Intellectual Trust and the Marketplace of Ideas

Allan Hazlett

Forthcoming in M. Lynch and A. Tanesini (eds.), *Polarization, Arrogance, and Dogmatism: Philosophical Perspectives* (Routledge).

Liberals have a problem with offensive speech. I don't mean "liberals" in the American sense of left-wingers, and I don't mean that they "have a problem" in the sense that they are offended. I mean "liberals" in the traditional sense, indicating those who posit a familiar suite of basic liberties, including freedom of thought and expression, freedom of religion and association, freedom of movement, and freedom to engage in consensual transactions with others, whether sexual, aesthetic, or economic. And I mean that they "have a problem" with offensive speech because offensive speech, which is tolerated by the principle of freedom of expression as it is standardly endorsed by liberals, has bad consequences.

My aim in this chapter is to show how liberals can respond to the problem of offensive speech using the familiar liberal concept of a "marketplace of ideas." I shall argue that intellectual trust among its participants is conducive to the functioning of a marketplace of ideas, and that when such intellectual trust is present, the bad consequences of offensive speech are mitigated.

1 Marketplaces of ideas

The idea of "the marketplace of ideas" has a rich history, including in the history of the interpretation of the United States Constitution's provision of freedom of speech. Although he doesn't use the phrase "marketplace of ideas," the idea is often associated with John Stuart Mill's defense of freedom of opinion and expression in *On Liberty*. Mill (1859/1989) posits four benefits of "freedom of opinion, and freedom of expression of opinion": (i) that censored opinions "may, for aught we can certainly know, be true"; (ii) that a censored false opinion "may, and very commonly does, contain a portion of truth"; (iii) that unless the truth "is suffered to be, and actually is, vigorously and earnestly contested, it will, by most of those who receive it, be held in the manner of a prejudice, with little comprehension of feeling of its rational grounds"; and (iv) that without such a vigorous and earnest contest of ideas "the meaning of the [true] doctrine itself will be in danger of being lost, or enfeebled, and deprived of its vital effect on the character and conduct." (p. 53)

The idea of a beneficial "contest" of ideas is central to the idea of a marketplace of ideas, where the implication is that a *market* is characterized by *competition* among sellers for the custom of buyers. Just as competition among bakeries will, if all goes well, create an incentive for them to produce high-quality bread, resulting in high-quality bread being offered for sale, competition among those who articulate ideas will if all goes well, create an incentive for them to produce high-quality ideas, resulting in high-quality ideas being articulated. It doesn't matter whether we think of a marketplace of ideas as a species of market, in which ideas are bought and sold, on sufficiently broad notions of buying and selling, or whether we think of a marketplace of ideas as merely analogous to an actual market. "Selling" an idea just means asking others to believe or accept it, or at least to take it seriously as a potential object of belief or acceptance, and "buying" an idea just means doing so. The "custom" sought by those who articulate

ideas is often the transfer of material wealth, but just as often the bestowal of status, honor, respect, or some other desirable social good.

The establishment of a market requires the establishment of two liberal principles of *free trade*, on which sellers are free to sell what they want and buyers are free to buy what they want. Without the establishment of these principles, there will be no competition among sellers for the custom of buyers, and thus none of the beneficial effects of such competition. Likewise, the establishment of a marketplace of ideas requires the establishment of two familiar liberal principles: *freedom of expression*, on which those who articulate ideas are free to articulate what ideas they want, and *liberty of conscience*, on which those to whom ideas are articulated are free to believe or accept or take seriously what ideas they want. To put this another way, a marketplace of ideas is established when speakers are free to say what they want and hearers are free to think what they want. Without this, there will be no competition among speakers for the "custom" of hearers, no vigorous and earnest contest of ideas, and thus none of the beneficial effects of such competition.

