Does Public Reason Undermine Social Justice?[1]

Andrew Lister, July 4, 2011. Comments welcome: andrew.lister@queensu.ca

One persistent worry about public reason is that demanding that the exercise of political power be publicly justifiable would rule out redistributive policies aimed at promoting social justice, not just perfectionist policies aimed at promoting controversial conceptions of the good. The main liberal argument for anti-perfectionism, Simon Caney explains, is that "principles of justice may only draw on moral considerations which cannot be reasonably rejected by anyone," and that since there is reasonable controversy about conceptions of the good life, principles of justice should not be predicated upon assessment of these conceptions of the good.[2] There is also reasonable disagreement about "issues of justice," however, which seems to imply that "the state should abstain from enacting principles of justice in many of the areas we would expect the state to act (in the economy, defence, and punishment)."[3] If we are not willing to admit perfectionism, the commitment to public justifiability seems to lead "to a fairly libertarian view of legitimate state action," in the words of Jonathan Quong.[4] Faced with this unpalatable consequence, we

---

1.   This paper is a draft chapter from a manuscript entitled "Public Reason and Political Community." One of the purposes of the book is to argue that public reason is a condition of community in a pluralistic society. This is only true, however, if we construe the demand for public justifiability as a constraint on reasons for political decisions, with a default of exclusion from consideration, as opposed to a constraint on coercive state action directly, with a default of inaction. So the paper fits into a more general project of distinguishing different varieties of public reason, criticizing the libertarian varieties, and (in other chapters) articulating the point of what for want of a better word I think of as the communitarian variety. I would like to thank Kevin Vallier and Matthias Brinkman for their comments on earlier drafts of this chapter.

2.   Simon Caney, "Liberal Legitimacy, Reasonable Disagreement and Justice," in *Pluralism and Liberal Neutrality*, ed. Richard Bellamy, and Martin Hollis (Illford, Essex: Frank Cass Publishers, 1999), 21.

3.   Caney, "Liberal Legitimacy, Reasonable Disagreement and Justice," 23. Justificatory liberals could of course deny that there is room for reasonable disagreement about justice, in the same way that there is about conceptions of the good, but many people find this asymmetry claim implausible. "While it is certainly true that people in modern democratic societies hold a variety of conflicting moral and religious views, it cannot be said that there is a 'fact of reasonable pluralism' about morality and religion that does not also apply to questions of justice;" Michael Sandel, *Liberalism and the Limits of Justice, 2nd Edition* (Cambridge: Cambridge University Press, 1998), 196, 203-04. See also Bruce Brower, "The Limits of Public Reason," *Journal of Philosophy* 91 (1994), 21; Joseph Chan, "Legitimacy, Unanimity, and Perfectionism," *Philosophy and Public Affairs* 29, no. 1 (2000), 21-22.

4.   Jonathan Quong, "Disagreement, Asymmetry, and Liberal Legitimacy," *Politics, Philosophy, Economics* 4, no.

might reject the fundamental liberal requirement of unanimous reasonable acceptability, but only at the cost of also abandoning liberal anti-perfectionism, generating what Caney calls the anti-perfectionist dilemma.[5] If there is room for reasonable disagreement about distributive justice, as well as about human flourishing, liberals must either water down their standard of public justifiability, allowing some perfectionism, or "bite the bullet," and sacrifice some justice-based policies.[6] The purpose of this chapter is to assess the most important responses to this dilemma.

The first section of the chapter characterizes the anti-perfectionist dilemma as a choice between the two pure models of public justification described in Chapter 1: the reasons-for-decisions / coercion model and the coercion / convergence model of public justification. The second section outlines the response based on higher-order unanimity, which involves aggregating or bundling separate binary decisions into one decision between policies **A**, **B**, and **0** (no state action), and then appealing to the nesting of *in*conclusive justification between **A** and **B** by conclusive justification of **A** and **B** as against **0**. Section 3 poses two questions about this response: (1) whether to apply the demand for unanimous idealized acceptability non-incrementally, to the choice between some coercion and none, or incrementally, to the choice between more and less, and; (2) at what level of aggregation or bundling of polices we should apply the acceptability requirement. The argument here is that if one is unwilling to adopt the incremental version of the principle, one must adopt a rule of maximal feasible disaggregation. Section 4 sets out a criteria for feasible disaggregation, the *unanimous unambiguous rankability* rule, but also argues that

---

3 (2005), 303; Jonathan Quong, *Liberalism Without Perfection* (Oxford: Oxford University Press, 2011), 193.

5.   Caney, "Liberal Legitimacy, Reasonable Disagreement and Justice," 24.

6.   Caney, "Liberal Legitimacy, Reasonable Disagreement and Justice," 29.

maximal feasible disaggregation can disaggregate too much, if the coerciveness of package policies is not proportional to the number of component policies. The conclusion of these sections is that in order to use the argument from higher-order unanimity to rule out perfectionism while not ruling out redistribution one must make quite specific assumptions about the set of reasonable moral preferences and about the measurement of coercion. Section 5 argues that reasonable people will naturally disagree about these matters, with the result that they will also disagree about which laws and policies are in the eligible set. Given this disagreement, the principle of public justification cannot be a principle that sets a threshold for laws and institutions to be authoritative. Public justification can only be an ideal of full justification, as envisaged by the reasons-for-decisions framing of the principle.

1.   Consensus, Coercion, and the Anti-Perfectionist Dilemma

The way the social justice objection is usually presented trades on an ambiguity about the object of public justification. Public justifiability requires that *principles* of justice be acceptable to all reasonable points of view, Caney says, but it is controversy about *laws and policies* related to justice that leads to the dilemma. There is undoubtedly controversy about what level of taxation is just, about the death penalty, about affirmative action, and so on. Such controversy is rooted in part in the fact that different comprehensive doctrines take different positions on these questions. Controversy also arises from the fact that even people who accept Rawlsian public reason may disagree about the precise weighting and interpretation of their shared, public reasons, in disputed historical and sociological contexts. A given principle might be acceptable to all reasonable points of view, even while its implications for a specific matter of public policy are open to reasonable disagreement. If public justification simply requires that political decisions

be made on the basis of such principles (public reasons), controversy about specific laws does not block state action. It is only when we frame the principle of public justification as a constraint on coercion directly, with a default of inaction, that reasonable disagreement threatens to undermine collective efforts to secure social justice.

The anti-perfectionist dilemma can be characterized as a choice between these two ways of framing the principle of public justification's unanimous acceptability requirement: as a constraint on coercive state action directly, with a default of not having any common rule on the matter in question, or as a constraint on the reasons for political decisions, with a constraint of exclusion from deliberation. The worry about social justice is clearly plausible if we think of public justification as applying to state action directly. If redistribution of wealth counts as coercive state action, then all it takes is for one reasonable view to object and the redistribution is not legitimate. If the demand for public justifiability applies at the level of reasons, however, the worry is much less pressing. The exclusion of non-public reasons from choices about the basic structure of society may well rule out some strongly egalitarian views, but it will also presumably rule out appeal to some highly inegalitarian views as well. The reasons-for-decisions or consensus approach to public justification is not a panacea for liberals, however, because it will not yield a strict anti-perfectionism, as we saw in Chapter 2. Thus, one way of presenting the anti-perfectionist dilemma is as a choice between a reasons-for-decisions / consensus account of public justification, which will permit some perfectionism, or a coercion / convergence account of public justification, which will rule out perfectionism but at the expense of also ruling out many reasonable but also reasonably contestable egalitarian policies.

