whereas subjective and objective consequentialism are radically different, as they locate the source of value very differently.²⁹

Driver's observation bears repeating. Kantianism and various subjective consequentialisms, such as *de se* consequentialism, share a deep and relatively neglected similarity; they explain facts about moral

²⁹ Driver (2001: 70).

rightness by facts about the moral significance of certain internal mental states.³⁰ Given this resemblance, it is no wonder that *de se* consequentialism exhibits features characteristic of non-consequentialist views like Kantianism, such as constraints. The main difference between subjective consequentialism and Kantianism lies in consequentialism's *telic* nature. Whether an agent does what's right, according to subjective consequentialists, depends only on whether they 'aim' at the right outcomes in action — or, equivalently, whether the agent has the right telic state — not on other states such as the 'maxim' of their action. I return to these observations in section 5.2.

4. *De Se* Consequentialism and its Competitors

De se consequentialism's import is mainly negative. It shows that the truth of constraint-supporting judgments does *not* entail the falsity of agent-neutral consequentialism, as is widely assumed. In particular, I have not defended the positive claim that constraints exist. Nevertheless, *de se* consequentialism provides indirect support for constraints by suggesting how they amount to more than mere 'rule-worship'. For example, if we concede that it's better that one is killed than that five are, we are faced with a familiar challenge: why *not* kill the one? Attempting to answer this question "leaves an ugly stigma."³¹ *De se* consequentialists, however, can dismiss the challenge by denying that it is better that I (myself) kill one than that the five are killed by others. Furthermore, rather than countervailing evidence against the truth of consequentialism, consequentialists can assert that the intuitive wrongness of killing one in order to save five tracks the badness of you yourself killing. This allows them to conscript constraint-supporting intuitions as *evidence* of *de se* value.

In sum, philosophers attracted to consequentializing feel the pull of both consequentialism's compellingness and of their constraint-supporting intuitions. *De se* consequentialism matters because it shows that these two claims do not pull in opposite directions. Rather, they pull towards accepting *de se* value.

³⁰ 'Internal' may have to be adjusted to account for factive mental states.

³¹ Dreier (1993: 24).

4.1 Contrasting Agent-Relative Consequentialism and *De Se* Consequentialism

Contrasting *de se* consequentialism with agent-relative consequentialism highlights the former's appeal. Thomas Hurka alludes to a subtle ambiguity when he describes the "concept of agent-relative goodness, or what is good from a person's point of view" (2003: 611). What's good relative to an agent and what is good from their point of view differ but can easily be confused. For example, if we meet for lunch, and I order a soup and you order a sandwich, it might appear that we disagree about what's good to order. Or it might appear that we're using an agent-relative concept, *good-to-order-relative-to-X*. However, if it turns out that the waiter has handed me a menu with only soups on it and that he has handed you one with only sandwiches on it, then what's good to order from our differing points of view differs, but that doesn't entail that we have different conceptions of what's good to order. We needn't be employing the concept *good-to-order-relative-to-X*. Rather, we can be thinking lucidly and consonantly about what's good to order relative to different menus.

The very same mechanism that explains our different lunch orders also allows agent-neutral consequentialists to enact agent-centered constraints. New Schoolers assume that incorporating some measure of agent-centering in consequentialism requires an agent-centered notion of value. This assumption is natural for consequentialists characteristically understand rightness through goodness. As a result, it's natural to assume that mimicking an agent-centered notion of rightness requires an agent-centered notion of goodness. But *de se* value relies on the inherent differences between agents' points of view, particularly with respect to *de se* thought, to enact constraints.

The mechanism is subtle. This is so partly because, at a certain level of abstraction, *de se* consequentialism and agent-relative consequentialism involve closely related formalisms. The standard recipe avoids contradiction in Schroeder's scenario by distinguishing the evaluative relations that H&F and J&F stand in, relative to Hans and Jens. Rather than implying the contradiction that $H\&F > J\&F$ and $J\&F > H\&F$, the standard recipe implies that $H\&F >_J J\&F$ and $J\&F >_H H\&F$. A process known alternately as Currying or Schönfinkelization entails that every distinction between relations, like the standard recipe's distinction in evaluative relations, is logically equivalent to a distinction in their *relata*, like *de se* consequentialism's distinction between outcomes like I_F and F . For example, we might

say that the silver medalist exemplifies two, distinct monadic properties: being-next-to-the-gold-medalist and being-next-to-the-bronze-medalist. Curryng offers an equivalent description. We could instead claim that the silver medalist stands in two instances of the same binary property *being-next-to*: once by being next to the gold medalist and again by being next to the bronze medalist. Curryng shows that these two descriptions are logically equivalent: necessarily, you stand in the ‘next to’ relation to the gold medalist iff you satisfy the monadic property of being-next-to-the-gold-medalist. Curryng the standard recipe distinguishes a particular killing as it appears as Jens’s outcome from how it appears as Hans’s outcome rather than distinguishing Jens and Hans’s evaluative standards.