Two comments are in order. First, just as what counts as "high-quality bread" is determined, in our toy example of a market, entirely by the actual preferences of the buyers, what counts as "high-quality ideas" is determined, in a marketplace of ideas, entirely by the actual preferences of the hearers – or, more exactly, by their dispositions to believe, accept, or take ideas seriously. The value of "high-quality ideas," like the value of "high-quality" goods of other sorts, is an entirely open question, depending as it does on the actual preferences of the relevant people. In his defense of the utility of a marketplace of ideas, Mill seems to assume a generic preference for true ideas well supported by evidence and arguments. But even if you thought that most actual human beings had such a preference, you might well think that they also have various other preferences – for sensational ideas, for comfortable ideas, for ideas that arouse partisan animosity – such that the establishment of a marketplace of ideas cannot be expected to have beneficial effects vis-à-vis knowledge generation and transmission. I don't have anything novel to say about this problem for Mill's defense of a marketplace of ideas. Briefly, what I think those sympathetic to Mill should say here is that human beings have a rotten track record of illiberal attempts to generate and transmit knowledge in the absence of a marketplace of ideas. But, as I said, that is nothing new.

Second, note that the present articulation of the idea of a marketplace of ideas is a consequentialist defense of freedom of expression and liberty of conscience, as opposed to a defense in terms of natural rights. We should thus expect exceptions to the principles of freedom of expression and liberty of conscience. Compare a consequentialist defense of free trade, which allows for various forms of market regulation and intervention for the sake of promoting competition, improving efficiency and productivity, consumer protection, and so on.

2 The problem of offensive speech

For my purposes here, we can define offensive speech as speech that causes offense. There is, of course, another, and more natural, sense of the "offensive," on which the offensive is that which *warrants* offense. In the present sense, the fact that someone is offended by some instance of speech is sufficient for its being offensive; in the other sense, even if someone is offended by some instance of speech, it remains an open question whether that speech really is offensive.

Given this definition, the problem of offensive speech is straightforward: offense (being offended, taking offense) is a negative emotion, in the sense that it is unpleasant, annoying, upsetting, or painful.[1] The establishment of a marketplace of ideas, comprising the establishment of freedom of expression and liberty of conscience, allowing for the production of offensive speech, is thus inherently risky. Because offense is a negative emotions, the causation of offense is pro tanto bad, because it is pro tanto bad to be displeased, annoyed, made upset, or pained. The problem of offensive speech is, therefore, something with which the consequentialist defender of free speech (§1) must contend. (Contrast the liberal who defends free speech as a natural right.)

We can bracket the question of whether offensive speech is harmful. For some liberals, societal restriction of offensive speech would be legitimate only if offensive speech is harmful, because only harmful conduct may legitimately be restricted by society.[2] If offensive conduct (in general) is defined simply as conduct that causes offense, then it seems that offensive conduct is not *per se* harmful.[3] You are not necessarily harmed when you are displeased, annoyed, made upset, or pained; negative emotions are not *per se* harmful, and there is nothing about offense, in particular, to suggest that being offended is per se harmful to the person offended. Of course, there is surely a sense of "harm" on which negative emotions are plausibly per se harmful: if a person is harmed whenever their well-being is diminished to some extent, and pain always diminishes a person's well-being to some extent, then negative emotions are per se harmful. But that is certainly not the sense of "harm" intended by those liberals who maintain that only harmful conduct may legitimately be restricted by society; they have in mind bodily injury and material disadvantage. Whether or not offensive speech is harmful, it causes something pro tanto bad, and therefore poses a problem for the consequentialist defender of free speech

I would like to illustrate the problem of offensive speech with an example, and I will be using an example of speech that offends *me*. What offends me, however, may not offend you, so I would encourage you to come up with your own example. My example is the slogan "all lives matter." In contemporary American political discourse, "all lives matter" is a response to and critique of the Black Lives Matter movement, an activist movement that opposes racist violence and systemic racism that arose in 2013 and came to prominence in 2014 through activism in response to the police killings of Michael Brown in Ferguson, Missouri, and Eric Garner in Staten Island, New York. Opponents of the Black Lives Matter movement respond by saying that "all lives matter." This, I find, is utterly infuriating. The slogan "all lives matter" offends me because it is offered to *negate* the idea that black lives matter, even though this negation is disguised as an egalitarian principle, superficially consistent with the proposition that black lives matter. What you are saying when you say "all lives matter" is that black lives do *not* matter, not in the way that they are asserted to matter by opponents of racist violence and systemic racism, and yet, rather than saying that, you cynically offer a universal humanist platitude in the service of an anti-humanist, anti-universalist message of white supremacy. "All lives matter" is coy, sarcastic, paternalistic – it casts black people as selfish and narrow-minded ("They only talk about how *their* lives matter") and white people as enlightened and impartial ("We care about *all* lives, not just our own"). Finally, the slogan "all lives matter" upsets me because it is an unsubtle part of a defining feature of the logic of the