This account of the anti-perfectionist dilemma in terms of consensus and convergence

justification raises a doubt about Jonathan Quong's recent attempt to show that political liberalism escapes the anti-perfectionist dilemma. As we saw in Chapter 6, Quong and Schwartzmann have shown that political liberalism is open to reasonable disagreement about matters of justice.[7] State action is legitimate so long as plausibly justifiable based on public reasons alone, even if there is reasonable disagreement about where the balance of public reasons lies. The unanimous acceptability test is applied to reasons, not to laws directly, allowing scope for differences in judgment about ranking, interpretation, and application of shared i.e. public reasons. Quong refers to disagreement based on conflicting rankings or applications of shared values or principles as "justificatory," and disagreement based on conflicting ultimate convictions as "foundational."[8] Only foundational disagreement blocks state action. Reasonable disagreements about justice are *necessarily* justificatory rather than foundational, Quong claims, while reasonable disagreements about the good *tend* to be foundational.[9] To count as "reasonable," views about justice need to accept the burdens of judgment, reasonable pluralism, the idea of society as a fair system of cooperation, and the ideal of public justifiability. Because they have these substantive commitments, reasonable arguments about justice will always appeal to legitimate public values. Reasonable disagreements about justice will therefore always be justificatory. The problem with this argument is that if we allow definitional filtering of arguments about justice by the criterion

---

7. Micah Schwartzman, "The Completeness of Public Reason," *Politics, Philosophy, Economics* 3 (2004), 198-203; Quong, "Disagreement, Asymmetry, and Liberal Legitimacy," 316; c.f. John Rawls, "The Idea of Public Reason Revisited,", The Law of Peoples (Cambridge, Mass.: Harvard University Press, 1999), 169.

8. Quong, "Disagreement, Asymmetry, and Liberal Legitimacy," 311-13. "Foundational disagreements are characterized by the fact that the participants do not share any premises which can serve as a mutually acceptable standard of justification. The second type of disagreement, justificatory disagreement, occurs when participants do share premises that serve as a mutually acceptable standard of justification, but they nevertheless disagree about certain substantive conclusions;" Quong, *Liberalism Without Perfection*, 204.

9. Quong, *Liberalism Without Perfection*, 215.

of reasonableness, we should allow the same kind of filtering of arguments about the good life.

It is true that reasonable disagreement about the good may be foundational, if we use the minimal, Larmorean criterion of reasonableness i.e. willingness and ability to reason sincerely with others about what is true and good and right. Yet the same is true of reasonable disagreements about justice, if we stick with the minimal sense of reasonableness. Once we add the more demanding, substantive criterion of reasonableness (acceptance of burdens of judgment, reasonable pluralism, public justification, etc.), reasonable disagreements about justice are necessarily justificatory rather than foundational, but then the same is true of reasonable disagreements about the good, in this more demanding sense. That is to say, when persons who are reasonable in the stronger sense disagree about the good, they do so on the basis of underlying conceptions of value that are not reasonably rejectable. If reasonableness involves a shared commitment to the principle of public justification, and this principle demands only plausible and not conclusive public justification, there is the possibility of state action in the context of conflicting reasonable judgments about the good life, where these judgments are about the ranking or interpretation of shared values. It is therefore not clear that Quong's response escapes the dilemma of either allowing more perfectionism or ruling out egalitarian redistribution.[10]

2.   The Higher-Order Unanimity Response

Egalitarians who are attracted to the idea of public justification have another way of

---

10. There will be legitimate questions about whether policies appealing to uncontroversial aspects of human flourishing are objectionably paternalistic, as for example in the case of the laws against smoking discussed in Chapter 2(laws based on the uncontroversial value of health, but also on a controversial ranking of pleasure relative to health). However, I want to separate the issues of paternalism and public justification. As argued in Chapter 2 one can be strongly opposed to paternalism without any commitment to public justification, and for some people, coming to accept the principle of public justification will make them less anti-paternalistic than they would otherwise have been.

avoiding the anti-perfectionist dilemma, which is to frame the principle as a constraint on coercive state action directly, but to make use of the argument from higher-order unanimity or nested inconclusiveness. Up until now, I have presented the problem of public justification in the context of binary choice, between a law **L** and **¬L** (no law). As soon as we allow for there to be more than one alternative, the idea of unanimous reasonable acceptability becomes ambiguous. Suppose we face a choice between **A**, **B** and **0** (**0** being the proposal of not having any law, on the subject at hand). In order to be adequately publicly justified, are we requiring that **A** be reasonably non-rejectable versus **0** only, or versus **B** as well? It would be irrational to conclude that the reasonable rejectability of **A** in favor of **B** forces us to settle for **0**, since in some instances reasonable people might agree that **0** is worse than **A** *or* **B**.[11] Our reasonable disagreement about which of the two possible laws is best is nested by our reasonable agreement on the preferability of either law to the alternative of none at all. A law will be proper, it would seem, if no one can reasonably deny that it is morally better than no policy at all, and if it is selected by some reasonable[12] procedure from the set of such policies. Thus Gaus is led to develop the idea of the "optimal eligible set" of proposals.[13] A proposal is optimal if not all members of the public agree

11. Gerald F. Gaus, *Contemporary Theories of Liberalism* (Thousand Oaks, California: Sage, 2003), 217.

12. Gaus argues, correctly I think, that the procedure used to select from the eligible set need only be reasonable, not invulnerable to reasonable rejection, in order for the law chosen to be legitimate; Gerald F. Gaus, *The Order of Public Reason: A Theory of Freedom and Morality in a Diverse and Bounded World* (Cambridge: Cambridge University Press, 2010), 392.. Requiring that a procedure be invulnerable to reasonable rejection as against all other procedures is equivalent to requiring conclusive justification as best. However, if we all favour one of some set of procedures over no procedure, it would be irrational to conclude that we must default to no collective decision procedure at all, simply because none of these eligible procedures is unanimously accepted as best. The reasoning is the same whether we are talking about policies or procedures.

13. Gerald F. Gaus, "Coercion, Ownership, and the Redistributive State: Justificatory Liberalism's Classical Tilt," *Social Philosophy and Policy* 27 (2010), 247-50; Gerald F. Gaus, "On Two Critics of Justificatory Liberalism: A Response to Wall and Lister," *Politics, Philosophy, Economics* 9, no. 3 (2010), 196-97; Gaus, *The Order of Public Reason*, 321-33, 391-409.

that it is morally inferior to some other proposal in the set. In others words, if we all agree that **A** is worse than **B**, **A** is not optimal, and there is no need to consider it. A proposal is eligible if all members of the public agree that it is morally superior to the default proposal of not having any law, on the matter at hand. In other words, if anyone reasonably thinks that **A** is worse **0**, **A** is ineligible. Gaus provides the following example:

	Members of the Public			
	Alf	**Betty**	**Charlie**	**Doris**
Proposals deemed accept-able by the individual	A	B	C	D
	B	A	A	A
	C	D	B	B
			D	
Default Proposal (Liberty; no law)	0			
Proposals deemed unacceptable	D	C		C

Either **A** or **B** is preferred by every member of the public to no law at all on the matter. **C** is deemed worse than nothing by Betty and Doris, while **D** is deemed worse than nothing by Alf. Thus only **A** and **B** are eligible. The four members of this public "all have a conclusive reason to select from this set, for both **A** and **B** are, from everyone's evaluative standards, improvements on the absence of legislation or the condition of liberty."[14]

---

14. Gaus, "Coercion, Ownership, and the Redistributive State," 250. I have changed Gaus's subscripted L's for letters, because it seems simpler, and this is the notation I will elsewhere in the chapter.

This idea of higher-order unanimity or nested inconclusiveness allows Gaus to avoid the conclusion that only a minimal, merely protective state is legitimate. Economic redistribution can be legitimate, if we apply the argument from higher order unanimity to economic systems. Nagel was uncertain whether egalitarian economic policies could be publicly justified, because it didn't seem *unreasonable* to reject taxation for redistribution, in the same way that it seems un-reasonable to reject having a common military policy.[15] It is unreasonable to reject the idea of having a system of enforced property rights, however, and if we consider welfare schemes and estate taxes to be simply a part of a more egalitarian system of property rights, 'redistributive' policies would be legitimate if chosen democratically, or by some other reasonable procedure. After all, property is backed by state coercion, the threat of use of collective force should some persons invade the property rights of others.[16] The escape clause for justificatory liberals who would be egalitarians thus consists in arguing that taxation for redistribution is part of a more egalitarian system of property rights, and that having one of a range of such systems is better than having none at all, and that because both egalitarian and a libertarian (i.e. a right-libertarian) systems of property rights belong to this eligible set, both are legitimate if chosen democratically.

