

Note: this is an article whose final and definitive version appears in *Politics, Philosophy, and Economics*; the official version of the paper is available via: <http://journals.sagepub.com/doi/abs/10.1177/1470594X17741228>.

Regulating offense, nurturing offense

Robert Mark Simpson

Abstract: Joel Feinberg's *Offense to Others* is the most comprehensive contemporary work on the significance of offense in a liberal legal system. Feinberg argues that being offended can impair a person's liberty, much like a nuisance, and that it is therefore legitimate in principle to regulate conduct because of its offensiveness. In this paper I discuss some overlooked considerations that give us reason to resist Feinberg's conclusion, even while granting this premise. My key claim is that the regulation of offense can inadvertently increase the incidence of offense, by nurturing offense-taking sensibilities. In the course of defending this claim and spelling out its implications, I explain why concerns about the inadvertent nurturing of offense are now more pressing, given the identity-political character of contemporary offense-based social conflicts, and I discuss why a reluctance to legally regulate offensive conduct need not be insensitive to the identity-political issues that animate those conflicts.

1. Introduction

The social politics of offense in Western liberal societies has transformed. Many of the major offense-based social conflicts of the 20th century were a result of gadflies, iconoclasts, and provocateurs – the James Joyces, D. H. Lawrences, and Monty Pythons of the world – self-consciously defying the mores of the religious and political establishment. This sort of thing still goes on, of course. But however much blasphemy and satire rile the evangelical fringe, Christian institutions overall have far less influence today

on Western societies' values and standards. The people crying foul, nowadays, in major episodes of offense-driven social conflict, generally aren't offended on behalf of ideals that are representative of the cultural mainstream. What we see instead, under this new social politics of offense, is a pattern of conflict in which the offended parties belong to minority religious groups, whose ideals and objects of reverence are of no special concern to the mainstream. The *Satanic Verses* controversy in the late 1980s, the Danish cartoons affair in 2005, and the Charlie Hebdo affair in 2015 are the best known examples of this. Each of these conflicts was instigated by expression that violated standards held by (some) members of the affected religious communities, but which was generally viewed as fair play within the political mainstream. The new social politics of offense is defined by the patterns of discord and polarization that occurs in these cases, in which the derision of minority worldviews is supported, or at least unopposed, by the majority.¹

In popular political discussion around these kinds of conflicts it is often assumed that liberal ideals are entirely on the side of the offenders, and not the offendees. The notion that people who cause offense to others must be protected against any government interference is invoked as if it were an uncontroversial axiom of liberalism.² But this is mistaken. The most important contemporary liberal work on the legal

¹ For philosophically-oriented discussion of (i) the *Satanic Verses* controversy, see e.g. Bhikhu Parekh, "The Rushdie affair: research agenda for political philosophy", *Political Studies* 38 (1990): 695-709; Jeremy Waldron, "Rushdie and religion" in *Liberal Rights: Collected Papers 1981-1991* (Cambridge: Cambridge University Press, 1993); (ii) the Danish Cartoons controversy, see e.g. Ronald Dworkin, "The right to ridicule", *The New York Review of Books*, 23rd March 2006; Philip Cook and Conrad Heilmann, "Two types of self-censorship: public and private", *Political Studies* 61 (2013): 178-96; and (iii) the Charlie Hebdo affair, see William A. Edmundson, "Charlie Hebdo meets utility monster", forthcoming in *The Critique*. Obviously all three cases involve offense to Muslims specifically, and I'll comment on the particular 'clash of civilizations' narrative that has been constructed around that pattern of conflict in §6. It is worth acknowledging right from the outset, though, that other religious minority groups are sometimes embroiled in this pattern of conflict. One such conflict in the UK involving the Sikh community is discussed by Monica Mookherjee, "Permitting dishonour: culture, gender, and freedom of expression", *Res Publica* 13 (2007): 29-52. One could also arguably characterize evangelical Christians in the UK in the 21st century as a minority religious group, in which case the controversy over *Jerry Springer: The Opera* in the mid-2000s would fit the pattern of conflict that I've sketched here as well. For discussion of this case see Russell Sandberg and Norman Doe, "The strange death of blasphemy", *The Modern Law Review* 71 (2008): 981-83; Stewart Lee, *How I Escaped My Certain Fate: The Life and Many Deaths of a Stand-up Comedian* (London: Faber and Faber, 2010).

² Many op-ed pieces along these lines were published in the wake of the Charlie Hebdo affair, e.g. Karl Sharro, "Charlie Hebdo and the right to be offended", *The Atlantic*, 10th January 2015; Dan Arel, "Charlie Hebdo has a right to offend you and it is not a double standard", *Huffpost Blog*, 18th February 2015; Nicole Hemmer, "In defense of the freedom to offend", *U.S. News*, 5th May 2015. The prevalence of the idea that liberalism is axiomatically opposed to any offense-based regulation is partly due to advocacy by civil libertarians, especially the type practiced by the ACLU in the American culture wars of the 1990s, see e.g. Nadine

significance of offense – Joel Feinberg’s *Offense to Others* – is in fact a measured defence of the permissibility of anti-offense regulations, and many legal theorists addressing this issue, post-Feinberg, have agreed with him that it is legitimate, in principle, for the liberal state to impose anti-offense regulations.³ Now, given that Feinberg’s account explicitly assigns special privileges to free speech, proponents of his account won’t be arguing for the censorship of Charlie Hebdo and the like. But they will be eager to correct the assumption that underlies popular debate in this arena, namely, that liberalism is always, or should always be, entirely tolerant of conduct that offends people.

In this paper I intend to position Feinberg’s account where it belongs, at the center of discussions about the legal significance of offense in a liberal society, but then make trouble for it. I will argue that there are good reasons – reasons that are made more salient by our new social politics of offense – to reject Feinberg’s view about the legal significance of offense. The crucial premise for Feinberg is that being offended can impair people’s liberty. He’s probably right about this. My contention here, however, is that it’s probably also the case that the legal regulation of offense helps to nurture offense-taking sensibilities and hence tends to increase the overall incidence of offense, thus contributing to the liberty-impairments that Feinberg aims to prevent. After discussing definitions and further background details of Feinberg’s account in §2, I’ll defend this key contention in §3, and then explore its broader implications in §4 and §5. I’ll finish in §6 by explaining why a reluctance to regulate offense needn’t be insensitive to the identity-political issues behind today’s offense-driven social conflicts.⁴

Strossen, “Regulating racist speech on campus: a modest proposal?”, *Duke Law Journal* 1990 (1990): 484-573. The clearest statement in support of this way of thinking that I know of by a legal philosopher is from Ronald Dworkin: “the *essence* of negative liberty is *freedom to offend*, and that applies to the tawdry as well as the heroic”; see “Liberty and pornography” in Susan Dwyer (Ed.), *The Problem of Pornography* (Belmont: Wadsworth, 1995): 117 (my emphasis).

³ Joel Feinberg, *Offense to Others: The Moral Limits of the Criminal Law Volume Two* (Oxford: Oxford University Press, 1985); for followers of Feinberg, see §2, note 14. The term ‘anti-offense regulations’ can read as encompassing both outright criminal prohibitions on offensive acts, and legal restrictions that police the time and place in which offensive acts are performed. Although the latter are still controversial, they are easier to defend, compared to the former, and my criticisms against Feinberg won’t be based on concerns that would arise in attempting to use his arguments to defend the former – more controversial – type of anti-offense regulations.

