

Moral Uncertainty and Public Justification

Jacob Barrett & Andreas T. Schmidt

Vanderbilt University

University of Groningen

© 2024 Jacob Barrett & Andreas T. Schmidt
This work is licensed under a Creative Commons
Attribution-NonCommercial-NoDerivatives 3.0 License.
<www.philosophersimprint.org/024003/>
DOI: [10.3998/phimp.3016](https://doi.org/10.3998/phimp.3016)

1. Introduction

Moral disagreement pervades our lives. We disagree about the rightness or wrongness of actions, the goodness or badness of outcomes, and the justice or injustice of institutions. These disagreements often seem quite reasonable – and equally intractable. Moral reasoning is hard, requiring us to navigate complex concepts and their intricate and often surprising implications. We come to this task with different life experiences, educations, and social networks, and so with different biases, priors, and evidence bases. And, even when we agree about which moral considerations matter, we often disagree about their weights. Moral thinking, in other words, is subject to the “burdens of judgment” (Rawls 2005, 55–57; compare MacAskill et al. 2020, 11–14). And it is a predictable consequence of these burdens that intelligent people, reasoning in good faith, will come to different conclusions about morality.

Given the many plausible moral views available to us, and their many capable and eager champions, it is difficult to know how to proceed. We must reckon both with the fact of our own uncertainty about morality, and with the fact that others inevitably reach different conclusions than we do. These two facts, although closely related, have spawned two different research programs in contemporary philosophy: *public reason liberalism* in political philosophy and *moral uncertainty theory* in ethics. Public reason liberals ask what laws governments should enforce, given individuals’ reasonable disagreements about morality. They argue that governments must take all reasonable positions into account: it is permissible to enforce a law only when it can be justified to all reasonable people. Moral uncertainty theorists are concerned with what individuals should do when they are uncertain about morality. Most argue that we should take all plausible moral positions into account: what you should do depends not only on the moral theory you find most plausible, but also on the verdicts of all other moral theories in which you place some positive credence.

Our goal in this article is to bring these research programs into contact. To frame our discussion, we investigate the hypothesis that moral

uncertainty theory lends support to public reason liberalism. Our tentative conclusion is that moral uncertainty theory cannot vindicate the stringent requirement that *all* laws be publicly justified, but that it nevertheless provides several reasons to take public justification seriously. Specifically, it supports governments adopting public justification as a weighty *pro tanto* principle – albeit one that can be overridden when the moral stakes are high.

Along the way, we also highlight some attractive features of our novel defense of public justification. For instance, some critics argue that existing defenses of the public justification principle fail to cohere with the principle itself because they assume controversial first-order views about morality or justice, either explicitly or in how they delineate the class of “reasonable” people. Moral uncertainty theory sidesteps this issue, because it permits uncertainty about all first-order views of morality and justice and relies on a thin and independently motivated notion of reasonableness. It also offers a fresh perspective on some intramural debates among public reason liberals – for example, those concerning who counts as reasonable, what it takes to justify a law to a reasonable person, and what role “shared reasons” should play in public justification.

We proceed as follows. In section 2, we outline public reason liberalism and moral uncertainty theory, and introduce our hypothesis that the latter provides support for the former. In sections 3, 4, and 5, we discuss arguments bearing on this hypothesis. We comment on intramural debates on public justification in section 6, and suggest some directions for further research in section 7.

2. Two Second-Order Approaches to Justification

2.1 *Public Reason Liberalism*

How should we justify institutions and state action? In this article, we focus on coercive *laws*. However, much of what we say applies, *mutatis*

mutandis, to other institutions and forms of state action, and perhaps even to informal institutions or social norms.¹

On what we call the *first-order moral approach*, political philosophy is about finding the correct moral theory and applying it in the political domain. To determine which laws are justified, we must figure out, first, whether egalitarianism, libertarianism, utilitarianism, or some other theory is the correct view and, second, what laws the correct view supports. This does not mean that utilitarians, for instance, must ignore moral disagreement. Instead, utilitarians factor in disagreement as an empirical regularity: what laws do the most good given that not everyone is a utilitarian?

Given moral disagreement and the pluralism characteristic of modern societies, many political philosophers nevertheless find the first-order moral approach unsuitable, preferring a mode of political justification that takes disagreement seriously. The most common proposal is the *public justification principle*: governments should only exercise their power in a way that can be justified to every reasonable person.² Call a law “publicly justified” when it can be justified to all reasonable people, and a law “publicly unjustified” when it cannot. The public justification principle imposes a prohibition against enforcing publicly unjustified laws, not a requirement to enforce publicly justified laws.³ It says that governments are prohibited from enforcing laws that some reasonable people reject.

1. Public reason liberals disagree about the object of public justification. For example, Rawls (2005, 215) holds that only issues of basic justice and constitutional essentials must be publicly justified, but that “it is usually highly desirable to settle political questions by invoking the value of public reason.” Quong (2011, chap. 9) extends the requirement of public justification to all laws; Gaus (2011, 490–97) extends it to social norms; Waldron (1993, 36–37) applies it to “all aspects of the social world.”
2. See, for example, D’Agostino 1996, Gaus 2011, Larmore 1990, Lister 2013, Nagel 1987, Rawls 2005, Vallier 2019, Waldron 1993. Vallier (2022) provides additional citations.
3. The public justification principle is often seen as a requirement on *legitimate* state action, which some, but not all, see as equivalent to *permissible* state action. We avoid that contested concept here and speak directly about what governments are permitted to do.

Theorists who accept the public justification principle are called “public reason liberals,” but the label is confusing because not all public reason liberals think that justification must proceed via “public reasons.” We can distinguish *consensus* from *convergence* liberals (D’Agostino 1996; Vallier 2011). Consensus liberals argue that we must bracket people’s private or non-shared reasons and ask whether laws can be justified to all reasonable people factoring in only the reasons they “share” or all see as carrying some justificatory weight (even if they disagree about this weight). Convergence liberals, in contrast, hold that whether a law is justified to a reasonable person depends on all of their reasons (or at least those reasons meeting a minimal standard of “intelligibility”). Accordingly, on the convergence approach, laws might be publicly justified (or unjustified) because all accept them (or some reject them) for non-shared reasons.

Public reason liberalism, then, is really a family of theories. We will comment on some intramural disputes later, but for now we merely note that, to remain a viable and distinctive position, public reason liberalism must navigate between two poles. First, because public reason liberals are not anarchists, they must show that at least *some* laws can be justified to all reasonable people: publicly justified laws should not form an “empty set.” There are two basic strategies to respond to this worry: either one can restrict the class of people who count as “reasonable,” say, to those who embrace core liberal commitments (Quong 2011, chap. 5); or one can lower the standard of justification so that a law is justified to a reasonable person, say, when they merely see it as “better than nothing” (Gaus 2011, 321–25). Second, to avoid their position “collapsing” into a first-order moral approach, public reason liberals must avoid implying that laws are publicly justified only when the correct first-order moral theory says so (Raz 1990, 46). For example, they should not restrict who counts as “reasonable” to those who accept the correct first-order moral theory, nor should they count a law as justified to someone only if it is justified by the correct first-order moral theory.

In this sense, public reason liberalism is a “second-order” theory of justification: a theory of how to justify laws given disagreements about first-order morality. And it claims that governments are prohibited from enforcing publicly unjustified laws.

2.2 *Moral Uncertainty Theory*

Whereas public reason liberalism asks how governments should proceed given disagreements about morality, moral uncertainty theory asks how individuals should proceed in light of their own uncertainty about morality: what ought someone to do when they don’t know what they morally ought to do (Sepielli 2009)? It therefore investigates the moral analogue to decision-making under empirical risk and uncertainty. Suppose you are sympathetic to utilitarianism but not fully certain; you also have some credence in other moral views, such as Kantianism and virtue ethics. You have three options to choose from — *A*, *B*, and *C* — and different moral views give you different prescriptions about what to do. Moral uncertainty theory aims to tell you what to do given your uncertainty about the correct moral view.