[1] Cf. Feinberg 1985, p. 1.
[2] This is Mill's (1859/1989) famous "harm principle." For a different use of the term "harm principle," see Feinberg 1973, 1984.
[3] Cf. Feinberg 1985, pp. 1-3.

political system of white supremacy: the denial of the reality of race. "All lives matter" is a short and simple way of saying that we should not talk about or think about race – and that anyone who does is ignorant or malevolent. And it just pisses me off to hear people say that, whilst casting themselves as occupying the moral high ground, using this high-minded moral truism – it is the *moralizing*, in the service of white supremacy, that is so awful about "all lives matter."

OK, so, I'm getting upset just thinking about it. And that is the point. Think of an example of speech that offends you, and dwell on it for a while, and you can experience the problem of offensive speech.

A few comments on the conceptions of offense and offensive speech that I am employing here. First, my example may suggest the following familiar idea: that "liberals" (in the sense of left-wingers) are particularly disposed to offense. Thus, the caricature of "liberals" as hysterical, "coddled," "crybullies," "snowflakes," etc. But this is no part of the present conception of offense. Moreover, "liberals" are not particularly disposed to offense. "Conservatives" (in the sense of right-wingers) are just as disposed to offense, but they are offended by different things: mistreatment of the American flag or National Anthem, women wearing headscarves and veils, trans women using the bathroom, descriptions of actions and institutions as "racist," etc.

Second, my example illustrates that it is no part of the present conception of offense that offensive conduct is represented as a wrong against the offended person. There is an alternative sense of "offense" on which to be offended is to represent yourself as wronged by the offensive conduct.[4] This is a common way of thinking about offense, where we imagine the offended party as in one way or another the *target* of the offensive conduct. This is what we are thinking when we think of racist conduct as offensive to people of color, sexist conduct as offensive to women, transphobic conduct as offensive to trans people, etc.[5] In this sense of "offense," those offended by the slogan "all lives matter" would, at least primarily and in the paradigm case, be black people or perhaps Black Lives Matter activists, not your white, politically moderate author. In the present sense of "offense," however, being offended is orthogonal to whether you are the target of the offensive conduct.

Third, I would like to set hate speech aside, as a special case of offensive speech, and my example is designed to be an example of offensive speech and *not* an example of hate speech. Of course, the scope of hate speech is controversial, and not all readers will agree with me on the classification of the slogan "all lives matter" as offensive speech but not hate speech. Those who disagree with me on this point will need to substitute their own example. Now, the definition of hate speech is likewise controversial. However, one illuminating way to distinguish between hate speech and merely offensive speech is to conceive of hate speech in terms of what it does to its target, such that, for example, racist hate speech is understood essentially in terms of what it does to people of color. Specifically, on the kind of view I have in mind, hate speech insults, demeans, subordinates, denigrates, threatens, disrespects, or attacks its target.[6] By bracketing hate speech, we can also bracket the question of whether hate speech harms its

[4] Cf. Feinberg 1985, p. 2.
[5] Cf. Feinberg 1973, pp. 102-3.
[6] Cf. MacKinnon 1987, Waldron 2012, Langton 1993, Langton and West 1999, Maitra 2009, 2012.

target.[7] It may be that hate speech harms its target, and is thereby legitimately restricted by society, whilst merely offensive speech does not.

How have liberals responded to the problem of offensive speech? The non-consequentialist liberal, as I have suggested, may acknowledge the pro tanto badness of offense, but see it merely as an unfortunate consequence of our natural right to free speech. One familiar kind of liberal response is to downplay the pro tanto badness of offense, dismissing this as merely a matter of "hurt feelings." In effect, this is to suggest that only harmful speech would present a problem for the liberal. Another familiar kind of liberal response is to argue that offense caused by speech is not the speaker's responsibility, but rather the result of weakness, over-sensitivity, misinterpretation, or confusion on the part of the offended hearer.[8] ("I'm sorry if anyone was offended," offensive speakers often say when they refuse to apologize.)

