

PATERNALISM AND EXCLUSION

Kyle van Oosterum

SOME PHILOSOPHERS believe that the distinctive wrong of paternalism has to do with taking a paternalizee's well-being as a *reason* for one's action.¹ This belief serves as a starting point for what I will call the *exclusionary strategy* for determining the wrongness of paternalism. The exclusionary strategy aims to show how some feature of the paternalizee's normative situation morally excludes acting for the paternalizee's well-being or benefit. In this paper, I explain what is wrong with the exclusionary strategy and offer an alternative "nonexclusionary" approach.

Before proceeding, I wish to highlight (and perhaps disappoint some readers in the process) that I will pay comparatively little attention to what *paternalism* means. That question merits its own paper and indeed has generated its own literature.² That being said, it will be helpful to have a rough idea of the phenomenon I have in mind. A useful starting point might be Gerald Dworkin's three *jointly sufficient* conditions for paternalistic intervention:

1. *Interference Condition*: An act *Z* (or its omission) interferes with the liberty or autonomy of *Y* (the paternalizee).
2. *No-Consent Condition*: *X* (the paternalizer) does so without the consent of *Y*.
3. *Improvement Condition*: *X* does so just because doing *Z* will improve the welfare of *Y* (where this includes preventing his welfare from diminishing) or in some way promote the interests, values, or good of *Y*.³

1 Groll, "Paternalism, Respect, and the Will"; Enoch, "What's Wrong with Paternalism"; and Parry, "Defensive Harm, Consent, and Intervention" and "What's Wrong with Paternalism?"

2 Kleinig, *Paternalism*; Feinberg, *Harm to Self*; Coons and Weber, "Introduction"; Dworkin, "Defining Paternalism"; and Bullock, "A Normatively Neutral Definition of Paternalism."

3 Dworkin, "Paternalism" and "Defining Paternalism."

All of these conditions have been criticized in some way, and other characterizations of paternalism offer useful refinements.⁴ Nevertheless, these conditions—or some suitably refined version of them—are often invoked not just in lay conversations about paternalistic policymaking but also in moral debates on the oft-assumed wrongness of paternalism. Consider an example where all three conditions come into play:

Fried Chicken: Frida is a normal adult who wants to eat delicious but unhealthy fried chicken. Her local government, being motivated by a concern for the physical well-being of its constituents, has decided to implement a tax on fried foods.

On Dworkin's account, Frida's government has discouraged her consumption of fried foods and, in so doing, interfered with her liberty, without her express consent, to improve her well-being. This seems like an instance of paternalistic intervention. Note, however, that this definition does not accommodate the assumption that paternalism is presumptively morally wrong. Indeed, in the example above, it might not be crystal clear whether the government in question has acted wrongly. This reflects Dworkin's assumption that we should generally prefer normatively *neutral* definitions and not smuggle in evaluative judgments about the concept we are defining unless, by not including those judgments, we fail to represent it adequately.⁵

Now, I will not take a position on Dworkin's methodological assumption, but the third condition in his definition of paternalism will be essential for what follows. This is because the group of philosophers initially mentioned believe that the distinctive wrong of paternalism has something to do with the "because" part of that improvement condition. If these philosophers are right, they will have vindicated the idea that part of our concept of "paternalism" consists of its pointing to a presumptively problematic practice.

The structure of my paper is as follows. In section 1, I spell out the details of the exclusionary strategy and its motivations. To set up my critique, I distinguish between two versions of the exclusionary strategy by borrowing from the literature on exclusionary reasons. The appeal to second-order exclusionary reasons (i.e., reasons not to act on our first-order reasons) offers a good way of characterizing views that fall under the exclusionary strategy. Section 2 tackles the "justificatory" version of the exclusionary strategy before turning to the "motivational" version. After examining several problems for how to

4 Shiffrin, "Paternalism, Unconscionability Doctrine, and Accommodation"; Grill, "The Normative Core of Paternalism" and "Antipaternalism as a Filter on Reasons"; and Groll, "Medical Paternalism."

5 Dworkin, "Paternalism."

develop these views plausibly, I turn in section 3 to a brief sketch of an alternative approach to determining the wrongness of paternalism. I argue that my nonexclusionary approach is a better way of obtaining the appealing aspects of the exclusionary strategy and cohering with the mainstream view that paternalism is *pro tanto* wrong. As Christian Coons and Michael Weber put it:

Normative debates about paternalism . . . don't usually concern *whether* it is problematic but *how* problematic it is. . . . There is always at least some *pro tanto* reason to avoid it.⁶

In this paper, I accept that paternalism is *pro tanto* wrong as this view is assumed (sometimes explicitly) by proponents of the exclusionary strategy. Of course, a good philosophical argument may convince us that paternalism is never permissible, but the justificatory bar for this will be high. Nevertheless, I show that both versions of the exclusionary strategy are inconsistent with this mainstream view, *contra* what its defenders claim. This is a surprising result that again motivates consideration of an alternative view that can accommodate the mainstream view. Correspondingly, this paper argues that the exclusionary strategy is problematic while suggesting a more familiar route for determining what makes paternalism wrong. Construing our normative reasons against paternalistic intervention in an exclusionary, second-order way creates many of the problems I cover in section 2. Instead, I argue that we have *first-order* reasons for and against intervention and that their weights can be discerned and balanced against one another to determine the wrongness of paternalism. As such, an overarching aim of my paper is to show that appeals to exclusionary reasons generate implausible implications and are unnecessary in debates concerning the (*pro tanto*) wrongness of paternalism.

1. THE EXCLUSIONARY STRATEGY

The exclusionary strategy, as I have called it, has been defended explicitly by at least three philosophers.⁷ These philosophers differ subtly in how they motivate and conceptualize the exclusionary strategy, but they can be grouped roughly into two subcategories. Borrowing from the literature on exclusionary reasons, we can say that there are *motivational* and *justificatory* interpretations of exclusion.⁸ Though not every proponent of the exclusionary strategy uses

6 Coons and Weber, "Introduction," 2–4.

7 Groll, "Paternalism, Respect, and the Will"; Enoch, "What's Wrong with Paternalism"; and Parry, "Defensive Harm, Consent, and Intervention" and "What's Wrong with Paternalism?"

8 Adams, "In Defense of Exclusionary Reasons."

the same terminology, the mechanisms described are essentially that of Razian exclusionary reasons.⁹ Where David Enoch and Daniel Groll defend a motivational account of exclusion, Jonathan Parry defends a justificatory account of exclusion to explain the wrongness of paternalism. I will outline these two views of exclusion before turning to specific problems in the next section.