Consider different versions of moral uncertainty theory. *My Favorite Theory* (MFT) says that you should choose the option favored by the moral view in which you have the highest credence (Gracely 1996; Gustafsson & Torpman 2014). So, if you think that utilitarianism is most likely to be right, you ought to follow utilitarianism. Although many seem implicitly to adopt this approach, most moral uncertainty theorists reject it. One reason is the problem of theory individuation (MacAskill et al. 2020, 41–44). Imagine you have 60% credence in consequentialism, 30% credence in Kantianism, and 10% credence in virtue ethics. It might seem that you should choose by applying consequentialism, but there are many forms of consequentialism: for example, you could be a classical utilitarian, or a prioritarian, or a consequentialist with a richer theory of the good. Say you have equal credence in three different types of consequentialism such that your credence in each of them is 20%. Instead of consequentialism, your favorite theory now turns out to be Kantianism!

An alternative that avoids this problem is *My Favorite Option* (MFO) (see Lockhart 2000, 26). MFO says that under moral uncertainty you ought to choose the option most likely to be right. So, to continue the above example, imagine that all three versions of consequentialism say to choose *A*, Kantianism says to choose *B*, and virtue ethics says to choose *C*. MFO tells you to choose *A*, as *A* has a 60% probability of being the right option and is thus the option most likely to be right.

Although MFO avoids the problem of theory individuation, most moral uncertainty theorists still reject both it and MFT for being *stakes-insensitive*. Because the stakes implied by different moral theories can vary greatly, the intuitively correct response is sometimes to hedge against moral risk by *not* choosing the option most likely to be right (MacAskill et al. 2020, 44–47). For example, imagine that you have 60% credence in consequentialism but 40% credence in a deontological theory on which killing — even to promote moderately good consequences — is very wrong. In this case, you should refrain from killing when killing has only marginally better consequences than not killing, because a small probability of doing something very wrong (violating a serious deontic constraint) can outweigh a large probability of doing something slightly wrong (producing marginally suboptimal consequences). Both MFT and MFO consider only the probability of different theories or options being correct while ignoring the stakes between them, so they fail to accommodate such “moral hedging.”

In this article, we thus assume that adequate theories of decision-making under moral uncertainty are stakes-sensitive. Specifically, where greater precision is needed, we assume:

Maximize Expected Choiceworthiness (MEC): *A* is an appropriate option iff *A* has the maximal expected choiceworthiness. (MacAskill et al. 2020, 48; see also Lockhart 2000, Ross 2006, Sepielli 2009)

MEC is the moral uncertainty equivalent of expected value theory: it sees moral theories as assigning choiceworthiness scores to options, multiplies the choiceworthiness score each theory assigns to an option

by your credence in that theory, sums together the ensuing products to determine an option’s *expected* choiceworthiness, and instructs you to choose the option with the greatest expected choiceworthiness. Much like expected value theory, MEC can be adjusted in various ways to account for issues such as risk-aversion (MacAskill et al. 2020, 48). Importantly, it is stakes-sensitive and so allows for moral hedging, as in our example above, where MEC would likely recommend not killing because of your credence in deontology.

Although MEC is the most popular approach to moral uncertainty, some object to its strong assumptions about intertheoretic comparability. It presupposes that we can make intertheoretic “unit” comparisons of the following form: the difference in choiceworthiness between *A* and *B* according to deontology is *n* times the difference in choiceworthiness between *A* and *B* according to consequentialism (Gracely 1996; Gustafsson & Torpman 2014).⁴ We see the force of this worry, but believe that it can be overcome, as proponents of MEC have provided several compelling responses (MacAskill et al. 2020, chap. 5, is a helpful survey). However, it would take us too far afield to discuss these responses here. Instead, we simply note that there are other approaches to moral uncertainty that also capture intuitions about stakes and moral hedging without relying on intertheoretic comparisons. These include bargaining approaches, which treat moral theories as if they bargain over what action to perform (Greaves & Cotton-Barratt 2024), and social choice approaches, which treat moral theories as if they vote on which action to perform (MacAskill 2016; Tarsney 2019).

In further work, it would be interesting to consider whether the case for public justification is stronger or weaker given different approaches to moral uncertainty or different assumptions about intertheoretic comparability. Indeed, bargaining approaches to moral uncertainty may be especially fruitful to investigate, given the central

4. In the only other discussion (we know of) connecting moral uncertainty to public reason liberalism, Valentini and List (2020, 203–4) raise this worry without developing it, writing only: “We suspect that, contrary to what some recent literature on moral uncertainty suggests, [intertheoretic comparisons] pose insurmountable challenges.”

role bargaining plays for some public reason liberals (Muldoon 2016, chap. 4). But, to keep things manageable and to keep the focus on public reason liberalism generally, we will assume MEC going forward.

2.3 *From Moral Uncertainty to Public Justification?*

Moral uncertainty theory has so far been primarily employed in applied ethics. Recently, however, there have been a few attempts to apply it in political philosophy (e.g., Bukoski 2021; Barry & Tomlin 2019). We welcome this development and believe that there is room for much work at the intersection of moral uncertainty and political philosophy. Here, we focus on public reason liberalism. Public reason liberals have good reason to take interest in, and to draw on, moral uncertainty theory.

First, moral uncertainty theorists and public reason liberals share various concerns and motivations: in moral matters, it is hard to know — and so we disagree about — which theory is correct, but at the end of the day we must still act. Public reason liberalism says that simply appealing to our favorite first-order moral or political theory will not do. Similarly, stake-sensitive approaches to moral uncertainty, such as MEC, say that deciding in light of one's Favorite Theory (or Option) in ethics will not do. Moreover, both approaches suggest that we should take other reasonable people's moral views seriously: public reason liberals explicitly focus on accommodating moral disagreement, and moral uncertainty theorists identify widespread disagreement as a key reason to be morally uncertain (MacAskill et al. 2020, 12–13). Likewise, as we will see, some public reason liberals already make incipient appeals to moral uncertainty, or to concepts in its vicinity, when justifying their approaches.

Second, public reason liberals are often faced with the challenge of justifying the public justification principle itself. Various responses are on offer. However, critics sometimes argue that these responses smuggle in controversial first-order moral theories (Enoch 2013; Wall 2002). For example, some public reason liberals argue that respect or political community can ground the public justification principle. But

is there not reasonable disagreement about these values too? And, if there is, why isn't public reason liberalism "but another sectarian doctrine" (Rawls 1985, 246) — a first-order theory competing with other first-order theories, rather than a second-order theory that stands above the fray? Others draw a stark epistemological distinction between conceptions of the good and morality on the one hand, and principles of justice and institutional justification on the other, arguing that only the latter have sufficient epistemic status to justify state coercion (Barry 1996, 169–71; Nagel 1987). However, it is far from obvious that theories of justice or institutional justification are epistemically stronger or more robust than views on morality and the good life (Clarke 1999; Enoch 2017). They, too, are subject to widespread disagreement and uncertainty.

Here, moral uncertainty theory may offer public reason liberals an underexplored route to avoid these problems.⁵ Specifically, it does not require a firm commitment to any first-order theory of morality or of justice, nor does it claim that political principles have a stronger epistemic status than ethical views and conceptions of the good. Instead, it permits uncertainty about first-order theories in political philosophy, including about the public justification principle itself.⁶

Finally, we will see that moral uncertainty theory provides interesting, external, and distinctive answers to some intramural debates among public reason liberals.

It is therefore worth investigating whether moral uncertainty offers support for public reason liberalism. To this end, we consider:

Hypothesis: Governments MU-ought to adhere to a rule of only enforcing publicly justified laws.

5. Of course, our discussion of existing strategies for justifying public justification is non-exhaustive. Public reason liberals attempt other responses to these worries, too. See Bajaj 2017 and Bepalov 2021, among others.
6. Admittedly, moral uncertainty theorists face an analogous challenge of how to deal with uncertainty about which approach to moral uncertainty to use (Weatherson 2014). But, again, moral uncertainty theorists provide several plausible responses to this "regress" problem. For discussion, see MacAskill et al. 2020, 33, and Trammell 2019.