As I mentioned (§1), the consequentialist defender of free speech should expect exceptions to the principles of freedom of expression and liberty of conscience. A marketplace of ideas is, at least in principle, compatible with the restriction of pornography, hate speech, seditious speech, and blasphemy – and all of these have been legally restricted, in various ways and at various times, in the United States, despite the First Amendment's provision of freedom of speech, and *liberal* defenses of the legal restriction of pornography and hate speech are familiar.[9] Nothing in the very idea of a marketplace of ideas precludes such restriction – nor does it preclude, for that matter, the restriction of *lèse-majesté*, uninformed speech, or vulgarity. Although the legal restriction of pornography and hate speech is intolerable to American civil libertarians and, according to the United States Supreme Court, inconsistent with the First Amendment, legislation restricting pornography and hate speech is common in other liberal democracies. It is an option open to the liberal to extend such restriction to merely offensive speech.

However, liberals have been reluctant to endorse the legal restriction of merely offensive speech. Even Joel Feinberg (1973), who is friendly to the idea of legally restricting offensive conduct, demurs when it comes to speech, arguing that the value of free expression is so great that expression of opinions should never be restricted on the grounds that the opinion expressed is offensive, although he allows for restriction on the grounds that the opinion is expressed in an offensive (e.g. vulgar) way (pp. 131-3). I said that I aim to show how liberals can respond to the problem of offensive speech using the familiar liberal concept of a "marketplace of ideas." The response I will articulate is different from those mentioned so far. It is an alternative to the societal regulation of offensive speech, with which many liberals are uncomfortable, but it takes the problem of offensive speech seriously, unlike more dismissive liberal responses.

3 Markets and trust

I am going to argue (§4) that intellectual trust among its participants is conducive to the functioning of a marketplace of ideas. Here, I will argue that certain kinds of trust among its participants are conducive to the functioning of a market.[10] When a market functions, in the present sense, high-quality goods are

[7] Cf. Waldron 2012.
[8] Cf. Lukianoff and Haidt 2013.
[9] Cf. MacKinnon 1993, Waldron 2012.
[10] I draw no distinction between trust and reliance (cf. Baier 1986, Holton 1994, Jones 1996); you might think that what I say here (and in §4) is plausible only if reliance is swapped for trust.

produced as a result of the presence of a competition-based incentive to produce such goods (cf. §1). From the consequentialist liberal's perspective, this is what markets are for – their function is to positively affect the quality of the goods being offered for sale, using a particular mechanism. When such "market forces" result in high-quality goods being produced, a market is functioning. To say that something is conducive to the functioning of a market, then, is just to say that, other things being equal, markets function better if that condition is satisfied and worse if that condition is not satisfied.

Note well: the claim that something is conducive to the functioning of a market is not inherently normative. When I say that certain kinds of trust are conducive to the functioning of a market, I am not saying that participants in a market *ought* to trust each other. All I am saying is that markets function better when such trust is present. And there is no reason to assume that we must always do whatever makes markets function better.

It is easy to overlook the fact that various things other than the mere establishment of the principles of free trade (§1) are conducive to the functioning of a market. Consider security. If it's dangerous to be on the streets, people may just buy their bread at the bakery closest to their home, and there will be no genuine competition among the bakeries. My claim is that certain kinds of trust among its participants is also conductive to the functioning of a market. First, I'll argue that participants' manifestation of certain dispositions is conducive to the functioning of a market. Second, I'll argue that participants' trusting each other to manifest said dispositions is conductive to the functioning of a market.

First, that participants' manifestation of certain dispositions is conducive to the functioning of a market. Consider, for example, consistency. Market forces will be diminished if bread-sellers can't predict what their would-be customers will want tomorrow (e.g. because they are too capricious in their preferences) or if bread-buyers can't predict what the bread being sold will taste like when they get home (e.g. because the bakers are inconsistent in their methods). Thus, sellers and buyers' manifestation of dispositions for consistency in their respective products and preferences is conductive to the functioning of a market. Consider, for another example, honesty. Market forces will be diminished if bread-sellers are deceived by their customers about what they want (e.g. they say they want healthy wheat bread, but are inclined to buy more tasty white bread) or if bread-buyers are deceived by their baker about the bread they buy (e.g. they fraudulently sell non-organic bread as organic). Thus, sellers and buyers' manifestation of dispositions for honesty about their respective products and preferences is conductive to the functioning of a market. Consider, for a final example, non-discrimination. Market forces will be diminished if bread-sellers discriminate against certain customers (e.g. they refuse to sell to gay people) or if bread-buyers discriminate against certain bakeries (e.g. they will only shop at the bakery owned by a fellow Presbyterian). Thus, sellers and buyers' manifestation of dispositions for non-discrimination against one another is conductive to the functioning of a market.