3.   Two Questions

The argument from higher-order unanimity depends on characterizing the space of policy

---

15. Thomas Nagel, "Moral Conflict and Political Legitimacy," *Philosophy and Public Affairs* 16, no. 3 (1987), 234.

16. Michael Blake, "Distributive Justice, State Coercion, and Autonomy," *Philosophy and Public Affairs* 30, no. 3 (2002), 277. Blake cites Jeremy Waldron, *The Right to Private Property* (Oxford [Oxfordshire] New York: Clarendon Press Oxford University Press, 1988)., as well as *Shelley v. Kraemer*, 334 U.S. 1 (1948), the U.S. Supreme Court case that found that enforcement of contracts involved state action, and hence that constitutional demands for equal protection ruled out enforcement of contracts involving racial restrictions.

choices as one involving two or more options each of which is superior to the third option of not

having any policy at all, on the matter at hand, instead of as two separate binary choices between

action and inaction. This aggregation or bundling of policies so as to create a 3+ option space

raises two questions. The first is whether the demand for unanimous acceptability should be ap-

plied only to the choice between some coercion and none, instead of to the choice between more

coercion and less. The second is what the proper level of aggregation of policies is, for purposes

of applying the principle.

3.1. Incremental vs. Non-Incremental Principles

The argument from higher-order unanimity is open to the following objection. Libertari-

ans and classical liberals may admit that the protection of property involves coercion, but claim

that their preferred system of property rights with minimal redistribution of wealth and income

involves less coercion than does the egalitarian system.[17] The egalitarian scheme of property

rights (including generous 'redistribution') is publicly justified with respect to a state of nature

but not with respect to the authentically liberal scheme, which must carry the day, as the less co-

ercive, and hence default policy. If we want *all* coercion to be justified, we need to show that a

given system involves no excess coercion – no coercion that can be reasonably rejected – not just

that a particular package of coercion is better than none at all.[18] To be a member of the eligible

set, it is not sufficient that an option be unanimously reasonably preferable to the inactive / non-

---

17. Gaus himself claims that the redistributive state championed by contemporary egalitarians is more coercive than the classically liberal state he favours; Gaus, "Coercion, Ownership, and the Redistributive State," 261-62; Gaus, *The Order of Public Reason*, 509-46.

18. A comparison with Rawls's difference principle may be helpful. To justify a given level of inequality, it is not sufficient to show that that the least well-off are better off than they would be in a state of nature. To justify all of the inequality, we would have to show that each extra bit of inequality helps, or at least doesn't harm, the least well off.

coercive policy of having no policy at all; instead, it must be unanimously reasonably preferable to all less coercive policies. The objection thus relies on an *incremental* version of the public justifiability principle. Conclusive public justification is required for each increment of coercion, permitting democratic choice only between equally coercive policies unanimously reasonably preferred to less coercive policies.

Egalitarians might accept the incremental version of the principle while asserting that their ideal system (including so-called "redistribution") is no more active or coercive than the libertarian system. The absence of redistribution simply involves the state using its coercive powers to exclude citizens from different properties, egalitarians can argue; if the government is not forcibly redistributing some part of your wealth, it is forcibly excluding others from that part of your wealth. Even if egalitarians are right about how to measure coercion with respect to property rights, however, in general the incremental version of the principle is demandingly libertarian. The more fine-grained one's measurement of degrees of coerciveness, the closer the most libertarian of reasonable citizens comes to being a dictator. As we make ever smaller distinctions between the levels of coerciveness of rival policies, we place correspondingly greater weight on the views of the barely reasonable libertarian.

Gaus rejects the incremental interpretation of the principle on the grounds that it confuses conclusive justification as eligible (better than nothing) with conclusive justification as best, which Gaus refers to as "simple conclusive justification."[19] He imagines a scenario in which a public composed of Alf and Betty is deciding between laws **A** and **B**, with **A** involving x amount of state coercion, and **B** involving $x + y$ amount of state coercion.

---

19. Gaus, "On Two Critics of Justificatory Liberalism," 197.

Alf says to Betty: "I hold that *A* is conclusively justified, but I disagree that the additional *y* coercion involved in *B* is justified. So *B* cannot be conclusively justified to me. On the other hand, if you think that x+y degree of coercion is conclusively justified, you must think that x coercion is conclusively justified, and so you must agree that A is conclusively justified. Thus, *A,* but not *B* is conclusively justified between us...[20]

Gaus comments that "Alf's crucial claim (that his less coercive alternative is conclusively justified in the simple sense to Betty) is false because she believes that the real benefits of coercion only set in when higher levels are reached."[21] It is true that Betty's belief that **B** is optimal does not by itself imply that she thinks **A** better than nothing. But if Betty doesn't think **A** better than nothing, **A** is ineligible, according to the incremental version of the principle. The incremental public justifiability principle holds that for an option to be eligible, relative to some set of possibilities, it must be unanimously preferred to all less coercive options, and not unanimously rejected in favour of any more coercive option. On this standard, if the set of options is {**0**, **A**, **B**}, and Betty prefers **0** to **A**, **A** is ineligible. Suppose to the contrary that Betty does think **A** better than nothing. When Alf applies the incremental version of the public justifiability principle, he does not claim that he has conclusively justified **A** as best. He recognizes the reasonableness of Betty's ranking **B** first, but simply points out that she in fact accepts that **A** is better than nothing, and then claims that the additional coercion involved in **B** must be acceptable from all qualified points of view (i.e. his as well as hers). The public justifiability principle imposes a constraint of multi-perspectival acceptability on the claim that a law's moral benefits outweigh its coercion costs. The incremental version of the principle simply applies this constraint continuously, to the decision to have more rather than less coercion, rather than as a step function, to the decision to
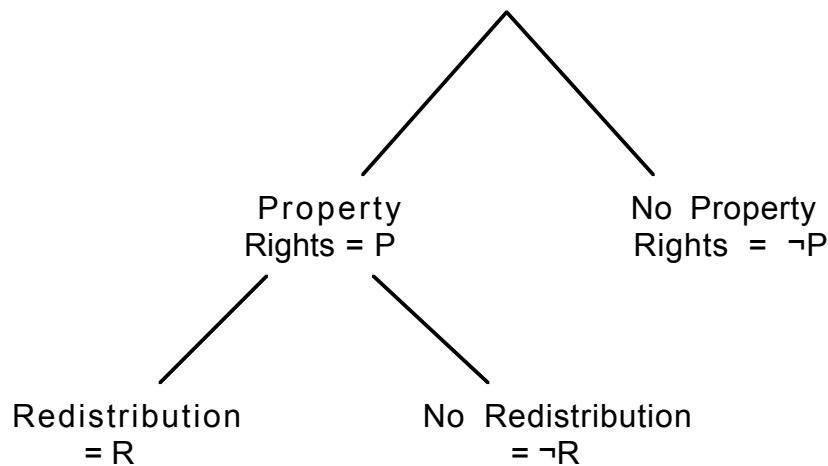
20. Gaus, "On Two Critics of Justificatory Liberalism," 201.

21. Gaus, "On Two Critics of Justificatory Liberalism," 201.

have some coercion rather than none.

3.2. The Aggregation or 'Zoom' Problem

      The demand for conclusive justification requires a default that obtains when justification is inconclusive. Yet this demand can be applied at different levels, to different descriptions of the space of choices, yielding different default policies. The argument from 'higher-order unanimity' makes use of this fact. It initially seems that redistribution is illegitimate because we apply the criterion of conclusive justification to two choices separately: whether to have a system of property rights (conclusively 'yes'), and whether to tax for redistribution (inconclusive, hence no).

```
                                /\
                               /  \
                              /    \
                             /      \
                            /        \
                           /          \
                Property              No Property
                Rights = P           Rights  =  ¬P
                  /\
                 /  \
                /    \
               /      \
  Redistribution      No Redistribution
      = R                 = ¬R
```

In this schema, we apply the qualified unanimity criterion first to P vs. ¬P, yielding P, and then to R vs. ¬R, yielding ¬R. The three-option description aggregates these choices:

      A : P, R = property rights plus redistributive taxation

      B : P, ¬R = enforcement of property rights, no redistributive taxation

      0 : ¬P, ¬R = no enforcement of property rights (hence no redistributive taxation either)

A and **B** are inconclusively justified against each other, but each is conclusively justified as opposed to **0**. Hence **A** and **B** are both members of the eligible set, and so legitimate if chosen

democratically. When the requirement of conclusive justification is applied to the single, three-option choice, it yields a different result than it does when applied to the two binary choices separately.