Note two features of this hypothesis.

First, we use “MU-ought” to refer to the “ought of moral uncertainty” or “what we ought to do according to moral uncertainty theory,” as opposed to what we ought to do according to the correct first-order moral theory. So, in our earlier example, stakes-sensitive approaches to moral uncertainty imply that a person sometimes MU-ought to choose a different option from what they ought to choose on the correct first-order moral theory (when they engage in moral hedging, for example).⁷

Second, our hypothesis concerns a rule that governments MU-ought to adhere to. This differs subtly from a hypothesis about what governments MU-ought to do in every token instance, such as:

*Hypothesis**: governments MU-ought only to enforce publicly justified laws.

An analogy might help. Suppose you are a utilitarian and think that under specific and unlikely circumstances it can be morally right for the government to torture someone (say, to save many lives). Still, you might think that governments should adhere to a *rule* of never torturing, because you think the risk of failing to engage in justified torture so minute, and the risk of unjustified torture so high, that adhering to a strict prohibition on torture outweighs its downsides. Imagine now that instead of being a committed utilitarian, you are uncertain about which first-order moral theory is correct. Even if many theories say that torture is always wrong, if you put some credence in utilitarianism, you might hold that in rare circumstances, governments MU-ought to torture someone (say, to save an enormous number of lives). However, these circumstances might be so rare and the moral risk of unjustified torture so high that, on balance, moral uncertainty theory favors governments adhering to a strict prohibition on torture. So, governments

7. Moral uncertainty theorists disagree about how best to understand “MU-ought.” For example, is this a rational or a moral ought? We remain neutral on this question here. For discussion, see MacAskill et al. 2020, 18–21, and Sepielli 2013.

MU-ought to adhere to rule of never torturing, even though there may be some cases where governments MU-ought to torture.

Now, similarly, moral uncertainty theory might support governments adhering to public justification *as a rule* even if, in some instances, governments MU-ought to depart from public justification. Put in terms of MEC: even if adhering to public justification sometimes leads to decisions that do not maximize expected choiceworthiness, adhering to public justification as a rule (instead of deciding whether to do so on a case-by-case basis) might still maximize expected choiceworthiness overall. So *Hypothesis* might be true even if *Hypothesis** is not.

Before moving on, we must dispense with one more ambiguity. Public justification concerns state action. Moral uncertainty theory, in contrast, is about what an individual ought to do, given their credences in different moral positions. But it is unclear *which* credences to appeal to when thinking about what governments MU-ought to do. So, moral uncertainty theory might not be a natural framework in this context.

We see, broadly, three ways of using moral uncertainty here, and in political philosophy more generally.

First, claims about what governments MU-ought to do could be addressed to a particular person who is morally uncertain. So, imagine that you are somewhat uncertain about morality. We could then use moral uncertainty theory to convince you that, given your own credences in what the government ought to do, you should accept that it MU-ought to adhere to the public justification principle. We might then repeat this process and seek to justify public justification to others who are morally uncertain with different credences.

Second, rather than addressing individuals’ actual credences, we might claim that only some credences over moral views are rational. We need not assume that there is one rational probability distribution over moral views, but could assume a range: for example, it may be irrational to have a high credence in a deeply racist view but rational to have a wide range of non-negligible credences in deontology, libertarianism, consequentialism, and so on. It is then based on these

rational probability distributions that we construct arguments about what governments MU-ought to do.

Third, we might address collective agents. So, when a government or society needs to decide on a law, we treat the relevant collective as having credences in different moral views. Here we could treat people's beliefs as represented in a population as analogous to an individuals' subjective probabilities. However, such a view encounters a challenge: how should we aggregate individuals' beliefs into collective credences? And whose beliefs should we aggregate: all people in society, all reasonable people (or those with rational credences), or only parties to the decision? For reasons of scope, we pass over such questions here.

The distinction between the first two approaches tracks a disagreement among moral uncertainty theorists about whether what agents MU-ought to do depends on their actual credences or their "rational" credences, which we cannot hope to resolve here (MacAskill et al. 2020, 4). Nor can we hope to resolve the question of whether to opt for the individual or the collective approach. Instead, we remain agnostic. If one prefers the "actual credence" view, then our ensuing arguments can be interpreted as addressing agents who are in fact substantially uncertain about morality, at least if they allot their credences among the sort of views generally seen as contenders in moral and political philosophy. If one prefers the "collective credence" view, then our arguments can be interpreted as addressing collectives, at least if such collectives allot their credences among the same sort of views. Alternatively, if one prefers the "rational credence" view, then our arguments can be interpreted as suggesting that, given a range of rational credence distributions over moral views, governments (rationally) MU-ought to adhere to the public justification principle. The rational credence view arguably provides the most "philosophical" grounding: it does not justify public justification *to* anyone in particular but instead grounds it in a range of credence distributions it is rational to have.

Let us now explore arguments for our hypothesis.

3. The Proto-Moral Uncertainty Argument

We begin with what we call the "proto-moral uncertainty argument" — an argument suggested, in incipient form, by two of public reason liberalism's best-known proponents: Rawls (2005, 125) and Gaus (2015, 1085). This argument is that, in general, even when our favored moral view suggests that we ought to enforce a law, there is a considerable probability that this is wrong. Governments that enforce laws according to our (or their) favored moral theory thus run a serious moral risk. Adhering to the public justification principle reduces this risk. Specifically, assume that for any potential law, at least one reasonable person gets the correct answer about whether this law is "objectively justified," that is, justified according to the correct moral theory. It follows that, if a government never enforces laws that some reasonable person rejects, it will only enforce laws that are objectively justified. Publicly justified laws are therefore "safe" options; publicly unjustified laws are not. To the extent that the above assumption holds, requiring public justification will ensure that we never enforce laws that are "objectively unjustified."

So, adopting the public justification principle reduces the probability of enforcing laws that are unjustified according to whatever moral theory turns out to be correct. Objectively unjustified laws may still slip through the cracks of the public justification test if no reasonable person has the correct view on some issue. But at least we know that, barring a blanket prohibition on enforcing any laws, we cannot reduce this risk any further.

However, this argument has two problems.

First, it assumes a strong asymmetry between wrongly enforcing laws and wrongly failing to enforce them. Let us say that when a government enforces a law that is objectively unjustified, this is a "moral false positive." When a government fails to enforce a law that is objectively justified, call this a "moral false negative." Adhering to the public justification principle reduces the risk of moral false positives. However, it also increases the risk of moral false negatives: objectively

justified laws will fail the public justification test whenever some reasonable person rejects them. So, the argument gives far greater weight to avoiding false positives than avoiding false negatives, and we need a justification for this asymmetry.

Second, the proto-moral uncertainty argument fails to be stake-sensitive. Bent on preventing moral false positives, it resembles MFO, which minimizes the probability that we choose an option that is objectively wrong. However, recall that plausible approaches to moral uncertainty focus not only on probabilities but also on stakes. As our example of killing to promote better consequences suggested, we sometimes MU-ought to hedge and prefer an option likely to be slightly wrong over an option with a smaller probability of being very wrong. So, arguments for our hypothesis need to go beyond the probability of moral false positives and negatives and consider the severity of such errors, too.

The Proto-Moral Uncertainty Argument alone will not do. We require reasons to think that adhering to the public justification principle will not merely reduce the risk of moral false positives but also appropriately balance the risks of false positives against the risks of false negatives, while being sensitive to stakes. We now propose four such reasons.

4. Four Considerations Favoring Public Justification

As we have seen, when balancing the risks of moral false positives against the risk of moral false negatives, we must consider two factors. First, we must consider the *probability* of either risk materializing. As the probability that a law is objectively justified decreases, the risk of a moral false positive goes up and the risk of a moral false negative goes down. Second, we must consider the *severity* of moral errors. The morally worse it would be to enforce some law that is objectively unjustified, the more severe the risk of a false positive. The morally worse it would be not to enforce some law that is objectively justified, the more severe the risk of a false negative.