⁴ In Timothy Garton Ash’s work on free speech in a global age, he indicates a concern with the cultural conditioning of a disposition towards taking offense that partly dovetails with my concerns. “Do we want to be the kind of human beings”, he asks, “who are habitually at the ready to take offence, and our children to be educated and socialised in that way?” See *Free Speech: Ten Principles for a Connected World* (New Haven: Yale University Press, 2016): 91. Also, although they don’t discuss the particular mecha-

2. Feinberg's case for the regulation of offense

The first thing to clarify is our definition of 'offense'. In the legal and philosophical literature on this topic it's customary to construe the term broadly, so that it encompasses a wide range of sub-harmful, adverse mental states. Offense thus includes disgust and sensory aversion (e.g. as in viewing a gruesome image), shock to one's sensibilities (e.g. as in a reaction to very rude behavior), certain forms of embarrassment (e.g. as in unwantedly witnessing nudity or sexual conduct), and affront to a person's values or beliefs (e.g. as when a devoutly religious person is upset by something that he regards as sacrilegious). Granted, we have some reason to be wary about treating all these states as species of the same genus, 'offense'. Given how much variety we see in the types of conduct that trigger these different states, there is room for doubt about whether it is appropriate for the law's stance on them to be governed by one overarching principle. But there is a shared feature here, insofar as these are all states in which the agent's attention is 'captured'. If the agent is forced to be in the vicinity of the triggering source, she finds she can't resist having her attention occupied by the cause of her adverse reaction. This is Feinberg's main reason for thinking disgust, sensory aversion, embarrassment, shock, and various forms of affront can all be bundled together under the banner of offense.⁵

There is a larger question here about whether it is possible to provide any principled demarcation of the class of mental states that are genuinely adverse but at the same time *sub-harmful*. Feinberg distinguishes offense and harm because he conceives of harms as wrongful setbacks to interest, and because he doesn't think that people have an interest as such in not being offended.⁶ This seems like a peculiar stance, on its face, and understanding why Feinberg adopts it would require us to unpack his theory of rights and interests in detail. But this probably isn't necessary for our purposes here. In order to make

nism for nurturing offense-taking sensibilities that I'll be putting forward here, a number of authors in popular political discussion have recently lamented what they see as a trend towards a culture of hyper-sensitivity and over-the-top offense-taking. See for instance Michelle Goldberg, "#CancelColbert and the return of the anti-liberal left", *The Nation*, 2nd April 2014; Jonathan Chait, "Not a very PC thing to say: how the language police are perverting liberalism", *New York Magazine*, 26th January 2015; Greg Lukianoff and Jonathan Haidt, "The coddling of the American mind", *The Atlantic*, September 2015.

⁵ Feinberg, *Offense to Others*: 10-13; *Harm to Others: The Moral Limits of the Criminal Law Volume One* (Oxford: Oxford University Press, 1985): 43-51. Note that some philosophers who are interested in the normative significance of disgust in particular have considered how we might develop more fine-grained distinctions between different kinds of adverse but sub-harmful psychological states; see for example Martha Nussbaum, *Hiding from Humanity: Disgust, Shame, and the Law* (Princeton: Princeton University Press, 2006); Daniel Kelly, *Yuck! The Nature and Moral Significance of Disgust* (Cambridge Massachusetts: MIT Press, 2011).

⁶ See Feinberg, *Harm to Others*: 38ff.

sense of treating offense as a meaningful normative category, and one that is distinct from the category of mentally- or psychologically-mediated harm, all we really need to say is this: there is a class of negative mental states – plausibly encompassing all of Feinberg’s sub-categories: disgust, sensory aversion, embarrassment, shock, and affront – with a common kind of adverse character which is in some sense distinctive. One distinctive characteristic, seemingly shared by all the negative mental states on Feinberg’s radar, is that they are necessarily experienced *as* adverse or unwelcome. It is possible in principle for someone to be harmed, including psychologically harmed, without consciously feeling dislike or aversion. But this isn’t the case with the kinds of negative mental states that Feinberg bundles together under the label of ‘offense’. Being offended necessarily means *feeling* offended. As Douglas Husak puts it, “we can be harmed without our knowledge, but we cannot be offended without having the appropriate mental experience”.⁷

The main theoretical question for Feinberg, and for us, is whether it is among the law’s legitimate purposes, in a liberal society, to try to protect people against these negative mental states. Liberals generally adhere to the Millian precept that the law’s *primary* legitimate purpose is to prevent acts that harm others, as opposed to self-harming acts or harmless wrongdoing. The question is whether this harm principle should be supplemented by an ‘offense principle’. Liberals can go two ways on this issue, broadly speaking. Some will say: “liberalism is committed to pluralism and diversity, and feelings of offense are a natural result of the interpersonal friction that diversity generates, and so it cannot be the law’s legitimate purpose to protect people from such feelings”. Others will say: “it *is* among the law’s legitimate purposes to protect people from offense, because being offended can itself impair the offended person’s liberty”. I will call those who favor the first line of thought *Hard-nosed Liberals*, and those who favor the second line, including Feinberg, *Balancers*. Balancers can agree with Hard-noses in thinking that offense is an inevitable counterpart to liberal pluralism. Where they will differ is in claiming that offense cuts both ways. Even if we cannot eliminate offense in a free society, this doesn’t entail that the effects of offense on the exercise of freedom are benign.

One might wonder just how *liberal* the Balancers are. Following Joseph Raz, we can classify a legal system or policy framework as a liberal one if it assigns normative priority to individual liberty, over the

⁷ Douglas Husak, “Disgust: metaphysical and empirical speculations” in Andrew von Hirsch and A. P. Simester (Eds.), *Incivilities: Regulating Offensive Behavior* (Oxford: Hart Publishing, 2006): 91-114, 94. A similar point is made by David Shoemaker, “‘Dirty words’ and the offense principle”, *Law and Philosophy* 19 (2000): 545-84, 550.

welfare or interests of collectives, and if its recommendation align with the ideals, policies, and institutions (e.g. the separation of powers, democratic government, free speech, secular pluralism) that have historically been markers of the liberal tradition.⁸ If we apply Raz's diagnostic criteria from a certain angle, Balancers might not even qualify as liberals. The challenge Feinberg takes up in *Offense to Others*, then, is to demonstrate that there is a credible rationale in support of the Balancer's position that can be located squarely inside the liberal tradition after all. Liberal societies routinely regulate nuisances. Nuisances are *harmless*, by and large. But as anyone with noisy neighbors knows, nuisances can still significantly impair the exercise of one's liberty. Regulating offense is the law's proper business, then, Feinberg argues, because being offended is relevantly similar to experiencing a sensory nuisance, like a rancid odor or an incessant, piercing noise. The experience doesn't effect a setback to any of the victim's interests *per se*, but it does inhibit her capacity to go about her personal business, and it thereby impairs her liberty. Acts that cause offense, for Feinberg, are

Objectionable for roughly the same kind of reason as the evils combated by nuisance law. Even when they are not harms, they are annoying distractions, unwelcome demands on one's attention, a bother that must be coped with however inconvenient it may be at the time to do so. They are, in short, themselves nuisances in a perfectly ordinary sense.⁹

In order to elicit intuitive support for this, Feinberg tells a series of stories about offensive but apparently harmless acts being performed on a public bus. Reflection on these stories helps us to see that some acts of this sort are intolerable enough to warrant the state's intervention. He invites the reader to imagine that while she is travelling to an important appointment, nearby passengers drool and belch, eat vomit and faeces, perform sex acts with an animal, or mutilate a corpse.¹⁰ The important thing about such acts "from the legislative point of view", Feinberg says, isn't that they're gross or aberrant, but that "they are nuisances, making it difficult for one to enjoy one's work or leisure in a locality which one cannot reasonably be expected to leave".¹¹ This sort of thing "commandeers one's attention from the outside, forcing one to relinquish control of one's inner state, and drop what one was doing in order to cope".¹²

⁸ See Joseph Raz, *The Morality of Freedom* (Oxford: Oxford University Press, 1986): 1-3.