To mimic constraints, consequentialists must hold that “the amount of value a state of affairs has [varies] from agent to agent.”³² Curryng limits what we can vary. In particular, since what an agent should do is a strict function of which of her available *outcomes* is *best*, to mimic constraints we either have to make available outcomes agent-centered or we have to make value agent-centered. While the standard recipe makes value agent-centered, *de se* consequentialism offers a natural and independently motivated account of outcomes that are, in a certain sense, agent-centered. In other words, agent-relative consequentialism varies the value associated with a state of affairs by varying the function from states of affairs to values from agent to agent. *De se* consequentialism, by contrast, varies the outcome associated with that state of affairs from agent to agent. Thus, the main difference between the two is that the former makes the evaluation of an outcome agent-relative whereas the latter makes the outcome itself agent-relative.

While perhaps formally trivial, this difference is philosophically substantive. Agent-relative consequentialism’s main challenge, as we’ve seen, is squaring agent-relative goodness with the ‘compelling’ consequentialist idea that it’s never wrong to do what’s best. *De se* consequentialism, in contrast, does not make evaluation agent-relative. It makes what’s evaluated agent-relative. So it uses ‘good’ in the sense expressed by the compelling consequentialist thought that it’s never wrong to do what’s best.

³² Oddie and Milne (1991: 44).

As a result, *de se* consequentialism does not involve a novel *kind of value*. It involves only a novel value *bearer*. My invitation to consider *de se* value thus resembles an egalitarian's invitation to consider the (agent-neutral) value of equality. Utilitarians may decline the invitation, but not because the value allegedly exhibited by equality is *unintelligible*. Utilitarians simply doubt that equality has familiar, intelligible agent-neutral value. Similarly, when I claim that *de se* killings are worse than their *de se* counterparts, I mean that the difference makes the former worse in the familiar way that other bad-making features do. Worries about intelligibility, such as those that afflict $>_x$ do not apply to *de se* value.

Moreover, *de se* consequentialism's sensitivity to presentational differences is a considerable advantage over agent-relative consequentialism in cases of *de se* or indexical confusion, such as the grocery case above. For example, forms of consequentialism imply that you should give away the grocery prize only if they recognize a difference between you getting the prize and the woman in the other queue (who happens to be you) getting the prize. But once we recognize that difference, we no longer need agent-relative value to reconcile consequentialism and constraints, as I've shown above.

Nevertheless, *de se* consequentialism's advantage in vindicating the compelling consequentialist thought can be doubted.³³ That thought is compelling partly because when something is agent-neutrally best, then it's best in a public or non-parochial way. However, if some outcomes are private, then their value, it might be thought, is private. So even if *de se* consequentialism complies with the letter of the compelling thought, it might seem to violate its spirit.

This impression is misled. For example, your qualitative experience of your own happiness may be inaccessible to me but that does not diminish its authority over what I should do in a utilitarian account of moral obligation. Though your experience of happiness is private, the authority of its value is public. Likewise, though, necessarily, a *de se* outcome is accessible to only a single agent, the outcome is not good or bad *only for her*. It is good or bad in a public, agent-neutral way.

³³ I thank a referee for *Ethics* for pressing me to clarify my position here.

Outcomes' privacy, however, is incompatible with a second appealing feature of consequentialism, beyond the compelling thought.³⁴ Consequentialism is appealing for social policy work because it offers a simple way for value to be added up across persons. However, if my *de se* outcomes are outcomes only for me, then their value cannot be straightforwardly aggregated from the policy-maker's point of view: from that perspective, your *de se* outcome is a third-personal outcome. As a result, *de se* value 'disappears' when considered from a third-personal point of view. *De se* value cannot, therefore, be part of policy or public decision-making.

This concern is a feature of constraints themselves, not of *de se* consequentialism *per se*. On a simple deontological framework, the wrongness of *my* killing matters to my decision making in a way that it doesn't matter to your decision making and *vice versa*. Thus, this shortcoming appears to be inherent to constraints themselves, afflicting any way of enacting them, not just *de se* consequentialism's particular method of doing so.

De se value might seem to violate the spirit of the compelling thought in a different sense. Following Scanlon's (1998) buck-passing analysis of value, suppose that, necessarily, something is agent-neutrally bad just when there is sufficient agent-neutral reason to disvalue it. Following Schroeder (2007), suppose that, necessarily, there is sufficient agent-neutral reason to disvalue something just when there is sufficient reason for *anyone* to disvalue it. Thus, if the outcome where I kill has *de se* (dis)value, then everyone has a reason to disvalue me myself killing. And that's very odd if only I can think the thoughts involved in that outcome.

While odd, this implication is not fatal. Although the *de se* (dis)value of me myself killing gives a reason for anyone to disvalue that outcome, only I am in a position to act on it.³⁵ The outcome falls under the sphere of my agency and no one else's. Insofar as the implication is odd, it is benign.

³⁴ I am grateful to [redacted] for discussion here.

³⁵ My position echoes those taken by Wedgwood and Broome on a similar issue, as reported in Schroeder (2011: 33).

4.2 *De Se* Thoughts, Centered Worlds, and Properties

There are, in principle, as many ways to Curry the standard recipe as there are ways of individuating outcomes by agents. For example, Campbell Brown and Jamie Dreier each explore mimicking constraints by distributing value over centered worlds or properties, which Lewis (1979) uses to model *de se* content. Yet I've chosen to Curry the standard recipe using *de se* thoughts, without mention of centered worlds or properties. It's natural to wonder whether that choice is sound.