Let us now say that a “publicly unjustified false positive” occurs when a publicly unjustified law is enforced when it objectively should not be (it is a type of moral false positive). A “publicly unjustified false negative” occurs when a publicly unjustified law is not enforced but objectively should be (it is a type of moral false negative). To defend our hypothesis, we need some considerations that either drive up the *severity* of publicly unjustified false positives relative to the severity of publicly unjustified false negatives, or increase the relative *probability* of publicly unjustified false positives. Specifically, these two types of considerations should show that not enforcing a publicly unjustified law typically has greater expected choiceworthiness than enforcing it. Now, some exceptions may occur, as we interpret public justification as a rule governments MU-ought to adhere to, rather than as a criterion of what governments MU-ought to do in all cases. Still, such exceptions should not be so frequent and severe that they undermine the expected choiceworthiness of governments adhering to public justification, nor should they arise so predictably that we could instead opt for an alternative rule with relevant exception clauses.

Here is a consideration of the second type, concerning the relative probability of publicly unjustified false positives and negatives:

Public Justification Tracks Objective Justification: When a law is publicly justified, it has a high probability of being objectively justified. When a law is publicly unjustified, there is at least a significant probability that it is not objectively justified.

The first half of this claim is from the proto-moral uncertainty argument above. The second half does not follow immediately from the first, but we think it plausible. If some reasonable people believe a law is unjustified, we should take seriously the possibility that it is unjustified. We could guarantee this result through a particular understanding of “reasonable”: when a reasonable person believes a law is not justified, then there is a significant probability they are right. This will hold if we stipulatively define a reasonable person with respect to

some law as someone who either holds a credible view with respect to that law or whose testimony we should take seriously, such that if they deny that a law is justified, we should think that there is a significant probability that the law is objectively wrong. Alternatively, and perhaps more promisingly, rather than defining a reasonable person in this way, we can stick with an intuitive notion of what a reasonable person is, and simply note that this notion overlaps quite considerably with the above stipulative definition: when someone disagrees with us about something, and we think that this disagreement is reasonable, we at least *tend* to believe that there is some probability that they are right. So *Public Justification Tracks Objective Justification* may hold either by definition or as a general tendency. Notably — and as will be discussed in greater detail below — this consideration appears to mesh better with a “convergence” than a “consensus” approach to public justification (the latter of which, recall, considers only “shared” reasons). When a reasonable person regards a law as unjustified, there is a significant chance that they are right — and this holds regardless of whether they believe this for reasons that all reasonable people share.

By itself, *Public Justification Tracks Objective Justification* does not get us far. It gets our foot in the door by suggesting that the probability of publicly unjustified false positives is worth taking seriously. But it implies nothing about either the severity of publicly unjustified false positives or about how probable they are relative to publicly unjustified false negatives. A second consideration further opens the door:

Public Reason Liberals Might Be Right: Public reason liberals offer several reasons why laws, in virtue of being publicly unjustified, are objectively unjustified. Under moral uncertainty, we should give some weight to these reasons.

Public reason liberals defend the public justification principle in several ways. For example, they propose various deontological reasons suggesting that publicly unjustified laws are wrong, because they are disrespectful (Larmore 1990) or authoritarian (Gaus 2011). And they propose various axiological reasons, suggesting that publicly

unjustified laws undermine social trust (Vallier 2019), political community (Leland & Wietmarschen 2017; Lister 2013), or a morally attractive notion of stability (Rawls 2005). Relatedly, they argue that publicly unjustified laws are less effective at securing whatever they aim to achieve because, all else equal, they are less stable, more likely to generate resistance and backlash, and so less predictable (Barrett & Gaus 2020). Under moral uncertainty, we should presumably give some credence to these arguments and, accordingly, to the idea that a law is wrong in virtue of being publicly unjustified.

This consideration has two effects. First, much like the first consideration, it drives down the probability of publicly unjustified false positives. To the extent that we think public reason liberalism might give us the right first-order theory of when it is permissible to enforce a law, we should think it less likely that we objectively should enforce any given publicly unjustified law. Second, and more subtly, it drives down the severity of false negatives. To the extent that we put credence in the considerations public reason liberals marshal against publicly unjustified laws — for example, those concerning disrespect — even laws we all-things-considered objectively ought to enforce become somewhat less choiceworthy (in expectation) when they are publicly unjustified. This makes the risk of not enforcing publicly unjustified laws less severe.

Public Reason Liberals Might Be Right gets a grip because of the role asymmetries play under moral uncertainty (MacAskill et al. 2020, 183–87). Some theories claim that public justification is in itself valuable; others claim it does not matter. But no plausible theory judges public justification in itself disvaluable. So, under moral uncertainty, we should at least treat it as *somewhat* valuable. Our third consideration concerns a similar asymmetry:

There May Be a Presumption Against Coercive Laws: There are many plausible moral theories on which there is a strong presumption against coercive laws, and no plausible moral theory on which there is the reverse presumption.

So under moral uncertainty, there is at least a weak presumption against coercive laws.

This consideration picks up on a common refrain — that coercion requires justification in a way failures to coerce do not — often invoked as a premise in arguments for public justification (Feinberg 1989, 9; Gaus 2011, 319–21; Rawls 2001, 44). Its relevance here is this. We are after reasons to believe that, when it comes to publicly unjustified laws, the risk of a moral false positive typically outweighs the risk of a moral false negative (where exceptions to this rule are relatively infrequent and insignificant). *There May Be a Presumption Against Coercive Laws* supports this claim: it drives up the severity of the risk of moral false positives relative to the risk of moral false negatives in general, and not just when it comes to publicly unjustified laws. Specifically, many moral theories claim either that all coercion is in itself *pro tanto* wrong or that unjustified coercion is in itself very wrong. Some theories deny this (e.g., Wall 2010), but no plausible moral theory claims that failures to coerce are wrong in virtue of being failures to coerce. Similarly, many moral theories claim that wrongful acts are, all else equal, worse than wrongful omissions. Some theories deny this, but no plausible moral theory claims that wrongful omissions are worse than wrongful acts. Under moral uncertainty, the above asymmetries drive up the severity of false positives but not of false negatives, since enforcing a law is a *coercive act*, whereas failing to enforce a law is a *non-coercive omission*. We thus have reasons to err on the side of not enforcing laws.

The above three considerations are of exactly the types we were looking for: they increase the probability and severity of publicly unjustified false positives relative to the probability and severity of publicly unjustified false negatives. However, unless one assigns high credences to public reason liberalism (or the values it relies on), the presumption against coercion, or the significance of the act–omission distinction, we doubt that they suffice to confirm our hypothesis. We

therefore present a fourth consideration that applies to at least some conceptions of public justification and gives them a significant boost:

Public Justification is a Low Bar: On many conceptions of public justification, a law is justified to a reasonable person even if they see it as highly suboptimal. It must meet only some low threshold of being “better than nothing” or “something they can live with.”

Public reason liberals typically view public justification as a standard short of optimality. A law is considered publicly justified if all see it as good enough, even if it is not everyone’s first choice (Gaus 2011; Valier 2019).

Now, the lower the bar one sets here — is a law justified to a reasonable person when they see it as pretty good, better than nothing, or merely non-disastrous? — the more severe the risk of publicly unjustified false positives becomes. After all, if a law fails to meet only a high bar of public justification because some reasonable people see it as slightly wrong, there is a moral risk of enforcing a slightly wrong law; but if a law also fails a low bar of public justification because some reasonable people see it as extremely wrong, then there is a moral risk of enforcing an extremely wrong law.⁸ Furthermore, the lower the bar, the fewer laws will fail the public justification test overall, meaning that publicly unjustified false negatives become less likely when governments adhere to public justification. If we set the bar quite low, many laws will be publicly justified, and those that are not will have a serious probability of being severely wrong (by *Public Justification Tracks Objective Justification*).