1.1. *The Motivational Account*

The motivational account of the exclusionary strategy focuses on the reasons *for which* an agent may act. Enoch and Groll both appeal to the idea of exclusionary reasons for action. An exclusionary reason is a reason *not* to act on some reason; it defeats or “excludes” a *first-order* reason to do some action but does not outweigh it.¹⁰ Let us take Joseph Raz’s “Ann the Banker” example to see how these types of reasons function in everyday deliberation about what we should do. Ann is a banker who, exhausted after a long day of work, nevertheless has to make an important decision about some financial deal. The fact that she is exhausted seems to give her a reason *not* to act on her best judgment of the reasons for and against making this important investment. In Raz-speak, Ann has an exclusionary reason.

With more of a grip on the concept of an exclusionary reason, I will introduce the context behind its specific application in these debates. One common thread among liberal or antipaternalist philosophers is their assertion that the *motive* behind paternalistic intervention is essentially *insulting* to the paternalizee, or potential target of our intervention.¹¹ It is not hard to see why they might think this. When a paternalist is motivated in this way, they believe they know what is best for a person, perhaps better than that person does themselves. This seems problematic insofar as it lines up with another popular liberal idea, which is that the individual essentially knows what is best for them. It is not in anyone else’s moral jurisdiction, if you will, to interfere with their choices unless they harm other people.¹² Perhaps then, if there is something wrong with paternalistic intervention, it resides in the *negative* beliefs and judgments we have about people’s choices and whether those are good for them to make.

However, there will be cases where a potential paternalizee does not know what is good for them. Enoch believes (and I agree) that there is nothing wrong with simply having a true belief about whether a paternalizee’s actions will

9 Raz, *Practical Reason and Norms*.

10 Clarke, “Exclusionary Reasons”; and Raz, *Practical Reason and Norms*.

11 Feinberg, *Harm to Self*; Shiffrin, “Paternalism, Unconscionability Doctrine, and Accommodation”; Quong, *Liberalism without Perfection*; Begon, “Paternalism”; Cholbi, “Paternalism and Our Rational Powers.”

12 Mill, *On Liberty*.

diminish their own well-being. If the paternalist knows that the paternalizee's choices will cause the paternalizee harm, what could be wrong with simply holding that belief? Enoch's suggestion is that the wrongness may consist of a paternalist being motivated to *act* on this belief about the paternalizee's choices.

This is where exclusionary reasons come into the picture. Let us take Enoch's example, which he borrows from Jonathan Quong. Your friend wants to borrow money that you are sure he will use to make himself worse off (perhaps by buying drugs). If you simply believe that he will use the money in a bad way, and you are probably right, there does not seem to be anything wrong with that. Where the wrong lies, Enoch argues, is in acting on that belief and ignoring your friend's questioning ("What's it to you what I do with this money?"), because in so doing you deny the value of your friend's autonomy over their own life.¹³ If there is something wrong with paternalizing here, it is because your friend's autonomy gives you an exclusionary reason *not* to act for the reason that it would be good for their well-being if you did not give them the money.

In a similar vein, Groll takes paternalism to be wrong because of how it treats the *will* of the potentially paternalized individual. Roughly, the idea is that a paternalizee's will is intended to silence, trump, or exclude the "reason-giving force" of the other considerations that might be at play when one (a potential paternalizer) is practically deliberating about what to do on behalf of the paternalizee.¹⁴ Groll imagines a medical scenario where a doctor performs some operation and considers a patient's wish not to have the operation as an ingredient in her deliberation about what would be good for the patient's well-being. Groll points out that the patient might be annoyed with the doctor's construal of their will as part of *her* deliberation and not itself the *decisive* factor about whether or not to perform the operation. In other words, as Groll puts it, the patient's will should have made "irrelevant" questions about whether it is good for them to have such an operation.

On both of these accounts, the thought is that a potential paternalizer acts wrongly in being motivated solely (or overridingly) by considerations of a paternalizee's good. They hold that the paternalizee's autonomy or will morally excludes such considerations as reasons for action. Importantly, neither Groll nor Enoch believes that paternalism is always wrong, and each has suggested

13 It is unclear whether Enoch has his own specific conception of autonomy in mind. See Enoch, "Hypothetical Consent and the Value(s) of Autonomy." For our purposes, we can interpret it broadly as a person's ability to make decisions in line with their values or conception of the good life. See Birks, "How Wrong Is Paternalism?"

14 Groll, "Paternalism, Respect, and the Will," 701.

that his is an account of the *pro tanto* wrongness of paternalism.¹⁵ Each maintains that paternalism is usually but not always wrong and believes that his account of the exclusionary strategy can vindicate that judgment. Recall that such a view is the mainstream one in philosophical writing about paternalism. It would be interesting if it turned out that their account was not consistent with this view (more on that in section 2.2).

1.2. *The Justificatory Account*

Let us turn now to the justificatory account. Parry's views on the wrongness of paternalism exist in the larger context of defensive harm, but I believe they fit well under the banner of the exclusionary strategy. Like Enoch and Groll, Parry is trying to figure out why it can be wrong to (paternalistically) act for someone's good or well-being. His response to this question appeals to the idea of a *moral power*, that is, the ability persons possess to change the moral or normative landscape around them (e.g., by changing what it is permissible to do to them).

For example, when a person consents to sexual intercourse, they make what is usually impermissible—another person interfering with their bodily integrity—into something permissible. Parry believes that just as we have the power to control our bodies and property (our material resources), we also have the power to control the use of our “good,” where “good” refers to reasons grounded in our well-being (our “normative resources”).¹⁶ To use someone's good, he claims, is to justify one's actions by appealing to the fact that it would be good for them if we did that. Let us return to the example offered by Quong above. Your friend has the power to make their good “inadmissible” as a justifying reason for action, such that declining to give them the money cannot be justifiable (for the reason that it would be good for them).

Notice that talk of the inadmissibility of a reason sounds very similar to the exclusionary reasons mentioned by Groll and Enoch. To my mind, it sounds similar because Parry is defending a justificatory account of exclusion. A justificatory account of exclusion holds that exclusionary reasons essentially change the “right-making” features of an action; they exclude or prevent ordinary moral reasons from standing in a justifying relation to actions.¹⁷ Let us consider a nonpaternalistic example of this phenomenon. Adams argues that laws can be thought of as (exclusionary) reasons that exclude reasons that might count

15 Groll's recent views on the wrongness of paternalism seem to involve much more rights talk than talk of exclusionary reasons. See Groll, “Paternalism and Rights.”