Similar reasoning in the case of restrictions on smoking could justify rather severe restrictions. Consider a proposed law banning smoking in bars and restaurants. Such a law is coercive, and reasonably rejectable, let us suppose, so it initially seems impermissible. Yet the government also bans smoking on airplanes and requires restaurants to have non-smoking sections. If we aggregate these various policies, as the property examples suggests we can, we might face a choice between the following options:

A = Severely restrictive laws vs. smoking
B = Mildly restrictive laws vs. smoking
0 = No laws vs. smoking

If **0** meant that people would be allowed to smoke on airplanes, it might be unreasonable to prefer this option to the mildly restrictive *or* the severe policy, either of which would then be permissible if chosen democratically. Notice that one might also argue that the stricter law against smoking is more coercive than the more permissive law, and therefore must be conclusively justified as against the milder anti-smoking regime, not just against the option of having no laws at all about smoking. If one adopts the non-incremental interpretation of the principle, however, one will have to rely on some criterion about the level of aggregation of policies in order to rule out the democratic legitimacy of the highly restrictive policies, while not thereby ruling out the legitimacy of redistribution as well.

The argument from higher-order unanimity might also licence significant restrictions in the case of abortion. Disaggregated, we face two choices: whether to have a law against abortion

(inconclusive, hence 'no'), and whether to have a law against murder (conclusively 'yes'). Yet we could consider prohibition of abortion to be an element of the prohibition of murder, and so describe our options as follows:

A = Broad law against murder (i.e. some abortions are prohibited, along with ordinary murders)

B = Narrow law against murder (i.e. no abortions are prohibited, only murders)

0 = No law against murder.

'No law' is now the non-coercive / inactive baseline, but this anarchical default is reasonably rejectable in favor of either the narrow or the broad murder law. Applied to this description of the space of choices, the requirement of conclusive justification permits some laws against abortion, so long as such laws are chosen democratically from among the set of murder laws that are rationally preferable to no law at all. Moreover, if we are allowed to aggregate policies to this level, why not even further? Instead of demanding conclusive justification for cultural and economic policies separately, we could apply the criterion to the choice between a perfectionist state (i.e. one that aims to improve people's lives for their own sakes, based on reasonably contestable ideals of a good life), a non-perfectionist state, and no state at all. If the inactive/non-coercive baseline is no state at all, then some fairly strong forms of perfectionism would be legitimate, if chosen democratically. As in the smoking example, one might invoke an incremental version of the public justifiability principle. The broader the state's definition of murder, one might argue, the more it is engaged in coercion, as discussed at the end of the last chapter. On the incremental version of the principle, every increase in the scope of the state's law against murder would have to be conclusively justified as against all narrower definitions. If we insist on the non-incremental version of the principle, however, the only way to resist the democratic legitimacy of highly restrictive abortion laws would be to appeal to an aggregation criterion that requires a lot of dis-

aggregation, again raising doubts about the legitimacy of aggregation in the property and redistribution case.

The problem with the argument from higher-order unanimity is therefore that there are different ways to divide up policies, and so different ways to specify the inactive or noncoercive baseline. When the qualified acceptability requirement is applied at a disaggregated level to separate choices between action and inaction – zoomed in, as it were – the principle has libertarian consequences. Conversely, it seems possible to legitimize a lot of very coercive state action by 'zooming out', and applying the unanimity criterion at a higher level of aggregation, to three-option cases in which the inactive baseline approaches the state of nature. It will be tempting for egalitarian liberals to apply the criterion of higher-order unanimity at a higher level of aggregation to permit redistribution, although not so high as to permit paternalist perfectionism, but then at lower level to rule out the prohibition of abortion. Yet this results-oriented reasoning seems unprincipled.

3.3. Incremental Public Justification and Disaggregation

There is a close connection between the question of whether to adopt an incremental version of the principle and the question of the level of aggregation at which we should apply the principle. The greater the level of aggregation, the closer the non-coercive default will be to a state of nature, the less demanding will be the requirement of unanimous acceptability, and the more tempting it will be to interpret the principle of public justifiability in incremental terms. Suppose we are considering binary policy domains X, Y, and Z. In the disaggregated space, X has to be unanimously acceptable with respect to ¬X, Y unanimously acceptable with respect to ¬Y, and Z unanimously acceptable with respect to ¬Z. Aggregated, we have a set of 8 options to

consider including the default: XYZ, XY¬Z, X¬YZ, ¬XYZ, X¬Y¬Z, ¬X¬YZ, ¬XY¬Z, ¬X¬YZ, any of which is eligible just in virtue of being unanimously superior to ¬X¬Y¬Z. As the level of aggregation rises, the default compared to which a package of policies must be invulnerable to reasonable rejection approaches anarchy, and the size of the optimal eligible set increases, since many different packages of policies, some highly coercive, will be better than not having any common rule in *any* of these areas of policy. The higher the level of aggregation, the weaker the criterion of public justifiability becomes, and the more the incremental version of the principle will look tempting, in order to exclude the case in which highly coercive perfectionism is democratically legitimate simply because it is part of a package of rules and policies that are together better than no package at all.

Consider a case of questionable aggregation: whether to have a system of property rights, and whether to use public money to support high culture. The qualified unanimity requirement applied to each question separately would rule in property rights, and rule out public support for culture. Aggregating these options (i.e. 'zooming out') yields to the following set of options:

A = Property rights, support for culture, or

B = Property rights, no support for culture

0 = No property rights (and so no support for culture either).

This bundling of policies seems arbitrary, since culture and property seem to be separate issues. The only reason the promotion of high culture passes the qualified acceptability standard is that we have bundled it with the protection of property, which is such an essential policy that if we had to choose between protecting property and promoting culture, on the one hand, or doing neither, it would be unreasonable to choose 'neither'. The question is why it should be sufficient, for the promotion of culture to be legitimate if chosen democratically, that state promotion of

culture is not such a serious bad that it outweighs the good of protecting property. Surely the standard should be whether it is reasonable to reject the promotion of culture holding the protection of property constant. Otherwise, we are not requiring that all coercion be invulnerable to reasonable rejection, but just that any given coercive policy be part of a package of coercive policies that is not reasonably rejectable as compared to the absence of the whole package. It is because we want all coercion to pass the qualified unanimity requirement that we should not arbitrarily lump policies together, allowing reasonably rejectable coercion of one kind (e.g. state promotion of high culture) to piggy back or free ride on the strength of the reasonably non-rejectable case for coercion of another kind (e.g. state protection of property rights). Arbitrary bundling of policies (i.e. over-aggregation) allows reasonably rejectable coercion to slip through the cracks, and gain legitimacy without being truly publicly justifiable, if we are using a non-incremental version of the principle. Conversely, if the incremental version of the principle were not justified, it would not be clear why there is any problem with the arbitrary aggregation of policies (i.e. over-aggregation, 'zooming out' too much).

Moreover, if the degree of coercion contained in a package policy is proportional to the number of component policies and disaggregation is feasible, disaggregation mimics the incremental version of the principle. Take the simplest case, where there are only three options and the degree of coercion is equal to the number of component policies:

Policy	Degree of State Coercion
A : Property rights, support for culture (P, C)	2
B : Property rights, no support for culture (P, ¬C)	1
0 : No property rights, no support for culture (¬P, ¬C)	0

The incremental version of the principle applied at the aggregated level requires that **A** be unan-

imously acceptable as against **B**, and that **B** be unanimously acceptable with respect to **0**, i.e. it must be the case that $A >_u B >_u 0$ (where "$>_u$" means "is unanimously preferred to"). Disaggregation requires that each of the two component policies P and C be unanimously acceptable as compared to the absence of that policy. Thus, state enforcement of a system of property rights would have to be unanimously acceptable as against not having a system of property rights, and state support for culture would have to be unanimously acceptable as against no such support. However, to require that P be unanimously acceptable with respect to ¬P is to require that **A** and **B** be unanimously acceptable with respect to **0**. Similarly, to require that C be unanimously acceptable with respect to ¬C is to require that **A** be unanimously acceptable with respect to **B** and **0**. Putting these conditions together yields the incremental criterion in the aggregated space, i.e. **A** must be unanimously acceptable with respect to **B**, and both must be unanimously acceptable with respect to **0**.