⁹ Feinberg, *Offense to Others*: 9.

¹⁰ *Ibid.* 10-13.

¹¹ *Ibid.* 22.

¹² *Ibid.* 22.

What follows from this, for Feinberg, is that there is a *pro tanto* reason to regulate such acts.¹³ The ‘*pro tanto*’ qualification needs to be stressed. Feinberg recognizes that some feelings of offense, just like some nuisances, may be relatively mild, or voluntarily encountered, or easily avoided. And moreover, some acts that cause offense may serve independently important ends. In light of these considerations, Feinberg readily grants that our *pro tanto* reasons for regulating certain instances of offensive conduct will frequently be outweighed by countervailing factors. An act that offends many people, that is maliciously intended, hard to avoid, and furthers no important end, is liable to restriction under his view. An act that serves an important end, while causing only mild offense, to only a few people, by accident, isn’t. The point is that while regulating offense impairs people’s liberty, leaving it unchecked sometimes does as well. The law’s legitimate purpose in regulating offense, then, derives from the liberal’s core ideal: protecting individual liberty.¹⁴ I will relate my discussion to other liberal accounts of offense in §5, but this view of Feinberg’s will be at the center of my analysis in what follows.

3. How could regulating offense affect its incidence?

In this section I’ll develop the main conjecture that underpins my response to Feinberg. As I stated above, I think there are reasons to believe that the legal regulation of offense in a society will increase the incidence of offense in a society. In making my case for this I’ll begin by introducing two pieces of analytical vocabulary that are helpful in talking about how offense originates and how it can be countered via social policy: first, a distinction between ‘raw’ and ‘processed’ offense; and second, a distinction between

¹³ Here is Feinberg’s precise statement of his offense principle: “It is always a good reason in support of a proposed criminal prohibition that it would probably be an effective way of preventing serious offense (as opposed to injury or harm) to persons other than the actor, and that it is probably a necessary means to that end (i.e. there is probably no other means that is equally effective at no greater cost to other values)”; *Offense to Others*: 1.

¹⁴ Other ‘Balancers’ who follow Feinberg include Judith Jarvis Thomson, “Feinberg on harm, offense, and the criminal law: a review essay”, *Philosophy & Public Affairs* 15 (1986): 381-95, and Raphael Cohen-Almagor, “Harm principle, offense principle, and the Skokie affair”, *Political Studies* 41 (1993): 453-70. Some critics of Feinberg support Balancing but say that only *reasonably taken* offense provides a justification for the regulation of otherwise harmless acts, e.g. see Andrew von Hirsch, “The offense principle in criminal law: affront to sensibility or wrongdoing?”, *King’s College Law Journal* 11 (2000): 78-89, and A. P. Simester and Andrew von Hirsch, “Rethinking the offense principle”, *Legal Theory* 8 (2002): 269-95. One recent Balancer appeals to the protection of cultural rights in making a case for the legitimacy of offense-based regulations; see Meital Pinto, “What are offenses to feelings really about? A new regulative principle for the multicultural era”, *Oxford Journal of Legal Studies* 30 (2010): 695-723.

interventions aimed at the ‘upstream’ minimization of offense, and those aimed at ‘downstream’ mitigation or redress.

3.1 Raw v. processed offense

There are all sorts of things we experience that can trigger instantaneous, unreflective feelings of aversion. There are things that repel the senses (e.g. foul odours, gory images), things that infringe norms of politeness (e.g. coarse language), and things that seem indecent, which we are liable to regard as properly confined to private spaces (e.g. naked bodies). The immediate feelings of aversion triggered by these stimuli can be contrasted with a different kind of aversion that is more durable and reflective, where the bearer of the feeling stewes over the trigger event, or replays it in her head, or dwells on the negative mindset that was elicited.

Here are two examples to illustrate the distinction. First, suppose that Ann shares her society’s normal aversion to public nudity, and suppose she’s unexpectedly exposed to the sight of a man’s genitals in public. Ann’s immediate adverse reaction to this will be roughly the same regardless of the context in which it occurs. But what trails in the wake of this immediate adverse reaction will differ dramatically depending on the context, and how Ann interprets the meaning of the episode in light the context. If Ann has been targeted by a ‘flasher’ exposing himself to her in an act of sexual aggression, she is much more likely to be seriously troubled by the episode, and to seek some form of redress, than if the unwelcome exposure resulted from an accidental ‘wardrobe malfunction’ on a sports field. Here is another example. Suppose that Bill feels a typical degree of aversion when he sees images of gore and bodily mutilation, and suppose he is exposed to an image of a person’s face that has been maimed by a gunshot wound. Again, the repercussions of his initial adverse reaction will mostly depend on the context of the experience and how Bill interprets its significance in light of that. If he sees the image in a book that he’s chosen to read, about the horrors of war, he’s less likely to be seriously aggrieved than if the encounter is due to the image popping up on his Facebook feed. In each case we can distinguish the immediate aversion, whose negative valence is independent of any judgement about the episode’s meaning, from a more enduring aversion, whose negative valence does depend, at least to some degree, on such a judgement.

Both kinds of states qualify as ‘offense’ by Feinberg’s lights. They are both adverse, sub-harmful mental states. But they seem different in an important respect, which our terminology should mark. I will

call the first kind of negative mental state *raw* offense, and the second, *processed* offense.¹⁵ This distinction makes it easier to articulate one misgiving about Feinberg's approach to offense. He says our *pro tanto* reason to regulate offensive conduct owes to the nuisance-like impairment to liberty that comes with feeling offended. Feelings of raw offense do indeed seem nuisance-like in a relevant sense. Like sensory nuisance, raw offense seems to just strike me. There is a stimulus that I dislike, and my encounter with it immediately elicits a feeling of aversion. By contrast, processed offense seems unlike sensory nuisance in certain important respects. It doesn't just strike me. It results from my reflection on, and interpretation of, the significance of the triggering stimulus. Because of this, states of processed offense cannot but be expressions, in some form, of the affected person's beliefs or ideals.¹⁶

Now, the way that nuisances impair liberty, for Feinberg, is by exacting a toll on our attention. And the rationale for trying to prevent offense, based on that concern, probably still holds when we're dealing with processed offense. Unwanted negative mental states remain a burdensome distraction even if they are partly a product of the subject's reflections and judgements. However, when offense is processed, rather than raw, the burden is effected through a more complex set of processes, and hence there is a wider range of interventions that could be used to counteract the burden. If we are seeking to prevent raw offense all we can do is try to stop the subject from coming into contact with the triggering stimulus. But in seeking to minimize processed offense, we can either intervene at that juncture *or* try to influence the interpretative processes through which the subject digests her reaction to the stimulus, in order to forestall the transmutation of raw offense *into* processed offense. Granted, this way of describing things tacitly builds in the assumption that offense is always a bad thing, to be minimized wherever and whenever it occurs. Some will insist that offense is a fitting reaction to certain things, and so, regardless of its liberty-

¹⁵ Others have proposed classifications on similar lines. R. A. Duff and S. E. Marshall use the terms 'immediate offense' and 'mediated offense' in a way that roughly maps onto my distinction between raw and processed offense; see "How offensive can you get?" in von Hirsch and Simester (Eds.), *Incivilities*: 57-90, 59ff. See also Judith Jarvis Thomson's discussion of what she calls 'belief-mediated distress' in *The Realm of Rights* (Cambridge Massachusetts: Harvard University Press, 1990): 253ff.