16 Parry, “What's Wrong with Paternalism?”

17 Moore, “Authority, Law and Razian Reasons”; and Adams, “In Defense of Exclusionary Reasons.”

in favor of law-breaking, such as pulling over on a highway to help a wounded animal.¹⁸ A distinguishing feature of exclusionary reasons is that they do not compete in weight with first-order reasons and generally have absolute priority over the reasons that they exclude.¹⁹ But if this is true, then even though we could have incredibly weighty reasons to help the animal, the law makes those seemingly weighty reasons play *no* justificatory role whatsoever in our deliberation. To my mind, the same thing is going on in Parry's account of the wrongness of paternalism. As he puts it, reasons to promote a person's well-being become "unavailable" as justifications for action by virtue of an exercise of our moral power (to exclude the use of our "good").²⁰

At this point, it might be helpful to distinguish between motivational and justificatory exclusionary strategies. The motivational account locates the wrong of paternalism in the well-being-related reasons that a paternalizer chooses to act on. Autonomy (or the will, in Groll's account) provides an exclusionary moral reason for the paternalizer not to act for the good of the paternalizee. The justificatory account makes no reference to a paternalizer's motivations for action. Instead, it focuses on how features of the situation make well-being-related reasons the wrong sort of reason to act on. This is because they are no longer part of the potential right-making reasons for justifying action. Whereas the justificatory account denies the ordinary justificatory role that well-being-related reasons play, the motivational account does not make this claim about reasons. Well-being-related reasons exist as strong reasons to act on, but it so happens that respect for autonomy or the will makes it so that such reasons are wrong to be motivated by. In short, the wrong lies either (i) in acting on a reason that no longer performs its function (the justificatory account) or (ii) in acting on a wrong yet functional moral reason (the motivational account).

2. PROBLEMS FOR THE EXCLUSIONARY STRATEGY

In this section, I will argue that both versions of the exclusionary strategy are problematic. I will show that both views struggle to accommodate the mainstream view of paternalism's *pro tanto* wrongness that also counts against them. Upon close examination, the justificatory account, while clearly specifying how well-being is to be excluded, delivers counterintuitively strong verdicts that seem never to countenance paternalistic intervention (when it seems

18 Adams, "In Defense of Exclusionary Reasons."

19 Raz, *Practical Reason and Norms*.

20 Parry, "What's Wrong with Paternalism?"

permissible). By contrast, the motivational account enjoys some intuitive advantages over the justificatory account, but it is unclear how to specify its exclusion of well-being in a plausible way. In section 3, I offer a general diagnosis of why these views go wrong, as well as an alternative view that outperforms them both. For now, if neither exclusionary account turns out to be plausible, this supports my contention that it is unnecessary and implausible to appeal to exclusionary reasons to explain the *pro tanto* wrongness of paternalism. This is because there may be alternative views, such as my own highlighted in section 3.2, that can vindicate much of the exclusionary strategy's appeal without a second-order level of reasoning and without the problems that such reasoning gives rise to.

2.1. Problems with Justificatory Exclusion

An important caveat to Parry's moral power account is that a person has to be able to *competently* refuse to be benefitted by others. More precisely, a person has to competently exclude the use of their good as a justification for someone's action toward them. This is a principled qualification inspired most likely by the oft-cited distinction between *soft* and *hard* paternalism.²¹ Though that distinction has come under fire, the thought is plausible enough: paternalistically interfering with someone seems less wrong if a person made their choice involuntarily. This involuntariness could be due to the individual not being an adult yet, being under the influence of alcohol or drugs, or perhaps suffering from some physical or mental ailment. Roughly, soft paternalists believe that whether a person's choice is voluntary is relevant to the justifiability of paternalistic intervention. By contrast, hard paternalists disagree that voluntariness should always matter.²² In practice, correct judgments of voluntariness can be hard to make, but it is *prima facie* plausible to include them as features that help justify paternalistic intervention. In any case, if a person "incompetently" refused a benefit, then this would lead to the intuitive verdict that we could still use their good as a justification for paternalistic action (assuming that such an action would count as "paternalistic" in the first place).

So far, so good. However, we might think cases of incompetent refusal are the low-hanging philosophical fruit for this debate. After all, some philosophers do not regard soft paternalism as a kind of paternalism at all.²³ The real challenge to Parry's justificatory account would be to identify one case where

21 Feinberg, *Harm to Self*.

22 This is an obvious caricature of a sophisticated debate that I am mentioning only to provide context for what follows. For evaluation of the distinction between soft and hard paternalism, see Hanna, "Hard and Soft Paternalism."

23 Feinberg, *Harm to Self*.

a competent refusal has occurred yet a paternalistic intervention would not be wrong. Consider an adapted version of Richard Arneson's famous case:

Pouting Young Adult: Tom is unreasonably distressed at some disappointment he has suffered. Perhaps he has been bested in competition for a job he coveted. . . . Perhaps a particularly charming rabbit he saw at the Humane Society pet adoption center and hoped to choose and make his pet was adopted by another person. Whatever the cause of his distress, he is unhappy, feels vaguely cheated by the world at large, and wants at the moment nothing more than to express his disappointment by committing suicide. In addition, Tom knows he will likely change his mind but right now has no interest in doing so. He is neither mentally ill nor incompetent as a decision-maker. He simply wants to commit suicide and has refused appeals by his friends to change his mind and think of his own well-being.²⁴

To my mind, this is a case where paternalistic intervention seems not only permissible but justified. Of course, a very staunch antipaternalist might just deny that it is intuitively permissible to interfere here. However, it is hard to see how if paternalism were not permitted here it would still be permitted in a similarly extreme case. It seems that the justificatory account, with its notion of a "competent refusal," makes such a paternalistic action unjustifiable. This may lead us to wonder how paternalism toward competent adults could ever be permitted on this account.²⁵

But this is far too quick. Proponents of this justificatory account might appeal to the distinction, captured nicely by David Owens, between *acting wrongly* and *wronging someone*.²⁶ Another way of putting this is that we might think we can commit a wrong without doing the wrong thing. When I break a promise to meet my friend to help another person who has been hit by a car, I have wronged my friend but not done the all-things-considered wrong thing. Here, my promise-breaking is intuitively justified, which suggests, as Owens puts it, that "committing a wrong can be the right thing to do."²⁷ This idea fits in well with "exclusionary reasons" terminology, because one can think of a promise as excluding the reasons *not to act* on or break the promise.

24 Arneson, "Joel Feinberg and the Justification of Hard Paternalism," 278–79.

25 This point has also been noted by Quong in his recent talk on Parry's account of antipaternalism, "Paternalism, Disagreement and Groups."

26 I thank Lorenzo Elijah again for pointing this distinction out to me.

27 Owens, *Shaping the Normative Landscape*, 45.

Perhaps then, in cases such as Pouting Young Adult, advocates of the exclusionary strategy can claim that although we have wronged the paternalizee, we have not acted wrongly in paternalistically interfering (in those extreme cases). I have not attended to this fact: just because someone has been wronged does not mean that what has been done is wrong or impermissible. Exclusionary strategy proponents can claim that permissible paternalistic intervention involves cases of “permissible wrongdoing,” so to speak, which allows them to maintain that paternalism is wrong but not always wrong. In other words, not all paternalistic wrongings are wrongs. This seems like a plausible enough conclusion to hold.