The situation changes only slightly if we move to a more general case, in which all four options are feasible, and where coercion is proportional rather than strictly equal to the number of component policies. We can label the options via a simple table:

	Y	¬Y
X	A	B
¬X	C	0

If coercion is only proportional rather than equal to the number of component policies, we know that $A >_c \{B, C\} >_c 0$, but we don't know whether **B** is more or less coercive that **C**. In this case, **B** and **C** have to be unanimously acceptable against **0**, while **A** has to be unanimously acceptable as against **B**, **C**, and **0**, i.e. $A >_u B, C >_u 0$. But this is what disaggregation implies, so long as it

is feasible:

1. If X $>_u$ ¬X, then **A** $>_u$ **C**, and **B** $>_u$ **0**

2. If Y $>_u$ ¬Y, then **A** $>_u$ **B**, and **C** $>_u$ **0**,

3. 1 and 2 imply **A** $>_u$ **B**, **C** $>_u$ **0**

So, to disaggregate policies is equivalent to applying the incremental version of the principle in the aggregated space, so long as coercion is measured by the number of component policies, and is feasible. If coercion is only proportional to the number of component policies, such that **B** could be more coercive that **C**, for example, then the incremental-aggregated version would require qualified acceptability for **B** with respect to **C**, whereas the non-incremental version would only require that both be acceptable with respect to **0**. Nonetheless, so long as coercion is proportional to the number of component policies, the incremental rule and the disaggregation rule accomplish much the same thing.

This equivalence is subject to one important condition, however, which is feasibility of disaggregation. To see why, consider the ranking **A0CB**. The person with this ranking favours X and Y together, but not one without the other, because (in this person's view) the benefits of the micro-level policies accrue only when both are in place. If this ranking is reasonable, the qualified unanimity requirement cannot be applied at the disaggregated level, because those with this point of view can't express an unambiguous or unconditional preference at this level. I turn to the aggregation problem in the next section.

4.   Maximal Feasible Disaggregation

The upshot of the previous section is that if we are committed to a non-incremental version of the principle, we will have to adopt a rule of *maximal* disaggregation. So far, I have as-

sumed that we can disaggregate or not, as we please. In some cases, however, disaggregation

may not be feasible. Why does the bundling of welfare policy with the protection of property

rights does not seem as arbitrary as does the bundling of property rights with the protection of

culture? One answer would be that whether or not one is in favour of property rights may reas-

onably depend on whether policies are in place to help the least well off and to redistribute

wealth in the direction of equality. After all, property rights bar those without property from

having access to the world's resources. Perhaps such rights are justified only if accompanied by

a system of welfare ensuring that everyone enjoys a basic minimum, or inheritance taxes, or

some other such redistributive policies. Of course, reasonable people will disagree about which

issues are connected with which. Gaus proposes that we should consider policies independent

only when all reasonable parties take them to be independent.[22] At first glance, it may not be ob-

vious why we should favour more holistic views in this way. However, there is an important dif-

ference between those who see issues as being disconnected and those who see them as being

connected. The person with the more holistic or 'joined-up' preference cannot express a prefer-

ence about issues separately, whereas a person with a modular preference can express a prefer-

ence about issues in the aggregated space. If the public justifiability criterion requires unanimity

of reasonable preferences and my preference is reasonable, I have to be able to give an answer.

If I can reasonably think that without redistributive policies it would be better not to have a sys-

---

22. For two issues A and B, with associated sets of proposals $\{\alpha_1...\alpha_n\}$ and $\{\beta_1... \beta_n\}$, Gaus defines independence as follows: "Issue A is independent from B if and only if no member of the public P has evaluative standards such that her identification of the eligible elements of $\{\alpha_1...\alpha_n\}$ depends on her decision of the eligible elements of $\{\beta_1... \beta_n\}$. A and B are fully independent if they are independent from each other." Gaus, "On Two Critics of Justificatory Liberalism," 26.. The "public" here is the set of qualified e.g. reasonable citizens. A proposal about an issue is eligible if and only if it is preferred by all members of the public no law at all on the issue. Issues are dependent, according to Gaus's definition, if someone can reasonably think that which proposals are better than nothing for one issue depends which proposal is chosen for another issue.

tem of property rights at all, the issues of redistribution and property are not independent. I can-not rank 'property rights' vs. 'no property rights' in general, because my preference over these options varies, depending on whether there will be any redistribution. If it is reasonable for one's moral assessment of systems of property rights to vary in this way, depending on whether redistributive measures are in place, we cannot apply the qualified unanimity criterion to the disaggregated set of choices. The contrast between redistribution and the promotion of culture is straightforward. No one reasonably thinks "I am in favour of state promotion of high culture, but if there is not going to be any state promotion of culture, then I prefer anarchy (i.e. no property rights)." Reasonable moral preference orderings do not contain this kind of dependence between the issues of property and culture, making the two issues separable. Thus we arrive at a first criterion of feasible disaggregation, or independence.[23] Our rule is disaggregate so as to avoid allowing coercion to slip through the cracks without having to secure unanimous approval of reasonable views, but only to the point at which all reasonable views can still rank the options on each disaggregated issue independently of the others. We can refer to this rule as *maximal feasible disaggregation*, where feasibility is defined by the criterion of *unanimous unambiguous rankability*.

The unanimous rankability criterion is hard to meet, with the result that in many cases we will have to aggregate issues so that we are choosing between packages with many component rules or policies. When feasibility is defined as unanimous rankability, the criterion of maximal feasible disaggregation will be highly permissive of state action (assuming for the moment that

23. To be clear, what we are aggregating here is not preferences but policy issues. The criterion specifies when we have to 'zoom out', so to speak, moving to consider sets of sub-policies due to a lack of independence of these sub-policies. The term "independence" is from Gaus's criterion of "justificatory independency;" Gaus, *The Order of Public Reason*, 495.

we are sticking with the non-incremental version of the principle). To see why this is, it will be

helpful to consider our simple two-policy case, in the version that includes option C:

	Y	¬Y
X	A	B
¬X	C	0

This format permits easy depiction of preference orderings. Here is one such ordering:

1	2
3	4

= A > B > C > 0 =

X, Y
X, ¬Y
¬X, Y
¬X, ¬Y

It is then straightforward to see which preference orderings involve reversals i.e. require aggreg-

ation. Consider first those preferences that rank A first, of which there are six[24]:

ABC0   AB0C   ACB0   AC0B   A0BC   A0CB

1	2
3	4

1	2
4	3

1	3
2	4

1	4
2	3

1	4
3	2

1	3
4	2

In only two cases are there no reversals:

ABC0   ACB0

1 < 2   1 < 3
^   ^   ^   ^
3 < 4   2 < 4

In the other four cases, there are reversals in one or both dimensions:

_____

24. I assume for simplicity that no one is indifferent.

AB0C			AC0B	A0BC	A0CB

$$\begin{array}{ccc} 1 & < & 2 \\ \wedge & & \wedge \\ 4 & > & 3 \end{array} \qquad \begin{array}{ccc} 1 & < & 4 \\ \wedge & & \vee \\ 2 & < & 3 \end{array} \qquad \begin{array}{ccc} 1 & < & 4 \\ \wedge & & \vee \\ 3 & > & 2 \end{array} \qquad \begin{array}{ccc} 1 & < & 3 \\ \wedge & & \vee \\ 4 & > & 2 \end{array}$$

In virtue of symmetry, the situation is the same no matter what option is ranked first; the only

preferences without reversals are the zig-zag patterns:

BA0C   CA0B   0BCA

2	1		2	4		4	2
4	3		1	3		3	1

B0AC   C0AB   0CBA

3	1		3	4		4	3
4	2		1	2		2	1

Thus, out of the 24 possible rankings, only 8 contain no reversals, and do not require aggrega-

tion. If the only reasonable preferences are in this set of 8, we disaggregate. If on the other hand

the set of reasonable preferences includes even one ranking containing a reversal, a ranking such

as **AB0C**.

The criterion of unanimous unambiguous rankability makes the rule of maximal feasible

disaggregation fairly lax, or weak, assuming we are using the non-incremental version of the

principle. So long as anyone reasonably finds themselves unable to rank policies independently,

we must apply the idealized unanimity criterion at the aggregated level. The concern here is that

the most holistic reasonable view will license democratic choice from amongst very broad,

multi-component "package" options, simply because many such options are worse than the default of no common rules at all in any of these packaged policy domains. If the worry about the incremental version of the principle was that it makes the barely reasonable libertarian a dictator, the worry about the maximal feasible disaggregation rule, when defined in terms of unanimous rankability, is that the barely reasonable holist is able to get very unpalatable options on the agenda of democratic choice, as elements of packages of policies that are better than the option of having no policy at all in any of these areas.