¹⁶ Disgust is one form of offense in relation to which the contentions I'm making here seem to be borne out in empirical research. As Martha Nussbaum says, summarizing relevant work by the psychologist Paul Rozin, although disgust seems to have an innate evolutionary origin focused on a particular set of stimuli, "teaching plays a large role in shaping the form that the innate equipment takes" and "societies have considerable latitude in how they extend disgust-reactions to other objects"; see Nussbaum, *Hiding from Humanity*: 94, 97. My suggestion, then, is that other forms of offense are relevantly similar to disgust, not because there's a common class of stimuli that trigger raw feelings of affront, annoyance, embarrassment, etc., but because the manner in which those feelings feed into people's judgements is significantly influenced by culturally-mediated attitudes and expectations.

impairing potential, we shouldn't always try to prevent it. I will say more about this line of reply to Feinberg's account in §5. My point for now is this: even if we grant for the sake of argument that Feinberg's nuisance-based rationale for seeking to prevent and deter offense extends to all types of offense, the mechanisms through which offense can be prevented are likely to be different for raw and processed offense.

3.2 Upstream v. downstream approaches to offense

This leads to the next piece of analytical vocabulary. Philosophers and legal theorists working on offense typically aren't interested in offense's social etiology. I don't just mean that they pay little attention to historical or sociological research on how offense-driven social conflict arises. (That's true of this paper as well.) I mean they don't even have a background conception of what causes offense to occur. Instead, it's commonly just assumed as a starting-point for the inquiry that, in societies like ours, the giving and taking of offense is routine and inevitable. The question is how the law should react to this, if at all. What's worrying here is that the lack of a backstory about the social origins of offense can end up standing in as an implicit backstory of its own, according to which offense-driven conflict is simply a brute fact of life, which can only be tackled remedially, since nothing can be done to prevent its occurrence in the first place. If offense is indeed a brute fact of life then it will happen all the same regardless of whether we have provocateurs intentionally seeking to offend others, or prudish moralists going out of their way to feel offended at things. It will happen irrespective of whether we have a political culture that incentivizes being offensive, or one that fosters hyper-sensitivity, or one that conscientiously tries to counteract these things. In treating offense as a brute fact of life we assume that the quantity of offense in society is something out of our control. There is no point trying to take proactive measures to ameliorate the unwelcome effects of offense for individuals, or to reduce the overall frequency of offense-driven social conflict.¹⁷

¹⁷ One way in which the under-theorization of the social origins of offense is revealed in the literature is in the kinds of cases of giving and taking offense that are treated as informative examples. Consider the example of the person who is *nauseated* by the sight of brightly-colored clothing, as discussed by Feinberg, von Hirsch, and Simester. This example involves an imaginary character with arbitrary, fictitious sensibilities. The use of this example subtly invites us to conceive of offense as something with inexplicable origins, which unpredictably pops into existence in social interaction almost like an allergic reflex. The picture of offense's origins that I'm recommending, by contrast, emphasizes the role of the offended person's judgements in interpreting the significance of their immediate (raw) feeling of aversion. This type of processed offense clearly doesn't pop into existence from nowhere. It results from people's views about which forms of behaviour are normal, decent, decorous, etc.

It is hard to deny that human beings have certain dispositions that incline us towards episodes of social discord. But even so, a conception of social intercourse on which offense is treated as a brute fact of life tempts us to view the factors that underlie offense-driven conflict as beyond anyone's influence. And this is a mistake. We may have some natural potential for discord, but the way this potential manifests in a society is affected by cultural and historical contingencies in the form of life prevailing in that society. When someone's experience of raw offense transforms into processed offense, this transformation is an expression of the judgement that the offending conduct was not only unpleasant, but somehow intolerable. And the factors that condition these judgements obviously aren't so primordial as to be fixed by the bare fact of our human nature. These sorts of judgements vary between societies, and within societies over time, and between sub-cultural enclaves that make up larger societies. For social groups in which overt class-stratification is viewed as normal and legitimate, people are habituated into offense-taking customs that are more sensitive to norms of status and deference to one's superiors. In cultures where puritanical ideas about sex are the norm, people are conditioned into offense-taking habits that are more sensitive to mores about public amorousness and exposure of the body. In social groups where nationalistic zeal is the norm, people are habituated into offense-taking customs that are more sensitive to the dishonoring of soldiers and great statesmen. From one culture to the next, people can be more or less inclined to take offense on a greater or lesser number of pretexts. And while it is highly implausible to suppose that the cultural and historical contingencies behind these dispositions are malleable enough to be engineered at will, it is also implausible to see these contingencies as fully immune to being shaped. The social agendas governments pursue in their institutional practices and policy-making can affect our habits and expectations vis-à-vis giving and taking offense.

Philosophers and legal theorists working on offense have little to say about how social policy could pre-emptively mitigate the effects of offense-based conflicts, because we generally treat offense as a *downstream* problem – a problem about how to *react* to certain bad states of affairs – and not an *upstream* problem – a problem about how to reduce the incidence of those bad states of affairs in the first place. Things could be approached in a different way. Offense-based social conflicts are destructive and hard to deescalate once in motion. As well as being bad for offendees experiencing negative mental states, they are also

bad for the rest of us, who have to inhabit the rancorous environment that trails in their wake. But as philosophers and legal theorists we needn't confine our focus to legal remedies; we can also give some thought to how social institutions can be used to reduce the incidence of offense upstream.¹⁸

3.3 Interactions between the upstream and the downstream

In response to the above one might argue that it is outside the remit of philosophy and legal theory to think about offense as an upstream problem. To see why, compare our question about the legitimacy of regulating offensive conduct with a parallel question about the scope and legitimacy of a Millian harm principle. We don't see it as the legal philosopher's job to ask how the state's full arsenal of institutional resources can be used to reduce the incidence of harmful acts upstream. She isn't expected to speculate about how measures might be taken in education or public information campaigns in order to deter harmful conduct. Rather, her job is just to explain when and why harmful conduct can be legally restricted. Granted, such restrictions will have a deterrent function that should help to reduce the incidence of harmful acts upstream. That's how deterrence works. But the legal theorist's concern for upstream harm-reduction only extends this far. By the same token, one might say, it isn't the job of legal philosophers examining the significance of offense to devise policies or institutional reforms aimed at reducing the incidence of offense. The theorist's job is to explain how, if at all, the law is justified in responding, downstream, to acts that cause offense, leaving it to others to consider how the state at large should address the wider social problems in this area.