While centered worlds or properties help make *de se* thought formally tractable, focusing on them in this context obscures *de se* consequentialism's most important feature: private outcomes. As we've just seen, *de se* consequentialism can mimic constraints only because some of its outcomes are private. Their privacy follows from the assumption that some outcomes are *de se* thoughts and that all *de se* thoughts are private. Once we understand how some outcomes are *de se* thoughts, private outcomes follow naturally.

However, rather than distinguishing directly between first- and third-personal thoughts, we could instead distinguish between ordered pairs taken to represent those thoughts, as Jamie Dreier (2011; 2020) does. While this paper owes much to his work, applying Lewis's centered-worlds approach to the project of consequentializing is more hindrance than help. Some (principally Cappelen and Dever (2013) and Magidor (2015)) have offered potent criticisms of Lewis's approach. These criticisms suggest that Lewis's account relies not only on centered worlds or properties but also on an allegedly cryptic *hyperintensional* notion of *self-predication*.

To see this, suppose for *reductio* that self-predication is merely intensional — that is, agent x predicating property F of x suffices for x to *self-predicate* F . Now suppose that Heimson is deceived about his own beliefs. Although he believes that he's Hume, the cognitive dissonance of that false belief has put Heimson in denial. As a result, Heimson does not believe that he himself believes that he's Hume. If H is Heimson's set of beliefs then a Lewisian centered worlds approach implies:³⁶

1. $\{ \langle w, x \rangle : w \in [\text{Heimson is Hume}] \text{ and } x = \text{Heimson} \} \subseteq H$
2. $\{ \langle w, x \rangle : w \in [\text{Heimson believes that he is Hume}] \text{ and } x = \text{Heimson} \} \not\subseteq H$

³⁶ Strictly speaking, the center must also include a time but that complexity does not rescue the account. See Liao (2012) on this point.

One day he sees a handsome stranger in line at the grocery store wearing a dashing pink velvet turban. Heimson recognizes the turban from the well-known portrait of Hume. The man must have chosen to don the turban, Heimson reasons, inferring that he must think that he's Hume. However, Heimson fails to notice that he's looking at his own reflection. The stranger is Heimson himself. Consequently:

3. $\{ \langle w, x \rangle : w \in [\text{He is Hume}] \text{ and } x = \text{Heimson} \} \not\subseteq H$
4. $\{ \langle w, x \rangle : w \in [\text{He believes that he is Hume}] \text{ and } x = \text{Heimson} \} \subseteq H$

Since 'he' refers to Heimson, for all worlds w , $w \in [\text{He is Hume}]$ iff $w \in [\text{Heimson is Hume}]$ and $w \in [\text{Heimson believes that he is Hume}]$ iff $w \in [\text{He believes that he is Hume}]$. As a result, (1-4) are inconsistent. Therefore, centered worlds, all by themselves, don't seem to capture the difference between first-personal and non-first-personal thought. As Liao (2012: 314) puts it, "there seems to be a fundamental conceptual distinction between ascribing properties to *oneself* and ascribing properties to an individual possessing a unique and exhaustive list of non-trivial properties."

If challenges like the one above from Cappelen and Dever (2013) and Magidor (2015) are sound, Lewis's account requires a special kind of self-predication. Working *only* with centered contents, at least as they are conventionally understood, seems to fall short of capturing the nuances of *de se* belief. And since private outcomes depend on the nuances of *de se* belief, employing centered worlds to model the outcomes involved in consequentialism at best obfuscates how to reconcile consequentialism and constraints. At worst, employing centered worlds rather *de se* thoughts fails to reconcile them altogether.

Indeed, because of this obfuscation, Campbell Brown (2011) is misled into denying that constraints and consequentialism are consistent. He begins by assuming that the outcome of an action is an entire world.³⁷ Then, to explore agent-relativity, he pairs worlds with individuals. He writes, "If the centered worlds $\langle x, i \rangle$ and $\langle y, j \rangle$ are centered on different people (i.e., if $i \neq j$) then they make no

³⁷ In Brown (2011: 753-4), Campbell Brown explicitly excludes consequentialisms that rely on "subjective probabilities", such as *de se* consequentialism, on the grounds that he is interested in objective, not subjective, rightness. But theories of objective rightness that incorporate subjective probabilities are perfectly consistent. So Brown's interest in objective rightness does not diminish *de se* consequentialism's relevance to his project. I am neutral on whether *de se* consequentialism explains objective rather than subjective rightness.

difference to whether [a theory of what's right maximizes value]" (759).³⁸ That's true: individuating outcomes by individuals does not, all by itself, reconcile consequentialism and constraints. After all, regardless of how we evaluate things, some set of centered worlds will have maximal disvalue and, other things equal, the consequentialist theory will permit agents to kill in order to avoid them. Only if the very worst outcomes are *private*, as *de se* thoughts are, can we be sure that we'll never have to kill in order to avoid them. But thinking of outcomes as centered worlds obscures their privacy, which is what misleads Brown.

He concludes with the following reflection:

This way of trying to accommodate agent-relative theories within consequentialism therefore fails. To be clear, I do not claim that agent-relative views are incoherent. I claim only that they are incompatible with the view that there is a single thing, the good, that we all ought to maximize. Such views are better thought of as saying that there are many things that ought to be maximized, a different one for each of us. (763)

Brown is, as I've just argued, mistaken to claim that views like *de se* consequentialism are "incompatible with the view that there is a single thing, the good, that we all ought to maximize". Moreover, his claim that "such views are better thought of as saying that there are many things that ought to be maximized, a different one for each of us" is ambiguous. *De se* consequentialism is *not* "better thought of as saying that there are many [kinds of value] that ought to be maximized, a different [value] for each of us." Its key claim is "that there are many [sets of outcomes whose value] ought to be maximized, a different [set] for each of us."