I will point out in passing that when applying the rule of maximal feasible disaggregation, it makes a big difference whether we consider a given policy to be impossible, instead of considering it simply a very bad idea. In the property example that gave us the labels P and R, I initially considered only three possibilities:

A = P, R
B = P, ¬R
0 = ¬P, ¬R

but not the fourth option:

C = ¬P, R.

The rationale for omitting **C** was that if there is no system of property rights, there can also be no redistribution. However, perhaps we can construe **C** = ¬P, R as a situation in which there is no system of property rights, but there is nonetheless a state that engages in periodic redistribution of de facto possessions. We might suppose that all reasonable preferences rank **C** last, because the idea of punctually redistributing de facto possessions seems even worse than a state of nature. Surprisingly, whether we count **C** as impossible or just very bad determines the content of the eligible set, because including **C** triggers the aggregation rule.

Consider the following set of preferences over the three possible combinations of disaggregated policies, assuming **C** to be impossible rather than just inadvisable. Which of these preferences is reasonable will determine whether we must aggregate, and thus the contents of the eligible set.

<#???#>

Far Left	Left	Right	Far Right
A : P, R	**A** : P, R	**B** : P, ¬R	**B** : P, ¬R
0 : ¬P, ¬R	**B** : P, ¬R	**A** : P, R	**0** : ¬P, ¬R
B : P, ¬R	**0** : ¬P, ¬R	**0** : ¬P, ¬R	**A** : P, R

According to the principle of maximal feasible disaggregation, the reason for not disaggregating these policies must be that some reasonable views cannot state an unambiguous or stable preference at the disaggregated level (whether to have a system of property rights, and whether to tax for redistribution). Suppose first that either Far Left or Far Right is a reasonable preference. Neither Far Left and Far Right have an unambiguous preference over P, taken alone; Far Left prefers P if R is in place, but ¬P without R, while for Far Right the situation is reversed. If either Far Left or Far Right are counted as reasonable preference orderings, we must aggregate. Under either of these assumptions about reasonable preferences, however, one or both of **A** and **B** are ineligible. If Far Left is a reasonable ordering, **B** does not pass the test of being universally preferred to the default, **0**; if Far Right is a reasonable ordering, **A** does not pass the test. Now suppose that only Left and Right are reasonable. In this case we must disaggregate, by the principle of maximal feasible disaggregation, since both Left and Right can state a preference for P vs. ¬P

and R vs. ¬R separately. Yet disaggregation yields an eligible set of **B**.   Since all qualified

points of view rank the 'no property rights' option last (envisaging a Hobbesian state of nature,

perhaps), then it is feasible to apply the qualified unanimity criterion *first* to the decision to have

a system of property rights, yielding an unambiguously positive response, *then* to the question of

whether or not to have policies that redistribute income in the direction of equality, where the

reasonable controversy over redistribution will force us to default to not redistributing.   The

surprising result is that the rule of maximal feasible disaggregation yields eligible sets that con-

sist of singletons; in no case do we get an eligible set that consists of **A** (P, R) and **B** (P, ¬R).

However, if we consider **C** to be possible but simply a very bad idea, we can get a differ-

ent result. Suppose that only Left and Right count as reasonable, leaving reasonable disagree-

ment only about the ranking of **A** and **B**:

Left	Right
A : P, R	**B** : P, ¬R
B : P, ¬R	**A** : P, R
0 : ¬P, ¬R	**0** : ¬P, ¬R
C : ¬P, R	**C** : ¬P, R

The Left ranking **AB0C** involves a reversal; R is ranked above ¬R if P, but below ¬R if ¬P.

Those with the **AB0C** preference cannot express an unambiguous or unconditional preference

about R, and we have to aggregate. **C** is commonly agreed to be worse than nothing, so the eli-

gible set is **A** and **B**. The non-incremental version of the principle permits democratic choice

between **A** and **B**, in this case. So it is possible to get an eligible set of **A** and **B**, if one keeps **C**

in the space of possibilities, since **C** forces aggregation.[25] But the results of the rule are highly sensitive to the way the set of options are specified, and to the set of preferences that are considered reasonable. What preferences are considered reasonable matters not only to what options are in the eligible set (because no one thinks them worse than no policy at all), but also to the level at which we apply the qualified acceptability rule (depending on whether or not some reasonable preference contains a reversal, at that level of aggregation).

The problem of choosing an appropriate level of aggregation when using the non-incremental version of the principle is closely related to the paradoxes explored by social choice theorists that arise when a group makes a set of logically connected decisions in a series of separate votes by majority rule. Philip Pettit gives the example of a party composed of three members or factions deciding how government will spend and tax.[26] All are agreed that the budget should balance. A wants to increase spending across the board while raising taxes. B wants to increase defence spending while holding taxes constant by reducing spending in other areas. C wants to increase spending in other areas while holding taxes constant by reducing spending on defence. Each holds consistent preferences, but if they consider the issues separately, they may get an outcome no one wants. Only A prefers raising taxes, but A and B want to increase spending on defence, while B and C want to increase spending elsewhere. The party votes in January by majority rule to promise not to increase taxes. In June, it also promises to increase defense spending.

---

25. One way to strengthen the independence criterion would be to apply it iteratively. Start with the aggregated space, eliminate options unanimously viewed as worse than nothing, then apply the disaggregation rule to the winnowed set of possibilities. This iterative application of the independence criterion would bring us back to our initial, three-option situation, in which the eligible set is always at best a singleton.

26. Philip Pettit, *A Theory of Freedom: From the Psychology to the Politics of Agency* (Cambridge: Polity, 2001), 112; Philip Pettit, "The Reality of Group Agents," in *Philosophy of the Social Sciences: Philosophical Theory and Scientific Practice*, ed. Chris Mantzavinos (Cambridge: Cambridge University Press, 2009), 77.

In September, it may also promise to increase spending in other areas, *even though each member remains committed to balancing the budget*. The reason that this can happen is that one majority favours holding taxes constant, a different majority favours increasing defence spending, and a third majority favours increasing other spending, even though no one want to run a deficit. Of course, whether or not this will in fact happen depends on what political calculations the members make, e.g. 'if I help vote down the tax increase, what are my chances of later rescinding the spending increase on defence?' It is sometimes said that such examples show that there is a problem with decision procedures that are responsive to individual preferences.[27] The problem in such cases is not that individuals are being consulted but that they are being asked to express preferences about disaggregated choices in cases in which their preferences have a conditional structure that can only be accurately captured in the aggregated space. Consider what Alf's preference ordering might look like over the joint options (ignoring the distinction between military and other spending, for the sake of simplicity):

1. Raise taxes, increase spending
2. Keep taxes the same, keep spending the same
3. Raise taxes, keep spending the same
4. Keep taxes the same, increase spending

If Alf is asked to state his preference on increased spending, his full answer should be "Well, that depends. If we're going to raise taxes, (1 or 3), then I'm in favour of increasing spending. But if we're not going to raise taxes (2 or 4), then I'm against." If Alf is asked to vote

---

27. "The unreliability of majority voting is revealed by the discursive dilemma (Pettit 2001a, ch. 5), a problem that generalizes the doctrinal paradox in juridical theory (Kornhauser and Sager 1993)... But the problem is not restricted to majority voting. It turns out that making a group responsive to its individual members in the manner that is exemplified by majority voting, but not only by majority voting, rules out an assurance that the group displays collective rationality;" Pettit, "The Reality of Group Agents," 77-78.

on a spending increase by itself, he has to try to predict whether a tax increase will follow. He has to make this prediction because his preferences in the aggregated space involve what I called a reversal. He is being asked to represent his moral preferences about taxation independent of his preferences about spending despite the fact that his preferences have a conditional structure. All he can do is to vote about taxation based on his best estimate of the likely outcome of future decisions about spending. Genuine responsiveness to individual preferences would require that individuals be able to express their preferences in the aggregated space.