Contrary to this line of thought, I believe it is productive to explore the upstream and downstream questions side-by-side, because there are *prima facie* good reasons to think that the downstream regulation of offense-causing conduct can have an upstream influence on the incidence of offense-driven social conflict. Making conduct liable to legal restriction because of its offensiveness can increase the overall incidence of feelings of offense. Let's assume for the sake of argument that the occurrence of raw offense is very difficult to prevent. Nevertheless, once feelings of raw offense exist, what becomes of them – whether they ebb away, or intensify into resentment and outrage – depends in various respects on how

¹⁸ One exception to my claim that legal theorists don't treat offense as an upstream problem is Husak, "Disgust: metaphysical and empirical speculations". Husak doesn't claim that the law contributes to the upstream incidence of offense. However, he thinks the legal regulation of offensive behaviour is often unjustified, and that this is more apparent once we take account of the ways in which people's propensity to feel offended by things (and in particular, *disgusted* by things) is due to their being conditioned into finding particular things disgusting.

people have been conditioned to understand and interpret those feelings. If our social milieu esteems and validates feelings of offense – if it encourages us to see them as a warranted reaction to certain stimuli – then it’s more likely that raw offense will solidify into processed offense, than if our milieu encourages us to see such feelings as fleeting and trivial. There are many factors that affect how feelings of offense are validated in a particular society. Whether the norm is to validate or withhold validation from an instance of offense will typically depend on various elements of its specific etiology, i.e. *who* did *what* to *whom*, and *when*, and *why*. And some of the factors undergirding these norms will probably only be influenced obliquely, if at all, by the formal social institutions and policies that exist in a given society at a time. Nevertheless, to the extent that any general trends are in effect, in determining whether and how feelings of offense are validated, one especially powerful way to confer esteem upon such feelings – to signal to people that their feelings of affront *are* of special importance – is to enact legislation that overtly identifies offense as a legitimate basis for the legal regulation of people’s conduct.

In short, if what converts immediate, *raw* offense into stable, *processed* offense, are social norms and expectations, by whose lights people interpret the significance of their feelings of raw offense, then laws that overtly regulate the offensive will be a factor that inadvertently fuels the incidence of processed offense. The legal regulation of offense *downstream* can therefore help promote an increase in offense *upstream*. Feinberg tells us that being offended is similar to a nuisance, in that it impairs people’s liberty by unwantedly capturing their attention. My contention here is that an approach to the regulation of offense that creates avenues for these impairments to be legally redressed is one that simultaneously tends to foster norms that will be partly instrumental in bringing about those very same kinds of impairments.

It may sound like I’m saying that anti-offense regulations encourage people to deliberately feel offended, or that it’s normal, in societies with anti-offense regulations, for people to experience processed offense in a disingenuous or knowingly hyper-sensitive way. But my argument isn’t premised on any such claim. I am saying that we all have understandings and expectations that influence how we process raw offense, and which are informed by our broader cultural resources. There is a standing temptation for all of us regard our own feelings of raw offense as substantial and worthy of other people’s special concern, and what anti-offense regulations do is signal to us that it’s reasonable to give in to this temptation. When I’m stewing over a feeling of raw offense, in a context where it could yet escalate into indignation, anger, or a demand for compensation, the fact that the law identifies feelings of offense as a legitimate basis for state intervention will, other things being equal, nudge me towards a reaction on which the feeling intensifies. No intention to take offense need be involved in this.

3.4 Law and social meaning

I will call the challenge to Feinberg's view that I have outlined in this section the *Backfiring Problem*. My claims about this problem are at least partly hostage to empirical fortune. Having said that, debates between Balancers and Hard-nosed Liberals typically don't hinge on empirical claims, so this doesn't differentiate my account from the accounts that I'm dialectically engaged with. Neither the Hard-nose who says offense is a minor irritant that can be ignored, nor the Balancer who insists that offense inhibits the exercise of people's liberty, ground their claims in social scientific data. Rather, they appeal to thought experiments and widely-shared experiences to try to elicit intuitive support for or against a legal theoretic proposition about offense's normative significance. My contribution to this debate is to point out that further considerations need to be factored into the theoretical accounting that underpins our judgements about offense's normative significance. The way that offense impairs people's liberty, if and when it does, depends on the causal process through which people come to feel offended. Given (i) the role of interpretative judgements in converting raw offense to processed offense, (ii) the role of social meanings in guiding such judgements about offensiveness, and (iii) the role of the law in establishing social norms and social meanings, the legal regulation of offense should be expected to have an effect on the overall incidence of offense. While empirical evidence could, in principle, be presented against any of these claims, I have tried to explain why they are *prima facie* plausible, concentrating on (i) in §3.1, and (ii) in §3.3.

The point that I have said the least about thus far is (iii). With respect to this claim I'm drawing from two related lines of thought in legal theory. The first is the notion – informed by French social theorists like Foucault and Bourdieu, and then influentially espoused by Lawrence Lessig, among others – that the acts of law-makers and governments “construct the social structures, or social norms, or... social meanings that surround us”.¹⁹ By social meanings, Lessig means the semiotic content that attaches to acts and entities – what those things convey, symbolize, and signify. One obvious way that the law imbues acts with social meaning is by *stigmatizing* them. Onerous regulations on cigarette smoking, for instance, transform the meaning of smoking from something appealing (or at worst neutral) into something seedy and shameful. Not all social meanings are actively constructed, and of those that are, not all of them are constructed via law and government action. But given the countless ways in which the rule of law shapes our

¹⁹ Lawrence Lessig, “The regulation of social meaning”, *The University of Chicago Law Review* 62 (1995): 943-1045, 951. For other influential legal theoretic accounts of the law's role in creating social meanings see Dan M. Kahan, “Social influence, social meaning, and deterrence”, *Virginia Law Review* 83 (1997): 349-95; J. M. Balkin, *Cultural Software: A Theory of Ideology* (New Haven: Yale University Press, 1998).

lives, and given how we routinely internalize law's demands, the law is a powerful source of social meanings. The second and related line of thought that I'm drawing on is one that stresses the law's expressive functions and powers. For example, the law has an important *informative* function. As Richard McAdams says, Legal restrictions convey information to those they govern about the imprudent, anomalous, or socially-disapproved nature of the acts that they restrict, and thus they have an ability to guide people's conduct that extends beyond their capacity to deter or persuade people of their legitimacy.²⁰ Given this kind of expressive theory of law, "debates over the appropriate content of law" can often be construed as "debates over the statement that law makes".²¹ There is of course plenty of room for disputes about the *extent* of the law's expressive powers, and its influence compared to other sources of information and normative guidance. Social scientific research may provide us with a more complete understanding of the relative magnitude of the law's role in establishing the social norms and meanings than guide judgements about offensiveness. But to assume that the law must play *some* such role, as I do here, isn't to venture out too far on a limb.

The conclusion that I'm drawing on the basis of these claims is that the imposition of anti-offense regulations will tend to increase the overall incidence of offense. One might challenge this by noting that there are other types of legal regulations, including anti-nuisance laws, that should give rise to a similar tendency, given the offense-generating process that I've proposed. But it seems unlikely that the legal regulation of nuisance-causing conduct increases people's propensity to be annoyed by nuisances. Why suppose that anti-offense regulations would be different? There is much that could be said here, but the most important thing to note, by way of reply, is that the paradigmatic sources of nuisance – cacophonies and rancid odors – are things whose power to frustrate and disturb is relatively culture-independent. The diversity in the range of things that provoke offense across different cultures is wider than in the range of things that cause sensory nuisance. This is indicative of the fact that offense, unlike most sensory nuisance, is highly sensitive to social meanings. Law is much less likely to have an effect on people's propensity to be annoyed by nuisances, then, since this propensity is much less sensitive to the kinds of social meanings that the law contributes to.

²⁰ Richard H. McAdams, *The Expressive Powers of Law: Theories and Limits* (Cambridge Massachusetts: Harvard University Press, 2015).

²¹ Cass R. Sunstein, "On the expressive function of the law", *University of Pennsylvania Law Review* 144 (1996): 2021-53, 2051.