Unfortunately, this appeal is unavailable to proponents of the justificatory account of the exclusionary strategy. Parry’s view renders well-being-related reasons counting in favor of paternalistic intervention disabled or unable to play any justificatory role for action.²⁸ Obviously, according to this view, if a paternalizer were to intervene on the basis of well-being-related reasons, they would naturally wrong the paternalizee. But what makes a paternalistic intervention in Pouting Young Adult “not” wrong? One might think the intervention is intuitively permissible and all-things-considered justified, but the *content* of this intuition and justification is surely the very same well-being-related reason that is disabled by exclusion. If some other non-well-being-related reason forms the intuitive justification for intervention, then we are not plausibly dealing with a case of paternalism anymore. After all, the exclusionary strategy’s account of paternalism relies on the notion that the justification for the intervention is well-being-related (see introduction).

So, we have something of a dilemma. Adherents of the justificatory account cannot defend the idea that a paternalistic wrongdoing would not be wrong. They cannot appeal to well-being-related reasons, and they need those very reasons to be discussing a “paternalistic” act in the first place. In other words, either they must accept that every paternalistic wrongdoing is indeed wrong—an extreme conclusion—or the act of intervention is “not wrong” but can no longer be described as “paternalistic.” Therefore, this Owens-style idea cannot be used to square the justificatory account with the *pro tanto* view of paternalism’s

28 Parry has suggested to me that there could be a positive and a negative way to read his view. On the negative reading, his view states that welfarist reasons are not there to justify the action. On the positive reading, the use of the paternalizee’s welfarist reasons is just a directed wrong to the paternalizee (e.g., a form of trespass). Perhaps a version of his view could be developed with only the positive reading. There are two problems here. First, the negative reading contributes to making it an “exclusionary” view in the first place. Second, and related, one might wonder how distinctive his view would be from other antipaternalist views without this negative claim.

wrongness. The result is that this view, while clear in its formulation, is a counterintuitively strong version of antipaternalism and cannot make room for intuitively permissible cases of paternalism.

Here is another concern with Parry's view. Recall Parry's analogy between material resources and normative resources. The inference drawn from this analogy is that "wrongable" paternalizees determine the moral status of paternalistic intervention because moral reasons *belong* to them. The idea that our reasons "belong" to us is mysterious. I think there is an importantly relevant distinction between claiming that these reasons are *about the paternalizee* and saying that these reasons *are theirs*. The former claim is straightforward and makes sense. After all, some philosophers think that it is wrong to act or be motivated by the reasons that *refer* to a paternalizee's well-being (e.g., what the motivational account of exclusion seeks to defend). The latter claim, namely, that moral reasons (i.e., reasons having to do with well-being) can be ours to control, strikes me as implausible and in need of further defense. Obviously, this taps into a deeper question about whether reasons can be "up to us" in a metaphysical sense that is admittedly not Parry's focus.²⁹ While Parry does offer a number of rationales in favor of having a moral power to exclude reasons, he has not shown that we have this power; in other words, it is still unclear how these reasons are (or become) *ours* in the way that material property is ours.³⁰ For now, this contestable analogy seems to be justifying the existence of this power and our supposed ownership of reasons. Therefore, Parry's argument is not only implausible as an account of paternalism's *pro tanto* wrongness, it seems also to be derived from implausible footings.³¹

2.2. Problems for Motivational Exclusion

2.2.1. A Prima Facie Problem and the Scope of Exclusionary Reasons

I want to suggest that the following insight can be gained from the justificatory version of exclusion: claiming that well-being-related reasons do not feature at all in a moral assessment of paternalistic intervention is unnecessary. It is unnecessary with respect to reaching the conclusion that paternalism is *pro tanto* morally wrong. Indeed, as I have just argued, the justificatory version of exclusion makes it difficult to render any paternalistic intervention permissible.

29 Moore, "Authority, Law and Razian Reasons"; and Chang, "Do We Have Normative Powers?"

30 I believe the strategy Parry pursues is to justify the power in virtue of how it serves the realization of some important value. However, this does not show that the power exists, nor does it dispel the mysterious claims about the ownership of reasons it seems to involve.

31 I thank an anonymous referee for the suggestion to elaborate the point in this way.

It also makes strange claims about our supposed ownership of moral reasons. Instead, advocates of the exclusionary idea could appeal to the motivational version of the exclusionary strategy (MES), which makes neither of those claims. The MES just argues that to be motivated to act on well-being-related reasons is *pro tanto* wrong.

Why would it be wrong to be motivated by these reasons? Part of what it is to respect autonomy (or treat one's will as structurally decisive, in "Groll-speak" now) morally excludes being motivated by what is good for the paternalizee's life. Importantly, exclusionary reasons, in this sense, are reasons for not being motivated in one's actions by certain "valid considerations."³² What seems more intuitive about this account than Parry's is that we are not making the extreme claim that well-being is *not* a valid reason-generating consideration and that it *could not* be part of the justificatory story. Instead, the thought is just that the importance of autonomy overrides or generally takes priority over well-being. The device of an exclusionary reason is one way of articulating that thought. This is how we get to the view that autonomy generates an (exclusionary) reason *not* to act on the reason that it would be good for the paternalizee's well-being to interfere.

But does this view do better in cohering with the verdict that paternalism is only *pro tanto* wrong? Enoch and Groll seem to think so, but I believe there are some ambiguities in their account that make this question difficult to answer affirmatively. The chief ambiguity consists in how *much* this exclusionary reason excludes. At the moment, the view looks like this:

First Pass:

- P1: Paternalistic interferences are wrong if there are unexcluded moral reasons that favor not paternalistically interfering.
- P2: There is an exclusionary reason that is grounded in the paternalizee's autonomy or will. It is an unexcluded moral reason not to interfere for the reason that it would be good for the well-being of the paternalizee (to interfere).
- C: Therefore, paternalistic interferences are wrong.

Of course, this statement of the view is far too general. Without qualification, it would rule out any case of paternalistic intervention (targeted at competent adults). This is because exclusionary reasons are generally thought to have *absolute* priority over the reasons that they exclude.³³ At first glance, this argument holds that autonomy (or the will), being the ground of an exclusionary reason,

³² Raz, *Practical Reason and Norms*, 185.

³³ Raz, *Practical Reason and Norms*.

always has priority over our first-order reasons to promote one's well-being, regardless of this reason's normative strength. Even very staunch antipaternalists will concede that this is a counterintuitively strong conclusion, which is why the widely held view is that paternalism is only *pro tanto* wrong. Now we can return to the question of whether Enoch and Groll's view is actually consistent with this widely held view despite the conclusion of this first-pass argument.