8. Sometimes, a law might be wrong even if no reasonable person views it as such, or extremely wrong even if reasonable people view it as only slightly wrong; not all moral risks are noticed by reasonable people. However, this possibility is consistent with our claim that laws that reasonable people view as severely wrong are riskier than laws they view as only slightly wrong, because such “unnoticed” risks typically cancel out in expectation: they are no more likely to attach to laws that reasonable people view as slightly wrong than they are to attach to laws they view as extremely wrong.

Lowering the bar also interacts with *Public Reason Liberals Might Be Right*, because most of the substantive considerations that public reason liberals raise seem to scale with where we set the bar. For example, if enforcing a law is disrespectful or authoritarian when people object to it, it is presumably more disrespectful or authoritarian in cases where they have a stronger objection. Similarly, if enforcing publicly unjustified laws undermines social trust or stability because individuals are less willing to comply with laws they object to, this effect presumably increases with the strength of their objections. Lowering the bar thus amplifies the two effects mentioned in *Public Reason Liberals Might Be Right*, further increasing the probability and severity of publicly unjustified false positives relative to the probability and severity of publicly unjustified false negatives.

Likewise, lowering the bar interacts with *There May Be a Presumption Against Coercive Laws*: laws that fail a lower bar tend to be more coercive, leading to a stronger presumption against enforcing them. There are two reasons for this. First, on some conceptions of coercion, enforcing laws against people who more stringently object to them is, in itself, more coercive.⁹ Some conceptions of coercion may deny this, but none say the opposite; so we should give this consideration at least some weight under moral uncertainty. Second, there is plausibly some empirical correlation between how coercive laws are and how strongly people object to them, both because people often object more strongly to more coercive laws and because it often requires more coercion to secure compliance with laws that are more strongly opposed. So, lowering the bar likely increases the presumption against coercive laws.

Once we lower the bar, however, we need to slightly reinterpret *Public Justification Tracks Objective Justification*. Laws that pass a low bar of public justification may often be “compromises” that are not

9. For example, all things equal, a law may be more coercive when, absent the coercion, people would have weaker reasons to comply with it (Nozick 1969, 464) or would more strongly prefer not doing so (Feinberg 1989, 204). And people typically have weaker reasons to comply with – and more strongly prefer not complying with – laws to which they more stringently object.

regarded as severely wrong by anyone but are widely regarded as slightly wrong. To maintain that publicly justified laws have a high probability of being objectively justified, we therefore should understand “objectively justified” laws not as morally optimal but as meeting the relevant bar of justification (for example, being better than nothing). Yet this is all to the good, since it makes room for the sort of moral hedging that plausible versions of moral uncertainty theory recommend. To illustrate this, consider two reasonable people with a choice between different laws on a particular subject:

	Moral assessment of reasonable person 1	Moral assessment of reasonable person 2
No law	Disastrous	Disastrous
Law 1	OK	OK
Law 2	Great	Disastrous
Law 3	Disastrous	Great

If the public justification principle only allowed laws that are everyone’s first choice, it would generate a disastrous outcome: the government must select “no law.” However, with a lower bar, Law 1 would count as justified to both reasonable people, so the government would be permitted to enforce Law 1 (but not Laws 2 or 3). Now, assume – for simplicity – that both people’s moral views represent views in which we have (rational) credences of roughly 50%. MEC would likely tell us to hedge and choose Law 1, even though it is guaranteed not to be the morally best option. And, more generally, a lower bar permits governments to engage in moral hedging in a way that a higher bar does not. This is an interesting connection between moral hedging and public justification. Moreover, it helps to illuminate how *Public Justification is a Low Bar* reduces the probability of publicly unjustified false negatives and increases the severity of publicly unjustified false positives: fewer laws are ruled out by a lower bar and those that are carry more severe risks.

Taking the various considerations together, then, we are left with the following picture. Given *Public Justification Tracks Objective Justification*, there is a significant probability that enforcing a publicly unjustified law is a moral error: when reasonable people think that a law is wrong, we should take this possibility seriously. *Public Reason Liberals Might Be Right* similarly drives up the probability of publicly unjustified false positives, but also decreases the severity of publicly unjustified false negatives: to the extent that we put credence in the considerations public reason liberals adduce, publicly unjustified laws become less likely to be objectively justified or at least more likely to carry moral costs. *There May Be a Presumption Against Coercive Laws* provides a general presumption against enforcing laws, thereby increasing the severity of false positives without similarly raising the severity of false negatives: under moral uncertainty, there is an asymmetry between the (greater) risk of coercive actions and the (lesser) risk of non-coercive omissions. Finally, by *Public Justification is a Low Bar*, the lower the bar we set for public justification, the more these other considerations are amplified: with a lower bar, the probability and severity of false positives goes up, and the probability and severity of false negatives goes down.

5. Are These Considerations Enough?

The above four considerations support our hypothesis. But are they sufficient to ground public justification as a rule to which governments MU-ought to adhere? Do they suggest that not enforcing publicly unjustified laws generally has greater expected choiceworthiness than enforcing them, and that exceptions to this rule involve only mild losses of expected choiceworthiness?

At first glance, with a sufficiently low bar, it seems plausible that governments will typically do what they MU-ought by not enforcing publicly unjustified laws. Under moral uncertainty, it is a good idea to avoid options with significant probabilities of very bad outcomes, which publicly unjustified laws carry (by *Public Justification Tracks Objective Justification*, *Public Reason Liberals Might Be Right*, and *Public*

Justification is a Low Bar), particularly when considering option-types with risk asymmetries — in this case, between coercive actions and non-coercive omissions (by *There May Be a Presumption Against Coercive Laws*). However, at this point we run into a problem: sometimes, failing to enforce a publicly unjustified law also carries a significant risk of producing a very bad outcome. For example, imagine that some reasonable people believe that a law is unjustified, but others think that its absence would be a moral disaster. In such cases, if we assign a high enough credence to the view that failing to enforce the law would be a grave enough disaster, enforcing the publicly unjustified law may maximize expected choiceworthiness.

But how often will such cases arise? Initially, they might seem rare. For consider:

Laws that Prevent Moral Disasters Are Typically Publicly Justified: Typically, if failing to enforce a law would be very morally bad, then reasonable people will recognize this, and agree that the law is justified (at least given a low bar of justification).

This consideration is intuitive. It follows from the idea that, when failing to enforce a law would be very bad, reasonable people tend to converge on this perspective — or at least, tend not to think that enforcing the law would be so bad that it would fail to meet a low bar. When something is very bad, it tends to be bad for many reasons, and so from any reasonable perspective. For example, no reasonable person rejects laws needed to maintain basic order in society and to avoid widespread carnage, death, and destruction. Thus, to the extent that this consideration holds, it drives down the probability of publicly unjustified false negatives that are severely bad. And this may be enough to save public justification as a rule, because it suggests that, even if it would sometimes maximize expected choiceworthiness to enforce a publicly unjustified law, failing to enforce such laws will not decrease expected choiceworthiness by a large margin.

Unfortunately for the public reason liberal, however, this consideration does not hold with sufficient generality. There are obvious counterexamples where a law is publicly unjustified but failing to enforce it would have a serious probability of moral disaster. Suppose you are confident that animal suffering matters, that imposing a tax on meat consumption would reduce animal suffering by a huge amount, and that even if you are wrong about this there is relatively little moral cost to enforcing the tax. Yet some reasonable people reject the law because they believe that animals have no moral status and that imposing the tax involves unbearable moralizing. In this case, the risk of a publicly unjustified false negative may outweigh the risk of a publicly unjustified false positive. A government that adheres to the public justification principle might err severely in not enforcing the law.

Can we avoid this result if we make justification a low enough bar that a tax on meat consumption will clear it even for those who deny animals have moral status? Perhaps — but this will not resolve the problem. Imagine a series of progressively more restrictive laws that each serve to reduce animal suffering. At some point, the law will become sufficiently restrictive that reasonable people who deny the moral importance of animal suffering will reject the law as passing even a minimal bar of justification. But if the law reduces animal suffering by a huge amount, and we have a sufficiently high credence that animal suffering matters, we will still see the law as much more choiceworthy in expectation than its absence — the objections of some reasonable people notwithstanding.