4. The implications of the backfiring problem

In this section I'll explain the main significance of the Backfiring Problem. In short, it shows how it can make sense for liberals to oppose the legal regulation of offense, even if they grant, in-principle, the key premises and inferences behind Feinberg's account from §2.

Hard-nosed Liberals – those who categorically oppose the legal regulation of offense – haven't provided an adequate response to Feinberg's in-principle justification for regulating offense. They think offense must be tolerated because it is an inevitable upshot of the social friction that occurs in the type of pluralistic society that liberalism committed to. But even if this is right, the Hard-nose's hard-line stance doesn't follow. If Feinberg is right that offense impairs people's liberty, then that is still a *pro tanto* good reason for the liberal state to regulate offense, even if it is true that offense is ultimately an inescapable result of the cultural friction that occurs in a liberal society. Can anything more be said for the Hard-nosed view?

One thought – descended from Mill in *On Liberty*, and elucidated by Waldron – appeals to offense's salutary potential.²² As an autonomous agent in a liberal society, seeking to assess different conceptions of the good in order to determine how I myself will live, it is a good thing to have my convictions attacked and disparaged, so that I can see whether, by my own lights, they withstand that challenge. Such attacks will of course typically provoke feelings of offense. But given that the attacks are ultimately beneficial, it is an error to regard their affective side-effects as 'pseudo-harms' that need remedying. Instead, we should see these feelings as a healthy sign that people are grappling with challenges to their ideals. As Waldron puts it, "outrage and disturbance" are "to be welcomed, nurtured and encouraged in the free society that Mill is arguing for".²³ The appeal of this line of thought is reflected in some attempts to defend deliberately provocative speech – in the arts, for instance, or in religious critique.²⁴ But it is hard to see how this line of thought could force us to accept the Hard-nosed position. After all, even if there ultimately are certain benefits to the offende'e's liberty that come with being offended, there is no good reason to suppose that these will always outweigh the short-term impairments to liberty that Feinberg's defence of Balancing adverts to.

²² John Stuart Mill, *On Liberty* (London: Penguin Books, 1985) (originally published 1859).

²³ Waldron, "Mill and the value of moral distress", *Political Studies* 35 (1987): 410-23, 413.

²⁴ See for example Amy Adler, "What's left? Hate speech, pornography, and the problem for artistic expression", *California Law Review* 84 (1996): 1499-572; David Edgar, "Shouting fire—from the nanny state to the heckler's veto: the new censorship and how to counter it" in Ivan Hare and James Weinstein (Eds.), *Extreme Speech and Democracy* (Oxford: Oxford University Press, 2009).

Another argument for Hard-nosed Liberalism opposes the regulation of mere offense because it construes this as a veiled form of legal moralism, that is to say, a way of restricting conduct merely on account of its purported moral wrongness. Liberals standardly reject such moralism, instead espousing a form of politics that accommodates a broad plurality of moral commitments, and whose purpose is to allow the bearers of those diverse commitments to cooperate and live autonomously in a well-ordered society.²⁵ Now, the *rationale* for Feinberg's offense principle may not appeal to any moral doctrine, besides its axiomatic prioritization of individual liberty. Having said that, his principle's *institutionalization* could facilitate a kind of *de facto* moralism, insofar as it allows that offensive acts can be regulated if enough people are severely offended by them.²⁶ But again, this doesn't force us to accept Hard-nosed absolutism. If and when anti-offense regulations operate in a *de facto* moralistic way, liberals have grounds for opposing them. But there is no reason to think that this will always be the case.

A third hard-nosed argument, pressed by Ronald Dworkin, is that there are species of expression like satire or some especially confrontational types of protest which offend by their very nature, but whose restriction would violate the kind of free speech principles that liberals are axiomatically committed to.²⁷ This concern about the priority of free speech is accounted for in Feinberg's analysis, though. When it comes to weighing-up the liberty-impairing effects of offense itself against the liberty-impairing effects of anti-offense regulations, Feinberg is clear that "no degree of offensiveness in the expressed opinion itself is sufficient to override the case for free expression".²⁸ On this count, then, he agrees with Dworkin. The content of an expressed opinion cannot be cited by the liberal state as a reason for restricting expression. Both authors agree, moreover, that the *manner* of expression *can* be cited by the liberal

²⁵ This is meant to be a sketch of the type of mainstream liberalism defended in, e.g., John Rawls, *A Theory of Justice* (Cambridge Massachusetts: Belknap Press, 1971), Thomas Nagel, "Moral conflict and political legitimacy," *Philosophy & Public Affairs* 16 (1987): 215-40, and Brian Barry, *Justice as Impartiality* (Oxford: Clarendon Press, 1995).

²⁶ Arguments along these lines are developed in Anthony Ellis, "Offense and the liberal conception of the law", *Philosophy & Public Affairs* 13 (1984): 3-23, and Michael Bayles, "Offensive conduct and the law" in Norman S. Care and Thomas K. Trelogan (Eds.), *Issues in Law and Morality* (Cleveland: Western Reserve University, 1973).

²⁷ See in particular "The right to ridicule" (note 1), "Liberty and pornography" (note 2), "Is there a right to pornography?", *Oxford Journal of Legal Studies* 1 (1981): 177-212, and *Sovereign Virtue: The Theory and Practice of Equality* (Cambridge Massachusetts: Harvard University Press, 2000): 380-81.

²⁸ Feinberg, *Offense to Others*: 44.

state as a reason for restricting expression, if the problematic element in the expression's manner isn't enmeshed with the expression's message.²⁹ The discrepancy between the two authors is ultimately quite subtle. In essence, Feinberg is more optimistic than Dworkin about the law's ability to disentangle substance and manner in a way that will sometimes make it acceptable for the liberal state to regulate a message because of the gratuitously offensive manner in which it is expressed.³⁰ It is possible that this difference will translate into differing prescriptions on specific free speech controversies at the margins. But the key point for Dworkin – that a commitment to free speech requires us to tolerate certain kinds of messages that are sure to offend some people – isn't merely *compatible* with Feinberg's view, it is expressly incorporated into his account. Balancers can fully endorse the liberal prioritization of freedom of speech.

All three arguments raise concerns that liberal lawmakers need to take seriously in any undertaking to regulate offensive conduct. But they do little to weaken Feinberg's case for the offense principle, because they don't properly confront, let alone rebut, his key premise, that offense is comparable to nuisance in its tendency to impair the exercise of liberty. That premise is hard for the liberal to resist, because it embodies the liberal creed that the final measure of a mid-level principle is its implications for individual liberty. What makes the Backfiring Problem significant, then, is that it shows how it could make sense, in theory, for a liberal to maintain a fixed opposition to the regulation of offense, as Hard-noses do, even while accepting and remaining cognizant of Feinberg's crucial premise about offense's negative effects. If feelings of offense really do impair people's liberty, much like nuisances, then the only way to defend a fixed opposition to the offense principle is if you can argue that permitting the regulation of offense will do more harm than good relative to the standards of assessment that Balancers themselves appeal to. Many liberals are instinctively attracted to a Hard-nosed view of offense. The Backfiring Problem shows what it would take, in the wake of Feinberg's decisive objection to that view, to re-establish a position that's functionally equivalent to it.³¹

²⁹ In Feinberg's words: "offensiveness in the manner of expression, as opposed to its substance, may have sufficient weight [to override the presumptive priority of free speech] in some contexts"; *Ibid.* 44.

³⁰ *Ibid.* 216-17.