³⁸ Brown assumes that a theory of what's right is consequentialized by a theory of value just when what's right (R) for an individual (i) given a set of options (A) — that is, when $R_i(A)$ — is the most valuable (grt_{\leq}) option in A — that is, when $R_i(A) = grt_{\leq}(A)$. However, if ought implies can in at least one sense of can, then the only options that can be right for an agent are ones choosable by that agent. So we should relativize A to individuals, yielding A_i , and instead hold that a theory of value grt_{\leq} consequentializes a theory of right (R) just when $R_i(A_i) = grt_{\leq}(A_i)$. If options vary between individuals in the way that *de se* consequentialism suggests, then it's possible that $R_i(A_i) = grt_{\leq}(A_i)$.

In sum, *de se* consequentialism achieves constraints by varying outcomes between individuals, not by varying kinds of value. *De se* consequentialism is thus best understood as the view that the absolute, agent-neutrally worst thing, to which all other outcomes are preferable, is a *de se* killing. But your *de se* killings are never my outcomes nor *vice versa*.

5. Fault Lines in Ethical Theory

5.1 Companions in Guilt

De se consequentialism is noteworthy not just for consequentialists sympathetic to constraints, but also for anyone interested in what we can call, following Nair (2014; 2020), the fault lines that divide various positions in ethics. For example, in his influential “Structures of Normative Theories” and in subsequent work, Jamie Dreier defends the claim, which I have defended here, that the agent-neutral/agent-relative distinction is orthogonal to the consequentialist/non-consequentialist distinction: accepting agent-relativity in ethics does not require rejecting consequentialism. However, this claim might be taken to diminish the importance of the agent-neutral/agent-relative distinction, which, philosophers largely agree, is one of the “greatest contributions of recent ethics.”³⁹ So it’s worth asking whether constraints trace a genuine fault line in ethics. I will argue that they do, just not the fault line between consequentialism and non-consequentialism.

I’ll focus in this final section on identifying and defending one claim in particular: that claims expressing agent-centered constraints typically — perhaps necessarily — involve indexical terms, particularly ‘I’. Doing so shows that *de se* value is noteworthy not only because it grounds a consequentialist account of constraints, explaining why they might amount to more than mere ‘rule-worship’. It also illuminates how any moral theory that posits constraints thereby attribute moral significance to differences like the one between first- and third-personal thought. As we will see, drawing out the essential indexicality of many agent-centered constraints supports Darwall (1986)’s distinction between ethical theories that start from the ‘point of view of the universe’ and those that

³⁹ Hurka (2003: 628).

start from the agent's point of view — or, as he puts it, theories that start from the 'inside-out' and those that start from the 'outside-in'.

This fault line merits attention not only for its own sake but also because it provides a 'companions in guilt' defense of *de se* value. *De se* value is at least *prima facie* puzzling. But it is merely the consequentialist aspect of the puzzle that agent-centered constraints pose for *any* ethical theory, whether consequentialist or not. Moral theories, as we'll now see, seem to require attributing moral significance to what I've been calling presentational differences to enact constraints. Insofar as *de se* value is odd, therefore, that's only because agent-centered constraints are odd. The oddity of *de se* value offers an argument against constraints themselves, not their expression in *de se* consequentialism.

Deontological and reasons-based ethical frameworks, which are often taken as competitors to consequentialism, express constraints as follows:⁴⁰

1. It is always wrong to kill.
2. There is never most reason to kill.

Neither claim explicitly mentions an agent. This is essential, for one important role of constraints is to offer practical or deliberative guidance about how to act. However, stating (1) with an agent, as in (3) below, undercuts its practical guidance:

3. It is always wrong *that anyone* kills.

For example, in cases like *Mafia*, it is certain that someone will kill. Consequently, (3) offers no practical guidance on whether to kill in *Mafia* since killing is unavoidable. In contrast, (1) tells us what to do: not kill, even if someone else will. Similarly for (2).

As a special case of (3), (4) also fails to play constraints' practical role:

4. It is always wrong for @AU to kill/ that @AU kills.

I've introduced '@AU' as a name for me, the actual author of the article. As such, it's natural to suppose that a constraint on killing, as it applies to me, is captured by something like (4). But (4) omits

⁴⁰ Hammerton (2020) discusses various logically distinct formulations of constraints. Although grammatically distinct, (1) and (2) above are roughly analogous to his formulation #4, with the *caveat* that I have assumed that the correct treatment of *de nunc* outcomes in a consequentialist framework parallels *de se* consequentialism's treatment of *de se* outcomes.

the first-person, making it susceptible to *de se* confusion, such as when I doubt that I'm @AU, which undercuts (4)'s practical guidance.⁴¹ For example, suppose that I accept (4), that it is always wrong for @AU to kill. But suppose that I don't know that I am @AU. When faced with an apparent choice between killing someone or having @AU kill them and I doubt that I am @AU, I will infer that I may kill since it is always wrong that @AU kills, which is inconsistent with a constraint on killing.