In cases like this, we may have to take a stand and insist that governments MU-ought to enforce a law that is publicly unjustified: the moral stakes of not doing so are simply too high. Now, if such cases were few and far between, this might still leave intact the hypothesis that governments MU-ought to adhere to public justification *as a rule*. But we worry that such examples are not isolated but represent a systematic problem. Specifically, *Laws that Prevent Disasters Are Typically Publicly Justified* appears to fail in many cases of reasonable

disagreement in which the disagreement is not over how to weigh conflicting considerations, but over whether some consideration matters at all.

For example, consider that there are also deep disagreements about whether laws should be justified by their effects on future generations, including those in the far future, or whether states have obligations toward faraway people in other countries. Such disagreements are particularly troubling for our hypothesis, because they involve *high-stakes* problems: if non-human animals, humans in other countries, or far-future people matter, failing to take them into account creates a grave moral risk, given the gigantic numbers of individuals involved, and given that many moral theories are scale-sensitive (or even hold that wrongness or badness scale linearly with the number of individuals affected). Accordingly, if we apportion significant credence to scale-sensitive theories that assign moral status to such beings, these theories will be highly influential in moral uncertainty calculations. So, if some reasonable people deny that such considerations matter to the justification of laws, insisting on public justification can come with grave moral risks. This makes us skeptical: moral uncertainty might not support public justification even as a rule. There will systematically be cases in which governments that adhere to public justification will act in ways that are very low in expected choiceworthiness.

Let us explore two responses to this concern.

5.1 Constraining Reasonableness

One response is to stipulate that anyone who denies, say, that animals or far-future people have moral status is unreasonable on that basis. By considering such people unreasonable, we might eliminate high-stakes counterexamples to public justification. However, we do not advocate this solution.

First, such an understanding would not chime with moral uncertainty theory, which, we have suggested, should interpret people as “reasonable” when there is a significant probability that they are right. Although we do not insist that reasonable people be defined this way,

we do note that, if public reason liberals want to invoke moral uncertainty (and particularly *Public Justification Tracks Objective Justification*), there must be a large overlap between people considered reasonable and those with a significant probability of being right. This leaves the set of reasonable people quite wide. Importantly, it includes those who deny that animals have moral status or who believe versions of the person-affecting view on which we lack obligations toward far-future people. Indeed, not only are such views held by regular people, but they are also serious positions in academic debates.

A potentially more promising response is that “reasonableness” can allow for people who believe animals do not have moral status, as long as they are uncertain about this: reasonable people must themselves be morally uncertain. Here, we might require only that reasonable people be “somewhat uncertain” and not necessarily that they fall within some range of rational credences over different moral views. This move could be used in defense of our hypothesis: any reasonable person will have some moral uncertainty over whether animals have moral standing. This implies they should largely support animal welfare reforms, as the vast scale of the problem drives up the moral risk of continuing the status quo. For example, imagine that you have only a 0.05 credence that it is bad or wrong to raise and kill non-human animals in factory farms. You should still find factory farming disastrous in expected moral choiceworthiness, as it kills more than 100 billion animals a year.

That reasonable people should be at least somewhat morally uncertain is a controversial but perhaps plausible way to interpret the Rawlsian idea that “reasonable persons recognize and accept the consequences of the burdens of judgment” (Rawls 2005, 488). But we doubt public reason liberals should make this move, because it presupposes a particular picture of moral uncertainty — namely, that reasonable people must employ MEC (or at least another stakes-sensitive approach). Imagine someone instead adopts MFT. They might maintain that we MU-ought not to take animals into account, because their favorite theory says that animals do not matter. Is such a person

unreasonable? To us, it seems more plausible that they are reasonable if their moral judgment about a law has some significant probability of being correct. Of course, we could define reasonableness so that only people who endorse MEC are included, but this would exclude enough people to render the public justification principle toothless. Many people seem implicitly to employ something like MFT or MFO, and even moral philosophers have only recently begun to consider alternatives. We do not want the “public” in “public justification” to consist only of a small minority, primarily composed of academic philosophers (compare Valentini & List 2020, 204).

These arguments are inconclusive, and perhaps others will support defining reasonableness to require some level of moral uncertainty and the endorsement of a stake-sensitive approach to moral uncertainty. This notion of reasonableness could then be used to avoid problematic cases. However, we find another route more promising.

5.2 *Going Pro Tanto*

The above proposal sought to avoid problematic cases by narrowing the set of reasonable people. The alternative is to keep this set broad but weaken the public justification principle. When formulating our hypothesis, we left this somewhat open:

Hypothesis: Governments MU-ought to adhere to a rule of only enforcing publicly justified laws.

So far, we have interpreted “adherence” as requiring governments to treat public justification as a strict prohibition. But perhaps we could rescue our hypothesis by interpreting this less stringently:

The Pro Tanto Interpretation: Governments adhere to a rule of only enforcing publicly justified laws when they treat this rule as a weighty *pro tanto* principle that can sometimes be overridden by other weighty concerns.¹⁰

10. Would adopting the *Pro Tanto Interpretation* leave us with a view that still qualifies as public reason liberalism? After all, some avowed critics of public reason liberalism admit that public justification has *pro tanto* moral weight — e.g.,

We could then view public justification either as a binary criterion such that a law is either publicly justified or not, or as a scalar criterion such that laws can be publicly unjustified to different degrees and this can count more or less against enforcing laws.

While either approach would likely avoid the problem we have raised, we believe that the scalar approach is especially plausible. On this view, the extent to which a law is publicly justified depends on several variables. Most obviously, it depends on what bar of public justification a law achieves. If a law fails a very low bar of public justification (say, because some reasonable people think it would be a disaster), there is a stronger reason not to enforce it than if it fails only to meet a high bar (say, because some reasonable people think it only slightly suboptimal). So, we may say that laws are publicly justified to a lesser degree if reasonable people have stronger objections to them, and that governments' *pro tanto* commitment to public justification should vary accordingly.

We might also distinguish between laws that are *fully* publicly justified, meaning that they are justified to every reasonable person, and laws that *approximate* public justification, say, by being justified to an overwhelming majority (Barrett & Gaus 2020, 224). We could then say that laws that are justified to fewer reasonable people are less publicly justified, and that this also modifies the strength of the *pro tanto* prohibition. This is for two reasons.

The first relates to *Public Justification Tracks Objective Justification*. Typically, when more reasonable people object to a law, the law is less likely to be objectively justified (especially when judgments are independent). If the majority think a law unjustified, this should decrease our confidence in the law's justification more than if a small minority

Enoch (2013) and Wendt (2019). We think that it would, because many public reason liberals themselves endorse only a *pro tanto* commitment to public justification — e.g., Ebels-Duggan (2010) and Leland (2019). Indeed, MacMullen (2023) even argues that there is now a consensus among public reason liberals that, sometimes, public justification can be outweighed. However, we will not press this point here, as we ultimately care more about our substantive conclusion about the relationship between moral uncertainty and public justification than about how views should be labeled.

objects, which in turn should have more impact than if a single person objects. So if a law fails to even approximate public justification — if it is not justified to a large majority — the risk of publicly unjustified false positives is higher than if it fails only to be fully publicly justified.

Second, most substantive considerations that public reason liberals raise against publicly unjustified laws seem to scale, not only with the severity by which a law fails the public justification test, but also with the number of people who object to it. For example, if it is disrespectful or authoritarian (or wrongfully coercive) to enforce laws against people who do not accept them, or if this undermines social trust or stability, then presumably this is a bigger problem if more people object to the law. The effects of *Public Reason Liberals Might Be Right* (and *There May Be a Presumption Against Coercion*) are therefore greater for laws that are not even approximately publicly justified.