If there is a problem with the first-pass argument, it resides in P₂, which is where some qualifications might be attempted. Perhaps P₂ can instead read:

P₂*: There is an exclusionary reason grounded in the paternalizee's autonomy or will. It provides an unexcluded moral reason not to interfere that is *usually undefeated* by the reason that it would be good for the well-being of the paternalizee (to interfere).

P₂* would allow us to say that there can be some cases where the exclusionary reason can be outweighed or defeated by the reason to act for the well-being of the paternalizee. This would seem to get the motivational account closer to the widely held view, but it unfortunately comes at the cost of distinctiveness. As Raz himself points out:

If [exclusionary reasons] have to compete in weight with the excluded reasons, they will only exclude reasons which they outweigh, and thus lose distinctiveness.³⁴

The problem with the P₂* move is that we lose what makes an exclusionary reason "exclusionary." Exclusionary reasons are reasons that refer to the balance of first-order reasons for performing some action and are *not* supposed to be *part of* that same balance of reasons. In other words, we would simply be saying that autonomy generates a first-order reason not to interfere that is often, but not always, stronger than the first-order reasons well-being gives us to interfere. However, this statement would not be consistent with the motivational account's commitments to the notion of exclusion.

In short, this view seems to fall prey to a dilemma. On the one hand, P₂ gives us a consistent statement of this view, but it generates the counterintuitively strong conclusion of the first-pass argument. On the other hand, P₂* allows these theorists to avoid this conclusion at the cost of a less distinctive view, which no longer seems exclusionary. Clearly, this view's proponents would not go for either horn of the dilemma. They believe they can coherently defend the

34 Raz, *Practical Reason and Norms*, 189.

view that paternalism is *pro tanto* morally wrong with the device of exclusionary reasons. How would they go about avoiding this dilemma?

An important feature that has been underspecified in the motivational account is precisely what *scope* such exclusionary reasons have—or should have, for that matter. What it means for exclusionary reasons to vary in scope is to say that they might exclude all or only some of the reasons that apply to some situation in practical reasoning.³⁵ For example, consider Raz's character Colin, who makes a promise to his wife to decide what to do about their son's education only on the basis of their son's interests. Here, Colin has an exclusionary reason not to act on reasons unrelated to his son's interests. However, the scope of that reason does not extend so far as to exclude considerations of justice to other people. Raz's notion of exclusionary reasons is complicated by, but also more faithful to, the circumstantial nature of practical reasoning because of these *scope-affecting* considerations. Indeed, the complication for practical reasoners consists in determining when these considerations narrow the scope of exclusionary reasons such that they no longer exclude conflicting first-order reasons.

How does this bear on the debate about the wrongness of paternalism? Recall that proponents of the MES only want to defend the *pro tanto* wrongness of paternalism. They may want to accommodate cases where a paternalistic intervention is intuitively permissible, such as *Putting Young Adult*.

Putting Young Adult was a dramatic case chosen to elicit the commonsense intuition that it is *prima facie* permissible to interfere with Tom's autonomous choice. Let us translate the details of the case into the MES framework as follows: Tom's autonomy (or will) generates an exclusionary reason not to act on the first-order reason (that is, that it would be good for his well-being if we prevented his suicide). Now, if we assume that MES proponents want to allow for paternalistic interference in this kind of case, what would they have to say? They could appeal to considerations that *affect* the ordinary scope of exclusionary reasons generated by a paternalizee's autonomy or will. Perhaps the scope of autonomy's exclusionary force might be limited to a paternalizee's non-self-annihilating decisions. So, while autonomy excludes acting for the reason that it would be good for a paternalizee's well-being, perhaps it does not exclude a first-order reason to prevent suicide.

However, reining in the scope of the exclusionary reason in this way is somewhat *ad hoc*, and it forces the MES proponent to unnecessarily defend a general prohibition against suicide. I believe that what is lurking in the background is some concern for Tom's well-being and a belief that it is sometimes permissible to act for such a reason. While we normally want to treat a person's autonomy

35 Raz, *Practical Reason and Norms*, 39.

(or will) as decisive in this exclusionary way, cases such as Pouting Young Adult make us hesitate because so much of Tom's well-being is at stake. However, in the case of your friend asking for money, you might feel more compelled to respect the exclusionary force of his autonomy. I think that the asymmetry between these cases might be explained in this way: a paternalizee's well-being sometimes seems to play the role of an excluded reason and sometimes seems to be unexcluded by their autonomy.³⁶ But how can this first-order well-being-related reason operate in both of these ways? Is there some principled way to distinguish when this well-being-related reason is plausibly excludable or nonexcludable?

2.2.2. *Different Ways to Identify the Scope of the Exclusionary Reason*

The answer to those questions depends on what account of well-being we are operating with. However, I am not convinced that applying any account of well-being could yield a nonarbitrary answer to the second of those questions. Let us plug in each of Derek Parfit's three accounts of well-being, one at a time, to see why this is the case.³⁷ First, objective-list theories claim that there is some list of goods, such as knowledge and friendship, that constitute well-being and make an agent's life good whether or not the agent desires them. This is a crude rendering of this theory, but it suffices for our purposes. Perhaps, using the objective-list theory, the MES proponent might suggest that autonomy excludes some of the goods on the objective list but not others. Those goods that autonomy does not exclude would provide a kind of unexcluded well-being-related reason that helps deal with certain cases of intuitively permissible paternalism.

The problem with this approach is that it will be difficult to determine *which* goods should not be excluded and in *which contexts* this ought to be the case. One general problem for objective-list theories is determining what goods plausibly belong on such a list. Here, we have a similar issue: how do we determine which goods belong on this list *and* how can we create a plausible separation between the excluded and unexcluded well-being reasons to which they give rise? Since the objective-list theory donates its conceptual baggage here, the MES proponent should probably not adopt this as their account of well-being.³⁸

Second, we could apply some form of hedonism to this question. Perhaps there is a *threshold* for the amount of pain to be prevented (or pleasure to be obtained) that could draw the line between excluded and unexcluded

36 This thought was suggested to me by Enoch on an earlier draft of this paper.

37 Parfit, *Reasons and Persons*, app. 1.

38 A further issue might be that implementing the objective-list theory conflicts somewhat with the spirit of autonomy's exclusionary scope. It might be strange that certain objective goods that I do not think are objectively good play some role in deciding when I am wronged by paternalistic intervention.

well-being-related reasons.³⁹ In Pouting Young Adult, we could say that Tom's death, being the ultimate loss of well-being, renders this decision unexcluded by his autonomy. Since this pain would surpass some threshold, it would be outside the scope of the autonomy-related exclusionary reason and thus defeat said reason. As a result, we could obtain the verdict that paternalistic intervention in that extreme case would not be wrong. For this to be consistent with the *pro tanto* view of paternalism's wrongness, the threshold would have to be very high. I think that this is certainly more plausible than applying the objective-list theory here.