³¹ One possible position would be to accept that the Backfiring Problem is indeed a problem, relative to Feinberg's standards of assessment, but to think that the costs to liberty that would result from having anti-offense regulations which backfire would be less than the costs to liberty that come with leaving offense unregulated, e.g. because anti-offense regulations deter more offense overall than they indirectly generate. Concerns about backfiring don't force someone who grants Feinberg's key premise to reject his conclusion. The significance of these concerns is that they show how it could be reasonable for someone who accepts Fein-

This also reorients the debate among liberals over how to deal with offense and offense-driven social conflict. The debate between those liberals who instinctively resist the regulation of offensive conduct, and those like Feinberg who are untroubled by it, can be reformulated as a debate about how the liberal state should define the social meaning of offense. It is implausible for Hard-noses to claim that offense has no negative effect on the exercise of liberty, and it is unprincipled for them to say that these negative effects don't matter. What is less obviously mistaken is the worry that semiotically equating injury and offense will contribute to an erosion in the kind of social ethos that a liberal society requires. There are uncertainties about precisely what the effect of any particular instance of anti-offense regulation will be with regards to the social meanings that surround the giving and taking of offense, as I noted in §3.4. But the worry that using the law to regulate offensive conduct will alter the social meaning of offense in ways that create problems for liberalism, on its own terms, makes sense even while those uncertainties remain. That is what the Backfiring Problem alerts us to. And these kinds of concerns are compatible with granting the key insight in Feinberg's defence of Balancing – that being offended can genuinely impair the affected individual's liberty.

5. What about appropriately-taken offense?

The Backfiring Problem is made more salient by the new social politics of offense. At historical moments where most people in a society takes offense at the same things – when there are mores around bodily exposure that are ostensibly universal, for example, due to an ethical monoculture, or convergent sensibilities across different groups in society – the role of regulations that police those things, in influencing people's propensity to be offended by them, is less evident and less of an issue. The Backfiring Problem arises as a concern in societies where different constituencies are offended by different things, and where policy that aims to placate one constituency's feelings of offense simultaneously stirs another constituency's ire, by regulating activities that they see as an ethical prerogative. Thus far I have been avoiding the question of whether some acts are offensive in an objective sense, such that it's right and proper to take offense in response to them. I have skirted around this because of the difficulties that it raises given the deep conflicts that define today's liberal politics. It's true, in principle, that if certain acts are objectively

berg's key premise to reject his conclusion. Note also that Feinberg's offense principle (note 13) has a qualification about *effectiveness*, such that one could accept all the implications of the Backfiring Problem that I've suggested and still endorse *the letter* of his principle.

offensive, then it is probably wrongheaded to worry about the law's potential to influence the transformation of raw offense into processed offense in response to those acts. However, I doubt that this point goes very far towards undermining the objection to Feinberg's account that I've been developing, built around the Backfiring Problem.

Consider a contrast with the regulation of sexual harassment. The introduction of such regulations in a society may nurture people's felt sense of grievance in response to experiences of harassment. It may legitimize adverse feelings that victims previously tended to downplay or ignore, in a way that resembles the processes I described in §3. But in this context that seems appropriate. Acts of sexual harassment should be regulated because they're seriously wrongful, and if this galvanizes the victims' felt sense of grievance then that is to the good, if anything, since it makes it harder for the justice system to ignore the problem. Now, sexual harassment involves wrongful acts that inflict harm, as opposed to wrongful acts that merely offend. But the rationale can arguably be carried over to apply in cases of offense as well. Simester and von Hirsch defend a version of the offense principle on which the justifiability of regulating offensive conduct owes not to its creating nuisance-like impairments to liberty, à la Feinberg's account, but to its *wrongfulness*, specifically, its disrespectful or inconsiderate treatment of others.³² Their core cases are insults and exhibitionism. Insults are disrespectful. Exhibitionism is inconsiderate. Acts of both kinds can be legally regulated by the liberal state, in principle, in view of their wrongfulness in these respects. And if one consequence of imposing such regulation is to help sensitize people towards being offended by these acts, that seems like a benign side-effect at worst, and a helpful incentive towards law-abiding conduct at best.³³

The limitation that one butts up against, however, in pursuing this line of argument, is that it is a matter of profound disagreement between different sectors of society whether there in fact *is* anything wrongful – unacceptably disrespectful, unreasonably inconsiderate, or objectively offensive – about the

³² See A. P. Simester and Andreas von Hirsch, *Crimes, Harms, and Wrongs: On the Principles of Criminalisation* (Oxford: Hart, 2011): 97ff.

³³ Simester and von Hirsch's view suggests that there is at least a *prima facie* reason why the intentional cultivation of offense-taking sensibilities mightn't necessarily be a bad thing, in the way that intentional cultivation of a susceptibility to experiencing nuisance probably *would* be necessarily bad. But this is a region of judgement in which the liberal state needs to tread very carefully. Even the core and supposedly clear cases of wrongful conduct that Simester and von Hirsch appeal to, namely, insults and exhibitionism, are things whose wrongfulness is a more controversial matter than they admit. It is true in some broad sense that insults are disrespectful and that exhibitionism is inconsiderate. But whether there is a discernible level of disrespectfulness or inconsiderateness at which such acts clearly qualify as wrongs is something that reasonable people can and do disagree about.

kinds of offense-causing conduct that lie at the heart of major episodes of offense-driven social conflict. Suppose we grant something like Simester and von Hirsch's position for the sake of argument. Let's say that for types of conduct that aren't just offensive, but also manifestly and uncontroversially wrongful, there is a *pro tanto* case for regulating them in view of their wrongfulness. And in such cases, let's say, it is true that the possibility of the regulation nurturing offense-taking sensibilities isn't a reason to oppose it, since being offended can be taken as a fitting response to such conduct. My point is that we could grant all this, and it still wouldn't give us any guidance about how the law should deal with the offense-causing conduct that generates major social conflict. Again, this is because there is deep disagreement about the ethical character of such conduct, which is of course precisely the reason *why* it is a flashpoint for offense-driven social conflict in the first place.

Part of Feinberg's reason for analogizing offense and nuisance is to try to identify a rationale for regulating offense that needn't be premised on the state taking one side in these kinds of disagreements. He says that only protecting people against offense that is appropriately-taken "would require agencies of the state to make official judgements of the reasonableness... of emotional states and sensibilities", which would be "dangerous and distinctly contrary to liberal principles".³⁴ As far as Feinberg is concerned, then, offense is a negative condition that the liberal state has reason to try to remedy regardless of any verdict about its contextual appropriateness. To take this position isn't to commit oneself to some sweeping and theoretically precarious notion of liberal neutrality. All that is being assumed here a fairly modest form of liberal *pluralism*.³⁵ Any recognizably liberal society must try to accommodate devoutly religious people alongside deeply anti-religious people, without sponsoring the agenda of one group over the other. In modern multicultural societies that means our legal system shouldn't encode any conclusions about whether it's appropriate to be offended by things like hard-edged religious satire. It still makes sense, then, to be concerned about the Backfiring Problem, and the law's potential to guide the transformation of raw offense into processed offense, because even if we grant that this *wouldn't* be something to worry about in cases of objectively offensive conduct, liberals will regard such cases as being beyond the law's radar.

³⁴ Feinberg, *Offense to Others*: 36-37.

³⁵ Feinberg's remarks about the inappropriateness of the state assessing the reasonableness of offense only pertain to the part of his account that is concerned with determining the intensity of people's feelings of offense. At other points his account seems to appeal to substantive values, beyond the mere axiomatic prioritization of individual liberty, and thus his account probably isn't strictly committed to a standard of liberal neutrality.