Overall, then, we find that adhering to public justification is often an effective way for governments to appropriately balance moral risks. However, certain exceptions (such as those involving animal welfare or far-future people) will systematically arise when governments incur a grave moral risk by not enforcing a publicly unjustified law. This calls into question our hypothesis that governments MU-ought to adhere to public justification as a rule. To rescue it, we have suggested weakening the public justification principle: moral uncertainty seems to support governments treating public justification as a weighty *pro tanto* consideration that can be overridden in high-stakes cases. This *pro tanto* requirement can be given either a binary interpretation or — more interestingly — a scalar interpretation. On the latter, the *pro tanto* commitment to not enforcing a publicly unjustified law should be stronger both when more reasonable people object to it and when they have stronger objections.

6. Intramural Debates

We have focused on whether moral uncertainty theory vindicates public reason liberalism. In this section, we briefly survey how our discussion sheds light on some intramural disputes among public reason

liberals. One such dispute, already addressed at some length, concerns who counts as “reasonable.” From the perspective of moral uncertainty, we have seen that the category of reasonable people must overlap with that of people who have a decent probability of being correct — either because we (rationally) assign significant credence to their moral view, or because we take them to be morally reliable and so take their testimony seriously. A central advantage of this notion of reasonableness is that it is independently motivated and avoids smuggling in so much specific normative content that it collapses public reason liberalism into a variant of the first-order moral approach.

A related debate concerns the constituency of public reason: to *which* reasonable people must governments justify themselves? Public reason liberals typically assume that public justification is owed domestically, that is, to all reasonable members of a society or state. But some argue that public justification is owed *globally* — either to individuals living in other states, or to other societies or “peoples” themselves (see Director 2019).

Moral uncertainty suggests that the objections of “outsiders” should carry some weight: following *Public Justification Tracks Objective Justification*, anyone who has some probability of being right is relevant. For several reasons, however, moral uncertainty also suggests giving greater weight to the objections of “insiders” (as far as the *pro tanto* commitment to public justification is concerned). First, insiders might better track the wrongness of local laws, because of “local knowledge” and because their own interests are involved. Second, following *Public Reason Liberals Might Be Right*, arguments from moral community, stability, respect, and so on, plausibly apply more strongly within societies. Third, *There May Be a Presumption Against Coercive Laws* suggests that justification is owed more strongly to those subject to coercion, particularly people bound by a shared coercive legal structure within jurisdictions. Finally, recall the familiar role of asymmetries under moral uncertainty: nearly all first-order reasons to value public justification suggest justifying ourselves to insiders; only some suggest justifying ourselves to others; so, given uncertainty among these theories,

we should give greater weight to justifying ourselves to a member of our own society than to someone outside of it.

Moral uncertainty also has implications for two other intramural disputes. First, one dispute concerns the conditions under which a law qualifies as justified to a reasonable person. Here, we have suggested that the importance of public justification is higher the lower we set the bar. If some person sees a law as somewhat suboptimal, the case for not enforcing it is weaker than if it is unjustified to them because they see it as morally terrible. Relatedly, although there is less discussion on this point, we have suggested that public justification might be weightier if we interpret the principle as less than perfectly stringent — as requiring a law to be justified not to each and every reasonable person, but approximately, say, to an overwhelming majority. When only one reasonable person objects to a law, the case against that law is weaker than when many reasonable people object.

Finally, recall the debate between “consensus” liberals who hold that public justification should only invoke reasons that all reasonable people share and “convergence” liberals who also include reasons that some reasonable people believe to carry no justificatory weight. Here, moral uncertainty theory provides no reason to endorse the shared reason requirement. To allow only shared moral considerations effectively treats non-shared considerations as if we have zero credence in them, which seems antithetical to moral uncertainty theory. Furthermore, appealing to non-shared reasons may sometimes bring additional epistemic benefits, because judgments reflecting non-shared reasons may be more independent.

In fact, our discussion suggests that the public justification principle may be most likely to lead us astray when people do *not* share reasons. For example, not enforcing a law might severely harm a class of beings (animals, far-future people, and so on) that some reasonable people see as lacking status, so that the harm done to them does not ground a shared or public reason. If so, consensus theorists do not locate the importance of the distinction between shared and non-shared reasons in the right place. From a moral uncertainty perspective, the

distinction is significant, not because we should only take shared reasons into account, but because cases where governments MU-ought to override the public justification principle tend to be those where there are strong non-shared reasons to enforce a law.

That said, public justification by shared reasons may still be somewhat desirable from the perspective of moral uncertainty. After all, consensus liberals argue that the values grounding public justification (grouped under *Public Reason Liberals Might Be Right*) become stronger under consensus: for example, laws that are justified by shared reasons may be more respectful or better promote political community (Leland & Wietmarschen 2017; Lister 2013). Convergence liberals deny this (Van Schoelandt 2019), but we can ignore such niceties here. Our point is only that, while moral uncertainty theory does not imply a shared reason requirement, it can hold that shared reasons play a valuable role in public justification — at least insofar as one puts credence in consensus liberals' first-order arguments to this effect.

7. Conclusion

We have argued that, while moral uncertainty theory cannot vindicate an exceptionless public justification principle, it implies that we should take public justification seriously. Given uncertainty about what moral or normative political theory is correct, and thus uncertainty about which laws are justified, governments MU-ought to adhere to a *pro tanto* version of the public justification principle. However, this is plausible only given a non-demanding interpretation of public justification: public justification must be a low bar. Moreover, the prohibition against enforcing publicly unjustified laws may need to be overridden in high-stakes cases where there is significant risk that failing to enforce a law is morally disastrous, even though some reasonable people disagree — for example, because the law protects a class of beings that some reasonable people deny have status. We have therefore provided a preliminary defense of public justification, not as an exceptionless principle, but as a *pro tanto* one.

However, our arguments are only first attempts at exploring the connection between public justification and moral uncertainty. We doubt that this will be the last word on the matter. Moreover, our hypothesis could be varied in a number of ways, and several robustness checks could be performed. For example, in future work, it would be interesting to consider what more might be said about the conditions under which an insistence on public justification is likely to lead us astray under moral uncertainty, and to explore how our conclusions might change given other approaches to moral uncertainty (especially those that make different assumptions about intertheoretic comparability). It is also worth considering what might change if we were to shift from the public justification of laws to the public justification of other objects in the public reason liberalism literature (such as principles of justice, constitutional essentials, or social norms), or from how governments MU-ought to proceed to how other potential subjects of public justification MU-ought to behave (for example, individual politicians or voters). Or we might compare the public justification principle to other principles, to see if we can find one that does even better from the perspective of moral uncertainty. Finally, we suggested above that the *pro tanto* public justification principle generates stronger reasons when a greater number of reasonable people reject a law and when they have stronger objections; future work could ask what political institutions would best realize such a principle, and how this relates to potential epistemic arguments for democracy under moral uncertainty. We think that there is much fruitful work to be done in this area and, more generally, at the intersection of moral uncertainty and political philosophy.¹¹

11. For helpful comments and discussion, we would like to thank Sameer Bajaj, Paul Billingham, Allen Buchanan, Alexander Motchoulski, Sarah Raskoff, Anthony Taylor, Teruji Thomas, Kevin Vallier, and especially Christian Tarsney. Thanks also to two anonymous referees, as well as to participants in the Global Priorities Institute Work in Progress Group and in the Centre for the Study of Social Justice Seminar Series, both at the University of Oxford.