Second, that participants' trusting each other to manifest said dispositions is conductive to the functioning of a market. It is good when sellers and buyers are consistent, but far better when they also trust each other to be consistent. If all goes well, bakers can expect the preferences of the bread-buying public to remain more or less constant, and can base their commercial decisions on that expectation, and their customers can expect the bread they buy tomorrow to be more or less like the bread they buy today, and can base their commercial decisions on that expectation. The same, mutatis mutandis, when it comes to honesty and non-discrimination. It is good when sellers and buyers are honest and non-discriminatory,

but far better when they also trust each other to be honest and non-discriminatory. This explains why the mere *appearance* of inconsistency, dishonesty, or discrimination can diminish market forces. Even if all the bakeries would be happy to have me as a customer, if I *think* that only the one will serve a Presbyterian, then there will be no competition among them for my business. Even if my would-be customers are normal people with relatively stable bread preferences, if they *seem* to be lunatics whose taste in bread may change dramatically from day to day, I will be unable to even attempt to compete for their business by baking bread that I think will please them.

Recall that the claim that something is conducive to the functioning of a market is not inherently normative. This jibes with my conclusions. It may not be conducive to the functioning of the bread market that I choose to shop only at the Presbyterian bakery, but so what? I am under no obligation to promote the functioning of the bread market; at best, it could be argued, I have some pro tanto reason to do so, given the value of a functioning bread market. More important, this applies equally in the case of the kind of trust described above. I argued that certain kinds of trust among its participants are conducive to the functioning of a market. This does not imply that anyone is ever obliged to trust anyone. Crucially, it does not imply that you are ever obliged to trust someone who is not trustworthy, e.g. someone not disposed to consistency, honesty, and non-discrimination, nor does it imply that you are ever obliged to trust someone who appears untrustworthy, e.g. someone with a track-record of inconsistency, dishonesty, or discrimination. The claim that certain kinds of trust among its participants are conducive to the functioning of a market is like the claim that certain kinds of trust among friends are conducive to the health of their friendship. This does not imply, crucially, that you are obliged to trust an untrustworthy friend, or one with a track-record of betrayal.

4 Intellectual trust

I argued (§3) that certain kinds of trust among its participants are conducive to the functioning of a market. Now, I shall argue that intellectual trust among its participants is conductive to the functioning of a marketplace of ideas. When a marketplace of ideas functions, high-quality ideas are produced as a result of the presence of a competition-based incentive to do so (cf. §1). For the consequentialist liberal, this is what marketplaces of ideas are for. When a vigorous and earnest contest of ideas results in good ideas being articulated, the marketplace of ideas is functioning as it is supposed to function. To say that something is conducive to the functioning of a marketplace of ideas, then, is just to say that, other things being equal, marketplaces of ideas function better if that condition is satisfied and worse if that condition is not satisfied.

Just as the claim that something is conducive to the functioning of a market is not inherently normative, the claim that something is conducive to the functioning of a marketplace of ideas is not inherently normative. This is important, for, while what I have to offer here is clearly a kind of defense of intellectual trust, I do not want to suggest that anyone is obliged to intellectually trust anyone.

In epistemology, "intellectual trust" is sometimes used to refer to a disposition to believe on the basis of a source, e.g. a disposition to believe on the basis of someone's testimony or to form beliefs using some faculty or method.[11] Here, I shall use the term a bit differently. To trust someone *intellectually* is to trust

[11] Cf. Foley 2001, Zagzebski 2012.

them to manifest certain intellectually virtuous dispositions. I shall proceed as before (§3) by arguing, first, that participants' manifesting these dispositions is conducive to the functioning of a marketplace of ideas and, second, that participants' trusting each other to manifest these dispositions is conducive to the functioning of a marketplace of ideas.

The dispositions I have in mind are not unlike the dispositions I discussed in connection with markets (§3). They include *open-mindedness*, i.e. a disposition to take alternative ideas seriously; *fair-mindedness*, i.e. a disposition to engage charitably and respectfully with interlocutors; *intellectual honesty*, i.e. a disposition to tell the truth and nothing but the truth about your ideas, evidence, and arguments; and *impartiality*, i.e. freedom from bias and prejudice. When these virtues are lacking among the participants of a marketplace of ideas it will tend not to function.