The need for a criterion of feasible disaggregation arises for the same reason that inconsistency of collective choice arises, which is that in both cases, individuals are being asked to express preferences about issues separately when their preferences at this level of disaggregation have a conditional structure. There are however three differences between the two scenarios. First, in the standard social choice cases, the voting procedure is majority rule in an actual public, whereas in the context of public justification we are using a unanimity rule in an idealized public. Second, in the social choice cases, individuals can vote on disaggregated but logically connected options, even though their preferences have a conditional structure, because they can vote based on their best guess about the likely evolution of policy on other issues. In the context of public justification, however, the criterion of unanimous reasonable acceptability is not well defined if some reasonable preference orderings contain reversals. Thirdly, in the social choice case the problem the group confronts is not just that some of its members may have to vote on the basis of predictions about votes on future cases. The problem is that if they decide on a piecemeal basis, they may (depending what predictions they each make about future votes) end up with what they all agree is the worst possible overall result in the aggregated space, e.g. the

deficit option may be chosen despite being Pareto inferior to all other bundled options.

Although unanimous rankability is a plausible criterion of feasible disaggregation, there is a fatal problem with the idea of *maximal* feasible disaggregation, which is that sometimes maximal feasible disaggregation will mean disaggregating too much, to a level at which it no longer makes sense to demand that one option be unanimously acceptable with respect to the other. The criterion of maximal feasible disaggregation will give the wrong results if the coerciveness of aggregated policies is not proportional to the number of component policies.

Return once more to the case of property rights, supposing that the only reasonable preferences are as follows:

Left	Right
A : P, R	**B** : P, ¬R
B : P, ¬R	**A** : P, R
0 : ¬P, ¬R	**0** : ¬P, ¬R

In this case, disaggregation requires that we apply the qualified acceptability test first to P vs. ¬P, and then to R vs. ¬R. However, requiring that R be unanimously acceptable as against ¬R is equivalent to requiring that **A** be unanimously acceptable with respect to **B** and **0**. If **A** and **B** are equally coercive, however, as our egalitarian earlier asserted[28], this requirement is arbitrary. It doesn't make sense to require that one option be conclusively preferred to another equally coercive option. In this case, we are disaggregating too much. The rule of maximal feasible disaggregation only makes sense to the extent that it mimics the incremental version of the principle.

---

28. See page 11

If **A** and **B** are equally coercive, or at least roughly so, it makes sense that they should both be in the eligible set. But if it is true that **A** is much more coercive than **B**, the logic of the property and culture example applies; why should the reasonably-rejectable coercion of R be allowed to piggyback on the non-reasonably-rejectable coercion of P? Thus, where the non-incremental disaggregated version of the principle diverges from the incremental aggregated version, it is the incremental aggregated version that makes more sense.

5. Reasonable Disagreement about the Eligible Set

The argument from higher-order unanimity was supposed to show that one can frame the principle of public justifiability as a constraint on coercion directly, so as to rule out perfectionism, while not simultaneously ruling out redistribution. In turns out that it is possible to use the argument from higher-order unanimity in this way, but that to do so one must make quite specific assumptions about the set of reasonable moral preferences and about the measurement of coercion. In order to set the level of aggregation at which to apply the qualified unanimity standard (i.e. how much to bundle policies), one must know at what level reasonable persons can provide unambiguous preference orderings (or whether disaggregated decision-making will yield sub-optimal results at the aggregate level). If one is using the incremental version of the principle (as I think one must if one accepts the framing of public justification as a direct constraint on coercion), then one needs to know how the various options rank in terms of coercion. There is room for reasonable disagreement about each of these issues. As a result, different people will consider different sets of policies to be eligible.

Take, for instance, the issue of how to measure coercion. Considered in isolation, having

a law forbidding action X is more coercive that not having such a law.[29] It does not follow, however, that the state necessarily engages in more coercion when it has a law against Y in addition to its law against X, than if it has a law against X alone. Laws can interact in ways that affect the total amount of state coercion, such that:

$$\text{Coercion}_{\text{TOTAL}} \neq \text{Coercion}_{\text{X}} + \text{Coercion}_{\text{Y}}$$

The example I gave in earlier papers concerned laws against assault. Having a law against assault is more coercive than not having such a law, but once we have a law against assault, adding a further law against domestic assault means that the state will start punishing domestic battery as assault, but stop punishing those that come to the defence of the battered spouse.[30] David Estlund pointed out in response that this example depends on it being the case that once domestic assault is illegal, private action to prevent it becomes legal. This assumption will not generally hold; the fact that it is illegal for you to water your lawn does not necessarily give me the right to turn off your sprinkler.[31] However, since my claim is only that the passage of one law *can* affect how coercive another law is (so that the total amount of coercion the state engages in *need not* be proportional to the number of component laws), the fact that this will not always be so is no problem. More significant is Estlund's further point that once we have distinguished the question of whether an action X is forbidden from the question of whether private actors have the authority to act to prevent others from X-ing, there is no ambiguity about what counts as more or less

---

29. That is to say, it involves the state in more coercion; whether or not the law reduces the level of coercion on the part of others, and so minimizes the overall level of coercion is neither here nor there; David Estlund, "Replies to Saunders, Lister and Quong," *Representation* 46, no. 1 (2010), 60.

30. Andrew Lister, "Liberal Foundations of Democratic Authority," *Representation* 46, no. 1 (2010), 27; Andrew Lister, "Public Justification and the Limits of State Action," *Politics, Philosophy, Economics* 9, no. 3 (2010), 160.

31. Estlund, "Replies to Saunders, Lister and Quong," 59.

state coercion. Other things equal, a law prohibiting watering one's lawn is more coercive than

the absence of such a law, and other things equal, a law prohibiting neighbours from turning off

each other's sprinklers is more coercive than the absence of such a law, Estlund suggests. This is

the claim I need to deny. I do so on the basis of the fact that to measure the extent to which a

state is engaged in coercion, one can't simply count the number of actions forbidden. One also

needs to have some sense of the significance these possibilities of action have. We cannot

simply count the number of actions forbidden in order to assess the level of coercion associated

with a particular law, for the same reasons that we cannot arrange our basic liberties so as to

maximize liberty *tout court*. A well known example: we should not say that there is less coer-

cion in a society with no traffic lights but no freedom of conscience than in a society with free-

dom of conscience and heavily regulated traffic, despite the greater number of discrete acts of in-

terference in the latter case.[32] If the coerciveness of a law against X is not simply a matter of

counting the number of cases in which someone would have liked to X and is not permitted to do

so, but involves qualitative judgments about the importance of the opportunities denied or made

available, then it may be the case that how important the liberty to X is depends on whether one

is also free to Y. For the same reason, how coercive it is to be denied the liberty to X may

depend on whether one is also denied the liberty to Y. One might think that a law against water-

ing the lawn that is not accompanied by a law forbidding private enforcement of the lawn-water-

ing rule is *more* coercive than a law against watering together with a law against private enforce-

ment. Having a law against watering but no law against private enforcement not only forces me

---

32. Charles Taylor, "What's Wrong With Negative Liberty," in *Philosophy and the Human Sciences: Philosophical Papers 2* (Cambridge: Cambridge University Press, 1985), 219; c.f. H.L.A. Hart, "Rawls on Liberty and Its Priority," *The University of Chicago Law Review* 40, no. 3 (1973): 534-55; and Ronald Dworkin, "What Rights Do We Have?,", Taking Rights Seriously (London: Duckworth, 1978).

not to water my lawn, it makes be vulnerable to what would otherwise be the trespass of my neighbour the environmental fanatic. At least if both laws are in place I know that I will only be fined by the city, if I am deemed to have broken the law, and that I will have a chance to dispute fine in question. In this situation, if we add a further law against private enforcement, it is not clear that the total level of state coercion goes up. Yes, we are now not permitted to turn off the neighbour's sprinkler. But the law against watering one's lawn is also less coercive, when private enforcement is banned. To be sure, when dealing with unrelated laws, coercion will be additive. Having a law vs. smoking and a law vs. jaywalking is more coercive than either alone. When laws are logically connected, however, this will not necessarily be the case. How coercive it is to enforce property rights that force people not to use or appropriate resources that are privately owned may depend on whether the state also forces people to contribute to welfare programs. I remain persuaded, therefore, that it is not generally true that the state engages in more coercion or is more coercive when it has two laws than when it has one.