Bibliography

- Bajaj, S. 2017. "Self-defeat and the Foundations of Public Reason." *Philosophical Studies* 174 (12): 3133–51. <https://doi.org/10.1007/s11098-016-0850-9>.
- Barrett, J., and G.F. Gaus. 2020. "Laws, Norms, and Public Justification: The Limits of Law as an Instrument of Reform." In *Public Reason and Courts*, edited by S.A. Langvatn, M. Kumm, & W. Sadurski, 201–28. Cambridge: Cambridge University Press. <https://doi.org/10.1017/9781108766579>.
- Barry, B. 1996. *Justice as Impartiality*. Oxford: Clarendon Press.
- Barry, C., and P. Tomlin. 2019. "Moral Uncertainty and the Criminal Law." In *Handbook of Applied Ethics and the Criminal Law*, edited by K. Ferzan and L. Alexander, 445–69. New York: Palgrave.
- Bespalov, A. 2021. "Against Public Reason's Alleged Self-Defeat." *Law and Philosophy* 40: 617–44. <https://doi.org/10.1007/s10982-021-09418-6>.
- Bukoski, M. 2021. "Moral Uncertainty and Distributive Sufficiency." *Ethical Theory and Moral Practice* 24 (4): 949–63. <https://doi.org/10.1007/s10677-021-10236-x>.
- Clarke, S. 1999. "Contractarianism, Liberal Neutrality, and Epistemology." *Political Studies* 47 (4): 627–42. <https://doi.org/10.1111/1467-9248.00221>.
- D'Agostino, F. 1996. *Free Public Reason: Making It Up As We Go*. Oxford: Oxford University Press.
- Director, S. 2019. "Global Public Reason, Diversity, and Consent." *Philosophical Papers* 48 (1): 31–57. <https://doi.org/10.1080/05568641.2019.1584541>.
- Ebels-Duggan, K. 2010. "The Beginning of Community: Politics in the Face of Disagreement." *The Philosophical Quarterly* 60 (238): 50–71. <https://doi.org/10.1111/j.1467-9213.2008.591.x>.
- Enoch, D. 2013. "The Disorder of Public Reason." *Ethics* 124 (1): 141–76. <https://doi.org/10.1086/671386>.
- Enoch, D. 2017. "Political Philosophy and Epistemology: The Case of Public Reason." *Oxford Studies in Political Philosophy* 3: 132–65. <https://doi.org/10.1093/oso/9780198801221.003.0007>.
- Feinberg, J. 1989. *The Moral Limits of the Criminal Law Volume 3: Harm to Self*. Oxford: Oxford University Press.
- Gaus, G. 2011. *The Order of Public Reason: A Theory of Freedom and Morality in a Diverse and Bounded World*. Cambridge: Cambridge University Press.
- Gaus, G. 2015. "On Dissing Public Reason: A Reply to Enoch." *Ethics* 125 (4): 1078–95. <https://doi.org/10.1086/680904>.
- Gracely, E.J. 1996. "On the Noncomparability of Judgments Made by Different Ethical Theories." *Metaphilosophy* 27 (3): 327–32. <https://doi.org/10.1111/j.1467-9973.1996.tb00212.x>.
- Greaves, H., and O. Cotton-Barratt. 2024. "A Bargaining-Theoretic Approach to Moral Uncertainty." *Journal of Moral Philosophy* 21 (1–2): 127–169. <http://doi.org/10.1163/17455243-20233810>.
- Gustafsson, J.E., and O. Torpman. 2014. "In Defence of My Favourite Theory." *Pacific Philosophical Quarterly* 95 (2): 159–74. <https://doi.org/10.1111/papq.12022>.
- Larmore, C. 1990. "Political Liberalism." *Political Theory* 18 (3): 339–60. <https://doi.org/10.1177/0090591790018003001>.
- Leland, R.J. (2019). "Civic Friendship, Public Reason." *Philosophy & Public Affairs* 47 (1): 72–103. <https://doi.org/10.1111/papa.12141>.
- Leland, R.J., and H. van Wietmarschen. 2017. "Political Liberalism and Political Community." *Journal of Moral Philosophy* 14 (2): 142–67. <https://doi.org/10.1163/17455243-46810052>.
- Lister, A. 2013. *Public Reason and Political Community*. New York: Bloomsbury Publishing.
- Lockhart, T. 2000. *Moral Uncertainty and Its Consequences*. Oxford: Oxford University Press.
- MacAskill, W. 2016. "Normative Uncertainty as a Voting Problem." *Mind* 125 (500): 967–1004. <https://doi.org/10.1093/mind/fzv169>.
- MacAskill, W., K. Bykvist, and T. Ord. 2020. *Moral Uncertainty*. Oxford: Oxford University Press.

- MacMullen, I. 2023. "Justified Coercion, Political Cooperation, and Exemption from General Laws." *The Journal of Politics* 85 (1), 153–65. <https://doi.org/10.1086/720328>.
- Muldoon, R. 2016. *Social Contract Theory for a Diverse World: Beyond Tolerance*. New York: Routledge. <https://doi.org/10.4324/9781315545882>.
- Nagel, T. 1987. "Moral Conflict and Political Legitimacy." *Philosophy and Public Affairs* 16 (3): 215–40.
- Nozick, R. 1969. "Coercion." In *Philosophy, Science, and Method: Essays in Honor of Ernest Nagel*, edited by S. Morgenbesser, 440–72. New York: St Martin's Press.
- Quong, J. 2011. *Liberalism Without Perfection*. Oxford: Oxford University Press.
- Rawls, J. 1985. "Justice as Fairness: Political Not Metaphysical." *Philosophy and Public Affairs* 14 (3): 223–51.
- Rawls, J. 2001. *Justice as Fairness: A Restatement*. Cambridge: Harvard University Press.
- Rawls, J. 2005. *Political Liberalism*. New York: Columbia University Press.
- Raz, J. 1990. "Facing Diversity: The Case of Epistemic Abstinence." *Philosophy and Public Affairs* 19 (1): 3–46.
- Ross, J. 2006. "Rejecting Ethical Deflationism." *Ethics* 116 (4): 742–768. <https://doi.org/10.1086/505234>.
- Sepielli, A. 2009. "What to Do When You Don't Know What to Do." *Oxford Studies in Metaethics* 4: 5–28.
- Sepielli, A. 2013. "What to Do When You Don't Know What to Do When You Don't Know What to Do ..." *Nous* 47 (1): 521–44. <https://doi.org/10.1111/nous.12010>.
- Tarsney, C. 2019. "Normative Uncertainty and Social Choice." *Mind* 128 (512): 1285–1308. <https://doi.org/10.1093/mind/fzy051>.
- Trammell, P. 2019. "Fixed-Point Solutions to the Regress Problem in Normative Uncertainty." *Synthese* 198 (2): 1177–99. <https://doi.org/10.1007/s11229-019-02098-9>.
- Valentini, L., and C. List. 2020. "What Normative Facts Should Political Theory Be About? Philosophy of Science Meets Political Liberalism." *Oxford Studies in Political Philosophy* 6: 185–220.
- Vallier, K. 2011. "Convergence and Consensus in Public Reason." *Public Affairs Quarterly* 25 (4): 261–80.
- Vallier, K. 2019. *Must Politics Be War?: Restoring Our Trust in the Open Society*. Oxford: Oxford University Press.
- Vallier, K. 2022. "Public Justification." In *The Stanford Encyclopedia of Philosophy* (Winter 2022), edited by E. N. Zalta. <https://plato.stanford.edu/entries/justification-public/>.
- Van Schoelandt, C. 2019. "Convergence in the Political Liberal Community." *Public Reason* 11 (2): 3–18.
- Waldron, J. 1993. *Liberal Rights: Collected Papers 1981–1991*. Cambridge: Cambridge University Press.
- Wall, S. 2002. "Is Public Justification Self-Defeating?" *American Philosophical Quarterly* 39 (4), 385–94.
- Wall, S. 2010. "On Justificatory Liberalism." *Politics, Philosophy & Economics* 9 (2): 123–49. <https://doi.org/10.1177/1470594X09345677>.
- Weatherston, B. 2014. "Running Risks Morally." *Philosophical Studies* 167 (1): 141–63. <https://doi.org/10.1007/s11098-013-0227-2>.
- Wendt, F. 2019. "Rescuing Public Justification From Public Reason Liberalism." *Oxford Studies in Political Philosophy* 5: 39–64. <https://doi.org/10.1093/oso/9780198841425.003.0002>.