The difference is that (1) and (2), in contrast to (3) and (4), use a non-finite clause, namely, 'to kill'. Non-finite clauses in environments like (1) and (2) tacitly involve first-personal thought by covertly expressing a generic first-personal pronoun tied to the *de se*, making them equivalent to 'it is always wrong *that you yourself* kill' and 'there is always most reason *that you yourself* not kill'.⁴² This is a general feature of *to*-clauses embedded under null grammatical subjects.⁴³ For example, 'It is good to see you again', as uttered by me, expresses that it is good that *I myself* see you again. Consequently, as others have noted, deontological and reasons-based theories don't incorporate agent-centered restrictions simply because of their apparent focus on actions.⁴⁴ Rather, (1) and (2) are agent-centered because they incorporate a *covert* indexical element tied to the first person. That is, what makes these restrictions agent-centered is that they're essentially indexical.

With (1) and (2)'s covert *de se* sense made explicit, we're in a better position to appreciate a deep similarity between the three accounts of agent-centered constraints:

5. *De Se Consequentialist*: It is always the worst outcome [that *you yourself* kill].

6. *Conventional Deontology*: It is always wrong [that *you yourself* kill].

⁴¹ Dreier (2020) discusses similar cases of 'indexical uncertainty'. He aims to show that views like *de se* consequentialism encode strictly more information than views like agent-relative consequentialism and that such information is not "junk information". I believe the arguments above support his claims.

⁴² It is widely acknowledged that *to*-clauses in environments like (1) and (2) express a covert anaphoric or deictic pronoun, PRO. Stanley and Williamson (2001), for just one example, exploit this observation to defend intellectualism about *know-how*.

⁴³ For a comprehensive defense, see Moltmann (2005), with important antecedents in Chierchia (1990). PRO generally inherits its referent from the sentence's grammatical subject, but in sentences with null subjects like (1) and (2) it takes on an *arbitrary* referent, which, Moltmann argues, is necessarily interpreted *de se*.

⁴⁴ For example, Dougherty (2013) articulates an ostensibly agent-neutral deontology. I doubt that such a view is possible, largely for the arguments given above and for those advanced by Hammerton (2017).

7. *Reasons-Based*: There is never most reason [that *you yourself* kill].

The three approaches are distinguished only by the normative property or concept that they privilege: value for consequentialists, norms or rules for deontologists, and reasons for the third. Each enacts the agent-centered constraint on killing by attributing the maximal negative moral status (*e.g.*, ‘always worst’, ‘always wrong’, ‘always most reason to not’) to the first-personal state of *you yourself* killing.

As a result, each faces the uncomfortable questions that may make *de se* value seem odd. In particular, if morality applies to each agent impartially, why should an agent attribute special moral significance to *her* killings? For example, (6) attributes wrongness only to killings that *you yourself* perform. Likewise, (7) implies that *your* killings are uniquely disfavored by the normative reasons that matter. These implications are puzzling because they indicate a profound and obscure moral asymmetry between oneself and others, which conflicts with morality’s apparent impartiality. But the conflict is rooted in the nature of constraints, not in the moral theories themselves.

De se value is, therefore, the species of a genus. A wide range of moral theories attribute special moral significance to first-personal thoughts in order to enact action-guiding constraints. Consequentialists argue that value exhausts moral significance. As a result, consequentialists attribute a novel kind of value to these thoughts, *de se* value. But *de se* value is not clearly more implausible than these other ways of interpreting the moral significance of the *de se*. To be sure, there may be special reasons for thinking that agent-neutral value is particularly ill-suited to capturing the moral significance of *de se* (and, more broadly, indexical) thought, but those reasons are not obvious. This is a question for future research.

The apparent essential indexicality of some agent-centered constraints, regardless of our favoured moral theory, also helps explain observations that others have made about their logical structure. Following McNaughton and Rawling (1995a; 1995b) and refinements by, *inter alia*, Portmore (2001) and Hammerton (2020), agent-centered constraints possess a distinctive distribution of agent-referring variables. This characterization of constraints harmonizes with my claim that constraints are essentially indexical. But beyond that harmony, the idea that agent-centered constraints are essentially indexical improves on the variable-based account in at least two ways. First, it shows that

having the appropriate agent-referring variables is necessary but not sufficient for being agent-centered; the referents of those variables must be understood first-personally or be undermined by cases of *de se* confusion. Second, the essential indexicality of some agent-centered constraints explains *why* such rules are unified by that particular distribution of variables: that distribution is necessary for the rule to be understood first-personally, which is necessary, as I've just argued, for interpreting it as an agent-centered constraint.

Finally, the essential indexicality of constraints would explain why “it is at least sometimes impermissible to violate [a constraint] in circumstances where a violation would serve to minimize total overall violations of the very same restriction, and would have no other morally relevant consequences”. If constraints are essentially indexical, there is an asymmetry between my killing and your killing, from my perspective, and between your killing and mine, from your perspective, which explains why a killing by one person cannot offset a killing by another. For example, I can't offset the (*de se*) badness of my killing by preventing the *de se* disvalue that would obtain from your killing. The *de se* disvalue of *your* killings is never part of my outcomes. Thus, I am never in a position to reduce the amount of *de se* disvalue (or failure to respect one's reasons or wrongness) that you put into the world, nor *vice versa*.