However, I am skeptical that a threshold approach identifies the right scope-affecting consideration for this exclusionary reason. My first concern is about how high the threshold should actually be. To my mind, the threshold approach seems more intuitively appealing the more ambiguously it is defined. Let us say the threshold was defined by the potential death of a paternalizee. One might think that though this is a concrete specification of the threshold, it seems somewhat arbitrary. Why should excruciating pain not satisfy the threshold? When the threshold is high yet ambiguously defined, this will lead to a lot of disagreement about if and when the threshold applies. Perhaps the MES proponent might reply that this is fine, because it mirrors the real-life complexities of practical reasoning about paternalistic intervention. However, insofar as this approach is used to try to distinguish between excluded and unexcluded well-being reasons, it raises more questions than it was intended to answer.

My second concern is that this hedonistic threshold-based approach might, depending on how we characterize it, start to resemble the "objectivist" tendencies of the objective-list theory. This is because the justification for a well-being threshold does not originate in the paternalizee themselves and seems to imply the view that pain or pleasure is worth avoiding or pursuing whatever else the paternalizee might want. No doubt this can be a plausible point of view, but the point of invoking exclusionary reasons is largely to bring such matters under the normative auspices of the paternalizee. That is, we want to let them determine the amount of pain and pleasure they want to receive over the course of their life. So, externally defining well-being thresholds for exclusionary reasons to apply seems troubling and inconsistent with the motivations for the exclusionary strategy.

Finally, we could try some kind of desire-satisfaction theory of well-being. Now, there are many different variants of this theory, so in principle, there are many ways MES proponents could deploy it. Perhaps, they could claim that there are *certain* desires whose satisfaction is not conducive to promoting well-being and that those desires might not fall within the scope of an

39 I thank Lorenzo Elijah for this way of formulating the point.

exclusionary reason. I think this move is already off the table, as we considered it *ad hoc* to rely on ruling out the desire to commit suicide in Pouting Young Adult as an unexcluded well-being reason. In general, it may appear arbitrary to rule out the satisfaction of certain desires just to obtain the intuitively right verdicts about cases.

Instead, we could rein in the scope of the exclusionary reason not by referring to certain desires but to certain *kinds* of desires. Perhaps uninformed desires would not be excluded by autonomy and thus permit paternalistic intervention, whereas informed desires ought to be excluded. The distinction between excluded and unexcluded well-being reasons could just be based on the distinction between the satisfaction of informed and uninformed desires. Again, I think that Pouting Young Adult shows that even on an informed desire-satisfaction theory of well-being, there seems to be some intuitively permissible well-being reason to act on and be motivated by. Arguably then, this way of identifying which well-being-related reasons are excludable fails as well.⁴⁰

In short, the MES cannot be given an articulation to accommodate the *pro tanto* view of paternalism's wrongness. On three plausible ways one could distinguish between excluded and unexcluded well-being-related reasons, the result was that the approaches were either arbitrary or counterintuitive. Though the motivational version of exclusion did not adopt the extreme approach of ruling out the justifying force of well-being-related reasons (as Parry's account seemed to do), it unfortunately could not neatly accommodate them into its framework.

3. MOVING AWAY FROM EXCLUSION: A SKETCH

3.1. Reflecting on Exclusion

Clearly, the proponents of the exclusionary strategy believe that we need to maintain the standard view that paternalism is often but not always morally wrong. The appeal to the normative exclusion of a paternalizee's well-being was thought to be one way to do this, but I have shown that neither version of the exclusionary strategy can be spelled out easily. There is something wrong with treating exclusion as a *constitutive* feature of the wrongness of paternalism rather than one that may explain paternalism's wrongness in some circumstances. It

40 Parry has suggested to me that we could fix the scope of exclusion in a simpler way without discussing different conceptions of well-being. For example, we might think that only a certain *quantity* of well-being can be excluded or only a certain *proportion* of well-being can be excluded. While these would be simpler, it is unclear to me how these views would differ from a threshold account once they are fully elaborated.

simply does not seem like we (always) wrong someone by taking their well-being as a reason for our action (or that we are acting on a reason that no longer plays any justificatory force for action). Another way to put what is going wrong here is to echo Scanlon's observation that invoking exclusionary reasons leads us to ignore the "substantive relevance" of the reasons we are excluding.⁴¹ These are reasons that have to do with a paternalizee's well-being. Such reasons are ordinarily good ones to be motivated by or justifying of action. However, it is possible that in the cases Parry, Groll, and Enoch identify, those reasons are not permissible to act on but perhaps only within a "nonexclusionary" framework.

It is worth stating what the exclusionary view gets right before considering an alternative way of accounting for the wrongness of paternalism. First, the exclusionary strategy can support intuitive verdicts about wrongful paternalism, as in Quong's money-lending case. Second, we might think, as Enoch does, that exclusion generates the correct moral phenomenology associated with paternalism. That is, when paternalizers act, they get involved in what is (morally) not their business, which makes it difficult to justify such actions in a way that is consistent with respecting the other person's autonomy.

So, the exclusionary strategy has these sorts of things going for it. However, the thrust of my paper suggests that going down this route is philosophically costly and onerous. The natural thing to do is to develop an alternative philosophical account—that is, an account that obtains the goods listed above and the verdict that paternalism is *pro tanto* wrong but does without talk of exclusionary reasons and the problems created by the exclusionary strategy. Importantly, this is not to say that we do away with reasons-talk for the wrongness of paternalism, but that we adopt a more familiar approach of reasoning on the *first-order* level. I call this the nonexclusionary approach.

3.2. *The Nonexclusionary Approach*

The view I have in mind is moderate without conceding too much to a position identified by Jason Hanna as "pro-paternalistic."⁴² Like Hanna, I think it is always a valid reason-generating consideration to act in someone's best interest or for the promotion of their well-being. Of course, just because that reason might be valid to act on does not mean that it will be *decisive* in all or even many cases. The idea on the table, then, is that well-being-related reasons (to paternalistically interfere) will normally vary in strength or weight. They will act in competition with reasons to *refrain* from interfering, which might be

41 Scanlon, "Reasons," 241.

42 Hanna, *In Our Best Interest*, 1.

autonomy or will-related (or some other antipaternalistic unit of concern).⁴³ The idea of balancing our (first-order) reasons for and against paternalistically interfering is not unfamiliar to the literature. However, remarkably little has been said about how to discern the strength or weights of these reasons. To that end, I think it would be helpful to turn to another idea in the literature on practical reasoning: *modifiers*.

Modifiers are facts that, though not themselves reasons, are capable of directly affecting the weight of a reason for action.⁴⁴ For example, imagine you have a desire to eat Kentucky Fried Chicken (KFC) which can plausibly give you a reason to go eat some KFC right now. However, the fact that it is rush hour and there will be traffic on the way to KFC might make you less keen to go eat some KFC now. Traffic is not itself a reason to not eat KFC, but it appears to *weaken* your reason to go eat KFC now, given that you do not want to spend so much time in traffic.