6. The politics of leaving offense unregulated

In major offense-driven controversies today the offendees belong to groups that suffer more than their fair share of disrespect and hostility. Certain liberals will instinctively reject any theory of offense that advises against the legal regulation of offense, then, because they recognize that in Western societies today, for members of religious minorities – Muslims in particular – offensive derision in public discourse adds insult to injury. It is understandable, in light of this, to want to counter the efforts of those who use offensive ridicule to add insult to injury. I have been discussing relatively abstruse theoretical concerns, about how the regulation of offense could condition a cultural ecosystem in ways that are unwelcome for Feinbergian liberals on their own terms. But does accepting this kind of argument against Feinberg's account mean that we have to downplay the costs for those targeted with offensive derision?

Not necessarily. First, note that adverse effects that are informally characterized as offense may be properly characterized as harms in a legal context. For example, when someone faces racist abuse in the workplace he may describe this by saying that he was subjected to offensive remarks, or that he felt offended. But the psychological toll of workplace abuse often goes beyond mere offense. Such abuse often effects a wrongful setback to the target's interests, and can therefore be characterized as a harm. Something similar is true for larger cultural conflicts around religion. Granted, in some cases, including some of the high-profile cases of offense-driven social conflict mentioned in §1, members of minority religious groups are merely offended rather than harmed. But some types of scornful expression about a person's religion – in particular, expression that is targeted against a specific individual – aren't just offensive, but rather, constitute harmful instances of harassment, threats, or vilification. And there is nothing in the debate between Balancers and Hard-noses that calls into question the overarching assumption that harmful communicative acts of these types are liable to regulation, both in principle and in practice. To oppose anti-offense regulations is not to license any hostile expression whose effects might be casually described using the language of offense.

With respect to anti-hate speech law specifically, I would stress that arguments for such laws generally aren't based on the charge that hate speech is offensive to its targets. As Waldron says, "it is not the function of racial or religious hatred laws to protect people against hurt feelings", and in debates about the case for anti-hate speech laws, discussion of the viability of an offense principle, he says, is little more

than a distraction (one which he thinks is often introduced into these debates by people aiming to discredit such laws).³⁶ Where the language of offensiveness is used in particular anti-hate speech statutes, some superficial reforms might be called for given the view that I'm defending.³⁷ However, for most defenders of anti-hate speech law, it is more appropriate for such laws to focus on discriminatory speech that's used to threaten, harass, or via some other mechanism inflict harm on its targets. In sum, debates about the case for regulating hate speech are orthogonal to debates about the legitimacy of anti-offense regulations. Expression that merely offends people isn't what advocates of anti-hate speech law are trying to combat, and as far as my argument goes there may still be a decisive case for restricting hate speech. It will just be one that doesn't invoke offense.

Another way in which we can maintain concern for the welfare of people on the marginalized side of offense-driven conflict, while still eschewing the regulation of offense, is in policy-making aimed at minimizing the incidence of offense-driven social conflict upstream. Suppose we are devising a public information campaign aimed at reducing offense-based social conflict by promoting an ethos of cross-cultural understanding. If we see offense as a negligibly bad experience, we will be indifferent between messages calling on offenders to be more conciliatory and culturally sensitive, and messages calling on *offendees* to thicken their skin. But such indifference would be harsh and unreasonable. Offense-driven conflict is bad for all of society, but it isn't bad for all of society's members in the same way or to the same degree. Methods of trying to reduce it that burden the already marginalized will rightly be resisted by those who recognize that offense can sometimes be a real hardship, not just a trifling irritation. And there is no reason why any of this need be denied or downplayed by the liberal who opposes anti-offense regulations because of the reasons that I've been presenting.

Here is one last reason why my argument against Feinberg's account needn't lead us towards a *laissez-faire* indifference about the costs of identity-political social conflict. Much political discourse about Is-

³⁶ Jeremy Waldron, *The Harm in Hate Speech* (Cambridge Massachusetts: Harvard University Press, 2012): 111.

³⁷ Although the leading defences of anti-hate speech law don't invoke offense, some anti-hate speech statutes in particular jurisdictions do advert to offensiveness among the identifying characteristics of the expressive conduct that they aim to regulate, as in section 18C of Australia's Racial Discrimination Act (1975), for example.

lam in the West today is structured by a ‘clash of civilizations’ narrative, which aims to promote and normalize the idea that Islam cannot be peacefully accommodated within liberal Democracy.³⁸ When a conflict that begins with the giving and taking of offense subsequently escalates into acts of violence and diplomatic crises, as in the Danish cartoons and Charlie Hebdo cases, the peddlers of this schismatic narrative get a welcome boost. Reactionary anti-Muslim ethno-nationalists *and* extremist Islamist preachers *both* get attention and a semblance of credibility in the aftermath of such events. The liberal policy-maker’s aim, in trying to negotiate these volatile dynamics, should be to defuse the rancor, and to encourage modes of social intercourse in which raw offense is less likely to be amplified – by *agents provocateurs* of whatever political persuasion – into overt belligerence and violence.

On the surface, Feinberg’s approach may seem like the one that best conduces to that aim. But I have argued that his approach, once it’s unleashed in the wild, will ultimately tend to support a cultural milieu in which it is easier for raw offense to be transformed into volatile and widespread acrimony. When the state symbolically ratifies offense, by treating it as a justification for the legal regulation of its originating trigger, it can inadvertently facilitate the agenda of those who would actively seek to foment offense-based social conflict in order to advance their own divisive political agenda.³⁹ Obviously it is a complex question how all the factors contributing to these social dynamics interact, and it would be implausible to assert that the regulation of offense is somehow the one paramount factor that allows the damaging escalations that I’m describing to occur. My point is that the liberal who opposes the legal regu-

³⁸ ‘The clash of civilizations narrative’ refers to accounts of politics that accord special explanatory significance to the allegedly vast differences between Western/Christian versus middle-Eastern/Islamic worldviews; the term is due to one of the popularizers of such accounts, Samuel P. Huntington, *The Clash of Civilizations and the Remaking of World Order* (New York: Simon & Schuster, 1996).

³⁹ Certain perverse incentives to foment offense were evident in the Danish cartoons controversy. The intensity of that conflict was partly due to a few individuals who circulated a dossier of the offending images to Muslim groups around Europe and the Middle East, with more extreme images – which hadn’t been published in the Danish newspaper at the center of the controversy – misleadingly mixed in. These actions meant that many more people were exposed to the images than would have seen them in the Danish press, and that many of these people were under a misapprehension about what had actually been published. For a detailed account of these events see Jytte Klausen, *The Cartoons That Shook the World* (New Haven: Yale University Press, 2009).

lation of offense in principle needn't be indifferent to these conflicts and their negative effects on minority religious communities. There are reasons to adopt this position that stem from concerns about these very effects, and the law's symbolic role in promoting them.⁴⁰

⁴⁰ This piece has benefited from the input of many people. Two anonymous referees from this journal provided detailed comments that improved the paper a great deal. I am grateful to audiences at the Australian National University, the University of Queensland, Monash University, the University of Melbourne, and the 2017 Australian Political Philosophy Conference for their feedback. Thanks to Seth Lazar, Kath Gelber, and Dan Halliday for speaking invitations at ANU, UQ, and Melbourne respectively. For other comments, criticisms, and suggestions, I'm also grateful to Susan Brison, Les Green, Toby Handfield, Katrina Hutchison, and Martha Nussbaum.