5.2 Ethics from the Inside-Out

Two points from the preceding discussion merit particular attention. First, enacting constraints requires the kind of information that appears only from the agent's point of view; it is therefore the universe's point of view, often assumed by consequentialists, that is incompatible with constraints, not consequentialism itself.

Second, more complex (but no less consequentialist) forms of consequentialism, like some subjective ones, allow constraints when they take the agent's point of view. Normative concepts like *good* are, therefore, more plastic than theorists typically suppose. A given ethical theory's first-order permissions and prohibitions are shaped more by its point of view than by its preferred flavour of normativity. Thus, the traditional fault lines between ethical theories traced by their privileged

normative concept (rights, goodness, virtuous character, reasons, etc.) may not carve at the most central ethical joints. Class ethical theories by whether they adopt the agent's point of view or whether they adopt what has been called, variously, Sidgwick's "point of view of the universe", Nagel's "point of view of nowhere", or Williams' "absolute conception of the world" cuts at a deeper joint.

A distinct route leads Stephen Darwall to a similar conclusion. He traces a fault line between ethical theories that start from the 'outside-in' and those that start from the 'inside-out', writing:

To put it in a rough and preliminary way, moral integrity involves a person's guiding his life by his own moral judgment, properly understood, and the fundamental responsibility of the moral life is the maintenance of integrity, so conceived. Instead of beginning outside the moral agent with a view about states of affairs that are intrinsically worthy of promotion, the alternative begins inside the moral agent with a view about moral character and integrity. The rationale for agent-centered restrictions is itself agent-centered.⁴⁵

Darwall's point is this. Some values don't appear in states of affairs, considered apart from the point of view of particular agents.⁴⁶ If we start theorizing about morality using only those states of affairs, and we are willing to recognize no values apart from those that appear in them — if we theorize about morality "from the outside-in", as Darwall puts it — then values that are visible only to agents, like integrity, seem puzzling. For example, if we start theorizing about a painting's beauty without the painting itself or a photo of it, but rather with a (very long) numerical readout specifying the location and velocity of each particle that makes up the painting, it's going to be very difficult, perhaps impossible, to judge whether the painting is beautiful using only the readout. The painting's beauty

⁴⁵ Darwall (1986: 305).

⁴⁶ Darwall's claims echo some of Locke's claims about secondary qualities. Secondary qualities are, roughly, dispositions of particulars to cause certain perceptual reactions in creatures like us. For example, properties like being loud or being crimson are secondary qualities but the sound and light wavelengths associated with them are not. Whether secondary qualities exist is controversial. But many philosophers, such as Wiggins (1987: 107), McDowell (1998: 112-150), Wright (1988), and Lillehammer (2007: 89-110) have found it fruitful to analogize moral value with them. I also pursue this analogy above.

won't appear when approaching it from the 'outside-in'. Rather, we must behold the painting with our own eyes, from our own perspective, to apprehend its beauty. The same holds true for *de se* value as a kind of moral ugliness. *De se* value does not appear when we examine what *de se* thoughts represent. It emerges only when we think the thoughts themselves. So it makes a great difference whether ethical theorizing starts from the point of view of the universe, which excludes or prescind from indexical thought, or from the point of view of a particular agent, which is constituted by it.

Objective consequentialism is a paradigmatically 'outside-in' approach to ethics. It limits itself to only those evaluative features that belong intrinsically to states of affairs. If Darwall is right, those features cannot provide a rationale for agent-centered constraints; we cannot justify constraints if we begin theorizing from the point of view of the universe and begin thinking about value from the 'outside-in'.

Consequently, to vindicate constraints we must approach them from the 'inside-out', and exploit indexical features that appear only from an agent's point of view. As I've argued, it should be unsurprising, then, that capturing constraints using outcomes requires positing evaluative features of outcomes that are not evaluative features of the states that those outcomes represent. Since representations mediate agents' engagement with the world, constituting their point of view on it, it is natural to exploit, for example, beauty, an evaluative feature characteristic of representations, to explain agent-centered constraints from the 'inside-out'.

Recognizing this fault line is especially important because it explains the mistaken insistence that consequentialism is incompatible with constraints. If we assume that every consequentialist theory is an outside-in theory, then constraints and consequentialism are indeed incompatible; constraints appear only from the inside-out. But consequentialism is not essentially an outside-in approach. If we begin consequentialist theorizing from the agent's point of view, in the manner of subjective consequentialism and of *de se* consequentialism in particular, then a consequentialist justification for something like agent-centered constraints is achievable through *de se* value.

Nevertheless, an apparent tension divides *de se* value and the axiology that standardly underpins consequentialism, which often regards value as, roughly, what's desirable 'from the point of

view of the universe' or from the outside-in.⁴⁷ 'The point of view of the universe' offers an appealing perspective on value because it fixes a standard of correctness or objectivity for evaluation. Objectivity might thus seem essential to agent-neutral value. Correspondingly, because *de se* value is invisible from that point of view, its objectivity might be doubted. And if *de se* value lacks objectivity, it might seem like a merely covert form of agent-relative value.