Consider the damage done by intellectually vicious speakers and hearers – consider a marketplace of ideas in which speaker's articulations of ideas are met with dogmatic dismissal, uncharitable interpretation, insincere objections, or bigoted refusal to engage and in which hearer's critical questions, objections, and articulation of alternative points of view are met with the same. And consider one of Mill's arguments (§1): that a censored false opinion "may, and very commonly does, contain a portion of truth." The idea here is that discussion between partisans of alternative positions on an issue can, if all goes well, lead one or both parties to revise their positions in the direction of the truth, by correcting error, incompleteness, or misunderstanding in their original position as a result of insights and lines of reasoning articulated by their interlocutor. But this can obviously only happen if both parties are more or less intellectually virtuous in the ways described above. The same, mutatis mutandis, when it comes to the other alleged benefits of a marketplace of ideas.

However, as above (§3), although it is good for participants to manifest these dispositions, it is far better for participants also to trust each other to manifest these dispositions. Like a bread market in which bread-buyers are inconsistent, dishonest, or discriminatory, the marketplace of ideas just described will not easily create a competition-based incentive for speakers to produce high-quality ideas. Consider a political journalist working on an analysis of some recent event. What reason do they have to create a carefully-researched, fact-checked, well-argued commentary, when what they publish is likely to be ignored or rejected with prejudice? They have none – or, rather, they have much less reason in the present case than they would have were their audience open-minded, fair-minded, intellectually honest, and impartial. And the same applies when it comes to hearers. What reason have you to engage with my ideas – to listen to my arguments, articulate criticisms, or offer your own take – if I am likely to ignore or dismiss with prejudice what you say? A vigorous and earnest contest of ideas requires those on both sides of the contest to have an incentive to engage. In the same way that inconsistency, dishonesty, and discrimination undermine market competition, the intellectual vices corresponding to the aforementioned intellectual virtues undermine competition in a marketplace of ideas.

Now, in this discussion I clearly have in mind, as the paradigm case of a marketplace of ideas, a marketplace of ideas in which the participants have at least some relatively strong preference for true and justified ideas over false and unjustified ideas – i.e. the kind of marketplace that Mill imagined. However, intellectual trust is conductive to the functioning of a marketplace of ideas, even when the participants are motivated less by truth and evidence and more by, say, the sensational. Prejudice and dogmatism will tend to undermine competition to produce sensational ideas just as much as it will tend to

undermine competition to produce true ideas. For a marketplace of ideas to function, ideas need to be given a fair hearing, whether this means sincere and unbiased consideration as to whether they are true or sincere and unbiased consideration as to whether they are sensational.

What we saw, above, in connection with markets (§3), is that a functioning market is not a lawless "state of nature," but rather a socially- and politically-constructed institution requiring a kind of cooperation among buyers and sellers. What I am now arguing is that a functioning marketplace of ideas is not a lawless "state of nature" either, not an unstructured space of hostility and enmity, but a collective enterprise, requiring a kind of cooperation. This is borne out in the way that social epistemologists describe contexts of knowledge generation and transfer. Miranda Fricker (2007) describes testimonial exchange as having "the special flavour of situations in which human beings treat each other as subjects with a common purpose," involving "respect" and inclusion in a "community" of "trusted" and active participants. (p. 132) The same, I want to say, of a marketplace of ideas. As José Medina (2013) argues, open-mindedness goes "beyond the testimonial," and covers other intellectual activities, e.g. "a philosophy debate, the brainstorming of scientists in formulating hypotheses, a discussion with friends about a movie or a novel, etc." (p. 79; cf. Fricker 2007, p. 50n, p. 60)

Cases in which the perspectives of marginalized people differ from the received narrative are illustrative here. For example, marginalized people will have no (motivating) reason to articulate their experiences if they expect that their privileged interlocutors will not listen, and privileged people will have no (motivating) reason to listen to such articulation if they believe that their marginalized interlocutors are being insincere. This is, in contemporary American political discourse, a familiar dynamic. It would be different if the marginalized trusted the privileged to be open-minded and the privileged trusted the marginalized to be honest. However, we must bear in mind that the claim that intellectual trust among its participants is conductive to the functioning of a marketplace of ideas is not inherently normative. What I have said does not imply that the privileged ought to trust the marginalized and the marginalized ought to trust the privileged. Rather, what I have said implies that there will not likely be a productive and beneficial dialogue between privileged and marginalized people – between *any* people – in the absence of the kinds of intellectual trust described here.