Modifiers come in two varieties. The example outlined above displays an *attenuator* in action, a fact that weakens the weight of a reason to do something. By contrast, *intensifiers* are facts that increase the weight of a reason to do something. For example, imagine you are walking around and notice a person who needs help. The fact that this person needs help presumably gives you some reason to help her. But the fact that you are the *only* person around who can help seemingly strengthens your reason to help.⁴⁵ That you are the only person around is not itself a reason to help, and the same would be true if you were one of many bystanders. However, that you are the only person around “intensifies” the weight of your reason to help if (and when) this reason exists. In short, modifiers can affect the weight of our reasons and play an important role in helping us to decide what action we should take or are justified in taking.

In the context of paternalism, there might be all sorts of facts that strengthen or weaken both our first-order reasons to paternalistically interfere and our reasons to refrain from interfering. For example, the amount of well-being that could be promoted (or prevented from being diminished) might intensify a reason to interfere. The significance of one’s autonomous choice might also modify the strength of a reason to refrain from interfering. Another potential modifier might be the *closeness* of the relationship a paternalizer has to a prospective paternalizee. Perhaps the more intimately related paternalizers are with paternalizees (i.e., paternalism between friends), the stronger a reason

43 This move is currently being considered by other philosophers too, and my sketch gestures at ways in which it can be made more precise. See Shafer-Landau, “Liberalism and Paternalism”; and Birks, “How Wrong Is Paternalism?”

44 Bader, “Conditions, Modifiers, and Holism.”

45 Dancy, *Ethics without Principles*.

becomes to paternalistically interfere.⁴⁶ There may be many more kinds of modifiers, and much more could be said in defense of these particular ones. I think these are helpful enough heuristics for discerning the strength or weights of these reasons in a variety of cases.

So much for my view. But what do we get when we couch the wrongness of paternalism simply in the terminology of first-order reasons and their modifiers? First, I believe we can already get intuitive verdicts about cases such as Pouting Young Adult. What we have there is a conflicting well-being-related reason to interfere and an autonomy-related reason to refrain from interfering. However, the well-being-related reason seems intensified by the amount of well-being at stake (i.e., the rest of Tom's possibly good life). Conversely, the autonomy-related reason seems attenuated by the fact that, by Tom's own lights, the choice does not seem that significant to him. Ergo, the well-being-related reason defeats the autonomy-related reason, which matches our intuitive verdict about this case being one of permissible paternalism.

Second, though paternalism is permissible here, we can still obtain the verdict that proponents of the exclusionary strategy want, namely, that paternalism is *pro tanto* wrong. In fact, the wrongness can still be tied to well-being. We should not claim that it is wrong to justify one's action on the basis of well-being (or be motivated by such a reason) because such reasons are excluded. Rather, the wrongness consists in acting on a well-being reason that has been *defeated* because it is weaker than a reason to refrain from interfering. To make this view consistent with the thought that paternalism is *pro tanto* wrong, one need only show how such well-being-related reasons might generally be weaker. They can appeal to a variety of the modifiers I suggested above to justify such a judgment. Crucially, we obtain a view of paternalism's *pro tanto* wrongness without the strong and counterintuitive commitments of the exclusionary views. That is, we need not say that these reasons belong to us, that some reasons can be made to have no justifying force, or that it is always wrong to be motivated by a certain class of reasons.

But what about the phenomenological point? Can this nonexclusionary approach still accommodate those strong (but not absolutely strong) antipaternalistic intuitions? One might worry that a paternalizer's determination of reason strengths and balancing of reasons already violates the "not your business" connotations of valuing a person's autonomy.⁴⁷ Another way to put it is that there is an important distinction between *recognizing* a conflict of reasons

46 There is an active debate in the paternalism literature about whether there is a morally relevant difference between paternalism that is practiced by one's intimates or by the state. See Tsai, "Paternalism and Intimate Relationships"; and Birks, "Sex, Love, and Paternalism."

47 This is another point made to me by David Enoch on an earlier draft of a paper.

and *imposing* a view as to how this conflict between reasons should be adjudicated.⁴⁸ Of course, there is no way to avoid an imposition about how to resolve such reasons conflicts. The antipaternalist in some sense “imposes” their view that reasons to refrain from interfering should generally prevail over reasons to interfere. Unsurprisingly, this is a “welcome” imposition in what is a generally antipaternalistic climate of philosophical writing.

I do not have the space to develop a comprehensive answer to this issue, as it is not the focus of my paper. However, I believe my view can affirm that it is generally not a paternalist’s business to interfere, but it can only do so if we are clearer on what it is to value autonomy. We could value autonomy in two ways: either we value its *possession*, or we value its *exercise*.⁴⁹ If the latter is so, I do not think the intuition that it is not our business to interfere will always be so strong. In cases such as Pouting Young Adult—or a variant of that case where Tom is a close friend of ours—we might think it utterly callous not to do something and get involved. Perhaps this is because what matters is not the fact that Tom is autonomous but that he is exercising his autonomy in a problematic way.

The idea, then, is that certain exercises of autonomy have more value than others and that it is those valuable exercises of autonomy that make us think it is not our business to interfere. For example, when an unconscious Jehovah’s Witness is given a blood transfusion, we might think this is problematic precisely because the expression of a religious belief is a valuable exercise of one’s autonomy. Here, we could plausibly think it is not our business to save their life. So, if autonomy’s value is linked somehow to its exercise, then we will not always think it is not our business (not to interfere). If that is true, then my view can still map onto this somewhat revised antipaternalistic phenomenological datum.

4. CONCLUSION

I hope to have shown that the exclusionary strategy is problematic due partly to how difficult it is to elaborate and because it does not square well with the mainstream view of paternalism’s *pro tanto* wrongness. Importantly, I do not think that problems with the exclusionary strategy should raise any concerns about the viability of exclusionary reasons in general. It should not do so because the application of exclusionary reasons to any domain of philosophy will come with its own unique intricacies and theoretical baggage. Nevertheless, perhaps because the exclusionary strategy is still being developed, the problems I have

48 Malm, “Feinberg’s Anti-Paternalism and the Balancing Strategy,” 198.

49 I am borrowing here from Raz’s discussion of autonomy. See Raz, *The Morality of Freedom*, 370.

raised may yet be resolvable. In that case, this paper can be read as an invitation to antipaternalists drawn to these ideas to deal with the complexities. It may also be that the exclusionary strategy is so appealing (for other reasons) that these complexities, if not resolvable, might be taken in stride.