The point I want to draw from this discussion is that the measurement of coercion is going to be controversial, as will the delimitation of the set of reasonable preference orderings. Such controversy will affect people's judgment about what the correct default is, deviation from which should require unanimous (idealized) acceptability, as the dispute between the libertarian and the egalitarian suggested. In such cases, it may be that no matter which option we choose, some reasonable persons will have objections to what they reasonably take to be the coercive exercise of political power. Reasonable disagreement about how to apply the public justifiability criterion cannot itself block state action, because what counts as state action is what is in ques-

tion.[33]   It can undermine the authority of the state, however, if we think that being a member of

the optimal eligible set is a criterion for a proposed law to be authoritative, rather than merely be-

ing a condition for the law to be fully just / justified. [34]

The fact of reasonable disagreement about the interpretation of a principle has different

implications for the question of authority than it does for the question of justice or full moral jus-

tification. If we ask what the right law, institution or policy is, we may look to a theory of justice

to provide the answer. For example, we may say that whatever else they may do, laws should

treat people as equals. If it is objected that there will be disagreement about the specification of

equal respect, and hence about its implications, our response is likely to be "yes, of course, but so

what? Each citizen should try to figure out the right conception of equality, and then make his or

her own decisions about what laws and policies to support based on the view he or she arrives

at." The situation is different if we ask what laws are authoritative, in the sense of meriting

obedience simply because they are the law. If to this question we answer "laws or institutions

are authoritative on condition that they people as equals," disagreement about the interpretation

or application of the standard of equality will pose a serious problem. Disagreement about

---

33. If I believe that I am entitled to vote for the egalitarian system of property rights on grounds of the argument
from higher-order unanimity, I am simply supporting what I reasonably take to be the best of the options that are
conclusively justified *as compared to what I take to be the appropriate default*. You may think that I am mistaken
about the measurement of coercion, that I have the wrong default, and that I should therefore be supporting **B**: P, ¬R
as the not-conclusively-defeated baseline. Yet our situations are symmetrical; I too have a reasonable objection to
your judgment about what the non-coercive default is. Without any reason to prefer your reasonable view over
mine, I have no choice, when it comes to deciding what kind of laws to support, but to apply the public justifiability
criterion relative to my own resolution of the aggregation/measurement issue. For this reason, reasonable
disagreement about the non-coercive default may not threaten the legitimacy of state action.

34.   By 'legitimacy', I mean the right of the state to enforce its rules, and derivatively the citizen's right vote for and
support the enforcement of such laws. By 'authority', I mean the state's right to make rules that impose obligations
on citizens. Ordinarily the two go together, but they can come apart. This can be the case in Hobbes's political
theory, for example, where the sovereign can without injustice (i.e. without breaking any agreements) do whatever it
wants, but subjects never have a duty to kill themselves or incriminate themselves.

specification and implementation of our abstract principle will mean that many laws will be judged by many citizens not to treat citizens as equals. Of course there can be good reasons for complying with laws that lack authority, based on one's assessment of the moral consequences of not complying. Without authority, however, non-compliance is more likely, and if our assurance about the compliance of others is undermined, we may not be willing to comply ourselves. Setting our standard for authoritative law too high (e.g. fully respecting people as equals), threatens to leave us without any effective law, simply because the existing law is not ideal, even if the absence of law is worse, morally speaking, than the existing law. Standards of authority must by their very nature be less subject to disagreement than standards of justice or full moral justification.[35]

The need for standards of authority to command broader agreement than do conceptions of justice means that the principle of public justification cannot be a standard of authority, but only an ideal of full justification. For the coercion variant of the principle, the crucial question is what it means for a law to fall outside (what I take to be) the optimal eligible set, due to being ineligible (worse than no law at all). Am I to conclude that such laws are merely not ideal, because not publicly justifiable, and hence not fully just or fully morally proper? Or, are ineligible laws lacking in authority? Gaus sometimes talks of public justification as a matter of justice or full

---

35. In this paragraph, I am spelling out what I take to be Waldron's reasoning, as for example in the following paragraph: "[A]n answer to the question of authority must really settle the issue. It is no good saying, for example, that when people disagree about rights, the person who should prevail is the person who offers the best conception of rights. Each person regards her own view as better than any of the others; so this rule for settling on a social choice in the face of a disagreement is going to reproduce exactly the disagreement that called for the rule in the first place. The theory of authority must identify some view as the one to prevail, on criteria other than those which are the source of the original disagreement;" Jeremy Waldron, "A Rights-Based Critique of Consitutional Rights," *Oxford Journal of Legal Studies* 13, no. 1 (1993), 32. See also Jeremy Waldron, *Law and Disagreement* (Clarendon Press, 1999), 1-17.Waldron, *Law and Disagreement*, 1-17.

moral rightness.[36] However, his main position is that failures of public justifiability rob laws of their authority.[37] David Estlund may be of the same view.[38] If the exercise of political power must be publicly justifiable i.e. reasonably acceptable, then it would seem that laws and policies could have authority over me only when I would be unreasonable to reject this authority. Presumably, if the question is whether I am obligated to obey, it should be my judgment about the non-coercive baseline that matters. You may be entitled to coerce me because you reasonably believe that the law in question is conclusively better than what you reasonably take to be the non-coercive default, but if I reasonably believe that the law in question is *more* coercive than what *I* reasonably take to be the non-coercive default, and I reasonably deny that the law is justified, the law has not be justified *to* me, and I am not obligated to obey.

As we have seen, however, attempts to apply the principle that coercive state action must be publicly justifiable are likely to lead to a wide range of divergent conclusions. Disagreement about which preferences are reasonable and about how to measure coercion are likely to mean

---

36.   He states, for example, that according to the "Political Liberty Principle," use of state force or coercion without adequate justification is an "injustice;" Gaus, "Coercion, Ownership, and the Redistributive State," 239. He also defines the "Public Justification Principle" in terms of what makes a "justified" coercive law, remarking that "An unjustified law fails to treat each person as free and equal;" Gaus, "Coercion, Ownership, and the Redistributive State," 244.

37. "Laws that fall outside the eligible set are laws that some member of the public have no reason at all to accept, according to their own evaluative standards;" Gaus, "Coercion, Ownership, and the Redistributive State," 249. "The aim is to justify an authoritative law that imposes a moral obligation to obey;" Gaus, "On Two Critics of Justificatory Liberalism," 196. A person will find an existing or proposed law L unacceptable, Gaus says, "if L is worse than its absence – a 'no-duty' to obey L, and 'no-coercion' based on L;" Gaus, "On Two Critics of Justificatory Liberalism," 196. In *The Order of Public Reason*, Gaus builds up to questions about law and policy based on reflection about the justification of the authority claims that he argues are a common part of our ordinary interpersonal morality. In this work, Gaus is explicit that being ineligible robs a moral rule of its authority. Proposals that are ineligible on Alf's ranking "are ineligible because he cannot accept claims to moral authority based on such rules. That is, he does not have reason to endorse this rule over no moral authority at all;" Gaus, *The Order of Public Reason*, 322.

38. "It is meant to be a necessary condition of permissible coercion or authority that there be a justification for such action that meets the requirement [the qualified acceptability requirement];"Estlund, "Replies to Saunders, Lister and Quong," 59.

that reasonable people will disagree about which option or package of options ought to constitute the non-coercive default, and (hence) about which options ought to be considered eligible. Understood in this way, as a criterion for state action to have authority, I think that the principle of public justifiability fails the disagreement criterion, i.e. the rule that principles of authority must not be subject to too much disagreement in application. Even if when correctly applied the coercion version of the principle of public justification is not libertarian (i.e. when applied according to one a particular specification, and a particular account of reasonable preferences and the degrees of coercion), the implications of its universal adoption would be philosophically anarchist. Public justifiability cannot be a criterion for rules to be authoritative. It can only be an ideal meant to guide decision-making.

6.   Conclusion

It is tempting to dismiss the demand for public justification when it is framed as a constraint on coercive state action directly, on the grounds that the principle has intuitively unacceptable implications for policy. The argument from higher-order unanimity or nested inconclusiveness can make this version of the principle less radically libertarian than it initially appears. Those committed to public justification can frame the principle as a constraint on coercive state action directly so as to avoid perfectionism, but appeal to higher order unanimity in order to permit redistribution in the interest of egalitarian social justice. Such an appeal *can* work, but the argument is very sensitive to how one specifies the set of reasonable preferences and how one measures the coerciveness of policies. The conclusion I draw is that the principle of public justification cannot function as a principle of authority, but only as an ideal of full justification, as is true for the reasons-for-decisions / consensus framing of the principle.