De se consequentialists can respond by appealing to an ancient source of objectivity in ethics: the *phronimos* or practically wise agent. One standard formulation of virtue ethics holds, very roughly, that you should act in some way just when and because the *phronimos* would act that way in your position. Practical wisdom provides the objective standard against which an act's rightness is judged. We can extend that standard to ground *de se* value's objectivity by holding that one outcome is agent-neutrally better than another just when and because the *phronimos* would prefer the first to the second, were she (*per impossibile*) in your shoes and knew only as much as you do now. To be sure, response-dependent accounts of value like this face well-known challenges, such as the Euthyphro. But these challenges are no more daunting than the challenge of explicating what 'the point of view of the universe' is. Consequently, *de se* consequentialists can vindicate the objectivity of *de se* value by grounding that value in the point of view of a privileged agent, rather than that of a universe.

6. Conclusion

De se consequentialism employs agent-neutral value and agent-centered outcomes to enact constraints. Consequently, it provides a counterexample to the claim that "rankings of outcomes will need to be agent-relative [...] in order to consequentialize certain non-consequentialist theories." To enact constraints, *de se* consequentialism implies that some *de se* outcomes are worse than their third-personal counterparts. This may seem puzzling. Nevertheless, I have argued that it is no more surprising than parallel implications from competing non-consequentialist theories and that it has precedent in a familiar evaluative property, beauty.

⁴⁷ I thank a referee for *Ethics* for pressing me to clarify my position here.

References

- Babb, M. (2016). The Essential Indexicality of Intentional Action. *Philosophical Quarterly*. 66 (264): 439-456.
- Bermudez, J. (2016). *Understanding "I"*. Oxford University Press.
- Bermudez, J. (2017). Yes, Essential Indexicals Really Are Essential. *Analysis*. 77 (4): 690-4.
- Bermudez, J. (2020). *Frame It Again*. Cambridge University Press.
- Brown, C. (2011). Consequentialize This. *Ethics*. 121(4): 749-771.
- Burge, T. (2000). Reason and the First Person. In C. Wright, B. Smith and C. Macdonald (Eds.), *Knowing Our Own Minds*. Oxford University Press.
- Cappelen, H. and Joshua Dever. (2013). *The Inessential Indexical: On the Philosophical Insignificance of Perspective and the First Person*. Oxford University Press.
- Castañeda, H-N. (1971). Intentions and the Structure of Intending. *Journal of Philosophy*. 68 (15): 453-66.
- Chalmers, D. (2011). Frege's Puzzle and the Objects of Credence. *Mind*. 120: 587-635.
- Chierchia, G. (1990). Anaphora and attitudes *De Se*. In R. Bartsch et al. (Eds.), *Language in action*. Dordrecht: Foris Publications.
- Crisp, R. (2005). Value, reasons and the structure of justification: how to avoid passing the buck. *Analysis* 65: 80-85.
- Darwall, S. (1986). Agent-Centered Restrictions from the Inside-Out. *Philosophical Studies*.
- Dreier, J. (1993). Structures of Normative Theories. *The Monist*, 76, 22-40.

- Dreier, J. (2011). In Defense of Consequentializing. In M. Timmons (Ed.), *Oxford Studies in Normative Ethics* (Vol. 1, pp. 97-119). Oxford: Oxford University Press.
- Dreier, J. (2020). Two Models of Agent-Centered Value. *Res Philosophica* 97 (3): 345-362.
- Driver, J. (2001). *Uneasy Virtue*. Cambridge University Press.
- Dougherty, T. (2013). Agent-neutral deontology. *Philosophical Studies*. 163: 527-537.
- Fitts, J. (2013). Chalmers on the Objects of Credence. *Philosophical Studies*. 170: 343–358.
- Foot, P. (1983). Utilitarianism and the Virtues. *Proceedings and Addresses of the American Philosophical Association* 57:273-283.
- Frege, G. (1956). The Thought: A Logical Inquiry. *Mind*. 65 (259): 289-311.
- Hammerton, M. (2017). Is agent-neutral deontology possible? *Journal of Ethics and Social Philosophy* 12 (3):319-324.
- Hammerton, M. (2020). Deontic Constraints are Maximizing Rules. *The Journal of Value Inquiry*.
- Hurka, T. (2003). Moore in the middle. *Ethics*. 113: 599-628.
- Hurley, P. (2013). Consequentializing and Deontologizing: Clogging the Consequentialist Vacuum. *Oxford Studies in Normative Ethics* (Vol. 3, pp. 123-153). Oxford: Oxford University Press.
- Hurley, P. (2017). Why Consequentialism’s ‘Compelling Idea’ Is Not. *Social Theory and Practice* 43: 29-54.
- Hurley, P. (2019a). New Consequentialism and the New Doing-Allowing Distinction. In Christian Seidel, Ed., *Consequentialism: New Directions, New Problems*. Oxford: Oxford University Press: 176–97.