That being said, I believe we can salvage the exclusionary strategy's appeal and maintain a similar antipaternalistic stance with a normative toolkit that is more familiar and on a run-of-the-mill first-order level of reasoning. In a way, the idea of balancing reasons for and against paternalistic intervention is a commonsensical one. What I hope to have added to this commonsense view is some more precision by adding modifiers to the debate. We should focus on not only the reasons for and against paternalistically interfering but also what might specifically influence the strength of those reasons. More can be said in defense of the view I have developed, but this sketch is an important step toward using normative reasons in the context of paternalism's wrongness in an intuitively better way.⁵⁰

Hertford College, University of Oxford
kyle.vanoosterum@philosophy.ox.ac.uk

REFERENCES

- Adams, N. P. "In Defense of Exclusionary Reasons." *Philosophical Studies* 178, no. 1 (January 2021): 235–53.
- Arneson, Richard J. "Joel Feinberg and the Justification of Hard Paternalism." *Legal Theory* 11, no. 3 (September 2005): 259–84.
- Bader, Ralf. "Conditions, Modifiers, and Holism." In *Weighing Reasons*, edited by Errol Lord and Barry Maguire, 27–55. Oxford: Oxford University Press, 2016.
- Begon, Jessica. "Paternalism." *Analysis* 76, no. 3 (July 2016): 355–73.
- Birks, David. "How Wrong Is Paternalism?" *Journal of Moral Philosophy* 15, no. 2 (April 2018): 136–63.

50 I am extremely grateful to Lorenzo Elijah and Jonathan Parry, who read multiple drafts and offered very useful feedback and discussion. I would also like to thank David Enoch, David Birks, Nikhil Krishnan, Jen Semler, Lewis Williams, Rose Macaulay, and the editors and reviewers of this journal for their helpful suggestions to improve this paper. A version of this paper was first presented at the Philosophy DPhil Seminar at the University of Oxford, where it benefitted from good feedback from the attendees. Finally, I owe thanks to the Institute for Ethics in AI at the University of Oxford for funding and supporting my research.

- . “Sex, Love, and Paternalism.” *Ethical Theory and Moral Practice* 24, no. 1 (March 2021): 257–70.
- Bullock, Emma C. “A Normatively Neutral Definition of Paternalism.” *Philosophical Quarterly* 65, no. 258 (January 2015): 1–21.
- Chang, Ruth. “Do We Have Normative Powers?” *Aristotelian Society Supplementary Volume* 94, no. 1 (July 2020): 275–300.
- Cholbi, Michael. “Paternalism and Our Rational Powers.” *Mind* 126, no. 501 (January 2017), 123–53.
- Clarke, D. S., Jr. “Exclusionary Reasons.” *Mind* 86, no. 342 (April 1977): 252–55.
- Coons, Christian, and Michael Weber. “Introduction: Paternalism—Issues and Trends.” In *Paternalism: Theory and Practice*, 1–24. Cambridge: Cambridge University Press, 2013.
- Dancy, Jonathan. *Ethics without Principles*. Oxford: Oxford University Press, 2004.
- Dworkin, Gerald. “Defining Paternalism.” In *Paternalism: Theory and Practice*, edited by Christian Coons and Michael Weber, 25–38. Cambridge: Cambridge University Press, 2013.
- . “Paternalism.” In *Morality and the Law*, edited by Richard Wasserstrom, 107–26. Belmont, CA: Wadsworth Publishing Company, 1971.
- Enoch, David. “Hypothetical Consent and the Value(s) of Autonomy.” *Ethics* 128, no. 1 (October 2017): 6–36.
- . “What’s Wrong with Paternalism: Autonomy, Belief, and Action.” *Proceedings of the Aristotelian Society* 116, no. 1 (April 2016): 21–48.
- Feinberg, Joel. *Harm to Self*. Vol. 3 of *The Moral Limits of the Criminal Law*. Oxford: Oxford University Press, 1986.
- Grill, Kalle. “Antipaternalism as a Filter on Reasons.” In *New Perspectives on Paternalism and Healthcare*, edited by Thomas Schramme, 46–63. New York: Springer International Publishing, 2015.
- . “The Normative Core of Paternalism.” *Res Publica* 13, no. 4 (December 2007): 441–58.
- Grill, Kalle, and Jason Hanna, eds. *The Routledge Handbook of the Philosophy of Paternalism*. New York: Routledge, 2018.
- Groll, Daniel. “Medical Paternalism—Part 1.” *Philosophy Compass* 3, no. 9 (March 2014): 194–203.
- . “Paternalism and Rights.” In Grill and Hanna, *The Routledge Handbook of the Philosophy of Paternalism*, 119–31.
- . “Paternalism, Respect, and the Will.” *Ethics* 122, no. 4 (July 2012): 692–720.
- Hanna, Jason. “Hard and Soft Paternalism.” In Grill and Hanna, *The Routledge Handbook of the Philosophy of Paternalism*, 24–34.

- . *In Our Best Interest: A Defense of Paternalism*. Oxford: Oxford University Press, 2018.
- Kleinig, John. *Paternalism*. Totowa, NJ: Rowman and Allanheld, 1983.
- Malm, Heidi. “Feinberg’s Anti-Paternalism and the Balancing Strategy.” *Legal Theory* 11, no. 3 (September 2005): 193–212.
- Mill, John Stuart. *On Liberty*. London: Penguin, 2010.
- Moore, Michael S. “Authority, Law and Razian Reasons.” *Southern California Law Review* 62, nos. 3–4 (1989): 827–96.
- Owens, David. *Shaping the Normative Landscape*. Oxford: Oxford University Press, 2012.
- Parfit, Derek. *Reasons and Persons*. Oxford: Oxford University Press, 1986.
- Parry, Jonathan. “Defensive Harm, Consent, and Intervention.” *Philosophy and Public Affairs* 45, no. 4 (September 2017): 356–96.
- . “What’s Wrong with Paternalism?” Unpublished manuscript.
- Quong, Jonathan. *Liberalism without Perfection*. Oxford: Oxford University Press, 2011.
- Raz, Joseph. *The Morality of Freedom*. Oxford: Clarendon Press, 1986.
- . *Practical Reason and Norms*. Oxford: Oxford University Press, 1999.
- Scanlon, T. M. “Reasons: A Puzzling Duality.” In *Reasons and Value: Themes from the Moral Philosophy of Joseph Raz*, edited by R. Jay Wallace, Philip Pettit, Samuel Scheffler, and Michael Smith, 231–46. Oxford: Oxford University Press, 2004.
- Shafer-Landau, Russ. “Liberalism and Paternalism.” *Legal Theory* 11, no. 3 (September 2005): 169–91.
- Shiffrin, Seana Valentine. “Paternalism, Unconscionability Doctrine, and Accommodation.” *Philosophy and Public Affairs* 29, no. 3 (July 2000): 205–50.
- Tsai, George. “Paternalism and Intimate Relationships.” In Grill and Hanna, *The Routledge Handbook of the Philosophy of Paternalism*, 348–59.