- Hurley, P. (2019b). Exiting the Consequentialist Circle: Two Senses of Bringing it About. *Analytic Philosophy* 60 (2): 130–163.
- Lenman, J. (2000). Consequentialism and Cluelessness. *Philosophy & Public Affairs*. 29: 342-370.
- Lewis, D. (1979). Attitudes *De Dicto* and *De Se*. *Philosophical Review*. 88(4): 513-543.
- Liao, S. (2012). “What Are Centered Worlds?” *Philosophical Quarterly* 62(242): 294-316.
- Lillehammer, H. (2007). *Companions in Guilt: Arguments for Ethical Objectivity*. Palgrave Macmillan.
- Louise, J. (2003). Relativity of Value and the Consequentialist Umbrella. *The Philosophical Quarterly* 54 (217): 518-536.
- Jackson, F. (1991). Decision-Theoretic Consequentialism and the Nearest and Dearest Objection. *Ethics* 101 (3):461-482.
- Magidor, O. (2015). “The Myth of the *De Se*.” *Philosophical Perspectives* 29: 249-283.
- Mahtani, A. (2017). The Ex Ante Pareto Principle. *Journal of Philosophy* 114 (6): 303-323.
- Mason, E. (2004). Consequentialism and the principle of indifference. *Utilitas* 16 (3): 316-321.
- Mason, E. (2008). Consequentialism and Moral Responsibility. In Seidel, C., Ed. *Consequentialism: New Directions, New Problems?* Oxford University Press.
- McDowell, J. (1998). *Mind, Value and Reality*. Cambridge: Harvard University Press.
- McNaughton, D. and Rawling, P. (1992). Honoring and Promoting Values. *Ethics* 102 (4):835-843.
- McNaughton, D. and Rawling, P. (1995a). Value and Agent-Relative Reasons. *Utilitas* 7 (2):31-47.
- McNaughton, D. and Rawling, P. (1995b). Agent Relativity and Terminological Inexactitudes. *Utilitas* 7 (2):319-25.

Moltmann, F. (2005). Generic one, arbitrary PRO, and the first person. *Natural Language Semantics* 14:257–281

Muñoz, D. (2020). The Rejection of Consequentializing. *Journal of Philosophy*.

Nair, S. (2014). A fault line in ethical theory. *Philosophical Perspectives* 28 (1): 173–200.

Nair, S. (2020). Fault lines in ethical theory. In Portmore, D., Ed. *Oxford Handbook of Consequentialism*.

Perry, J. (1979). The Problem of the Essential Indexical. *Philosophical Review* 86: 874–97.

Peterson, M. (2013). *The Dimensions of Consequentialism: Ethics, Equality, and Risk*. Cambridge: Cambridge University Press.

Peterson, M. (2010). A Royal Road to Consequentialism? *Ethical Theory and Moral Practice*, 13, 153-169.

Portmore, D. (2001). McNaughton and Rawling on the Agent-relative/Agent-neutral Distinction. *Utilitas* 13 (3): 350-356.

Portmore, D. (2005). Combining Teleological Ethics With Evaluator Relativism: a Promising Result. *Pacific Philosophical Quarterly* 86, 95-113.

Portmore, D. (2007). Consequentializing Moral Theories. *Pacific Philosophical Quarterly*, 88, 39-73.

Portmore, D. (2009). Consequentializing. *Philosophy Compass*, 4, 329-347.

Portmore, D. (2011). *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford: Oxford University Press.

Portmore, D. (2014). Replies to Gert, Hurley, and Tenenbaum. *Philosophy and Phenomenological Research* 88: 241-255.

Railton, P. (1984). Alienation, Consequentialism, and the Demands of Morality. *Philosophy and Public Affairs* 13 (2): 134-171.

Scanlon, T. (1998). *What We Owe to Each Other*. Harvard University Press.

Scheffler, S. (1985). Agent-Centred Restrictions, Rationality, and the Virtues. *Mind* 375: 409-419.

Schroeder, M. (2006). Not So Promising After All: Evaluator-Relative Teleology and Common-Sense Morality. *Pacific Philosophical Quarterly* 87: 348–356.

Schroeder, M. (2007). Reasons and Agent-Neutrality. *Philosophical Studies* 135 (2): 279-306.

Schroeder, M. (2008). Teleology, agent-relative value, and ‘good’. *Ethics* 117: 265-295.

Schroeder, M. (2010). Value and the right kind of reasons. *Oxford Studies in Metaethics* 5: 25-55.

Schroeder, M. (2011). *Oughts*, agents, and actions. *Philosophical Review* 120 (1): 1-42.

Schroeder, S. A. (2017). Consequentializing and its Consequences. *Philosophical Studies* 174 (6):1475-1497.

Setiya, K. (2018). Must consequentialists kill? *Journal of Philosophy* 115: 92-105.

Smith, M. (2003). Neutral and relative value after Moore. *Ethics* 113: 576-598.

Smith, M. (2005). Metaethics. In F. Jackson and M. Smith (eds.), *Oxford Handbook of Contemporary Philosophy*. Oxford: Oxford University Press, 3-30.

Smith, M. (2009). Two kinds of consequentialism. *Philosophical Issues* 19: 257-272.

Stanley, J. and Williamson, T. Knowing How. *Journal of Philosophy* 98: 411-444.

Suikkanen, J. (2008). Consequentialism, constraints and the good-relative-to: a reply to Mark Schroeder. *Journal of Ethics & Social Philosophy* 3: 1-8.

Suikkanen, J. (2020). Consequentializing Moral Dilemmas. *Journal of Moral Philosophy* 17 (3): 261-289.

Tenenbaum, S. (2014). The Perils of Earnest Consequentializing. *Philosophy and Phenomenological Research* 88: 233-240.

Wiggins, D. (1987). *Needs, Values, Truth*. Oxford: Blackwell.

Wright, C. (1988). Moral Values, Projection and Secondary Qualities. *Proceedings of the Aristotelian Society, Supplementary Volumes* 62: 1-26.