5 Offensive speech

I argued that intellectual trust among its participants is conductive to the functioning of a marketplace of ideas (§4). What has this to do with the problem of offensive speech (§2)?

I shall argue that when intellectual trust, of the sort described (§4), is present, the bad consequences of offensive speech are mitigated. In short, offensive speech is easier to cope with under conditions of mutual intellectual trust. I said that offensive speech is unpleasant, annoying, upsetting, or painful (§2). However, it is less unpleasant, annoying, upsetting, and painful when it comes from someone with whom you enjoy mutual intellectual trust, because contexts of mutual intellectual trust are conductive to engagement with offensive speech. Let me explain.

I think the very idea that offensive speech is unpleasant, annoying, upsetting, or painful sounds right primarily when we think of offensive speech coming from someone we do not trust intellectually. We are distinctively irked, I think, by *anonymous*, or nearly anonymous, offensive speech, where the possibility

of intellectual trust seems to be precluded. (Hate speech – the burning cross, the vandalized home, the slur shouted to a stranger from a passing car – is paradigmatically, although not always, anonymous.) When the speaker is a complete stranger, about whom we know nothing other than that they have come out with such-and-such offensive utterance – the random social media poster, the microaggressor about whom we read in the newspaper, the protester we see shouting on TV – there is no possibility of a vigorous and earnest contest of ideas – and, I want to suggest, that is a big part of what makes such situations unpleasant, annoying, upsetting, or painful. Another distinctive kind of case is the case of offensive speech coming from someone we positively *distrust* intellectually, like a familiar politician, media personality, or celebrity. Here, too, there is no possibility of a vigorous and earnest contest of ideas, and we are left only with the anger that we associate with offensive speech.

However, because we tend to think of cases of offensive speech in contexts in which mutual intellectual trust is absent, we can miss the fact that offensive speech can occur in contexts in which mutual intellectual trust is present.

Consider what it feels like when a *friend* says something offensive, as compared to what it feels like when an anonymous stranger or a hated talking head says something offensive. Offensive speech coming from a friend, in this kind of case, is more *surprising* than it is infuriating. We want to know *why* our friend said what they said. If a friend of mine tells me that they don't think we should say "black lives matter" because "all lives matter," I am potentially in a position to ask them why they think that and to tell them why I disagree. When I am in such a position, the crucial reason is that I trust my friend to be open-minded, fair-minded, intellectually honest, and impartial. Or, more precisely, I trust them to be *more-or-less* virtuous in these respects, to be *relatively* virtuous in these respects – and, to the extent that I do, I am in a position to engage with their offensive speech. Unlike in the case of the anonymous stranger or the hated talking head, I think there is a chance I can convince my friend to change their mind. When this happens, I expect that, whatever I say, my friend will respond with dialogue rather than with vitriol. I do not suspect that my friend is just saying "all lives matter" to win votes or get clicks. I do not think that they have arrived at their views because of explicit racism. Given all this, a vigorous and earnest contest of ideas is possible.

Now, a friend is someone whom you trust not only intellectually, but in general. And, more important, a friend is someone whom you not only trust, but *love*. Thus, we cannot draw any conclusion about coping with offensive speech under conditions of mutual intellectual trust from the present case. At best we might conclude that offensive speech is easier to cope with under conditions of mutual affection and amity – which is not surprising. However, we can strip the features of friendship away and the central point remains plausible: when you intellectually trust someone, it is less angering and more surprising when they say something offensive. I have in mind someone with whom your relationship is primarily intellectual – think of the relationships between teachers and students, academic collaborators, or authors and their readers – and whose virtuous intellectual dispositions are familiar to you. When you trust someone to be open-minded, fair-minded, intellectually honest, and impartial, even if they are not your friend, you are in a position to respond to their offensive speech with questions, challenges, criticism, and discussion. Because of that, I am arguing, offensive speech is easier to cope with than it is when such intellectual trust is absent.

Medina describes the intellectual vice of “meta-blindness,” which is centrally an insensitivity to your own intellectual limitations, as involving “a failure to relate to others affectively” and a “lack of empathy.” (p. 81-2) The case of offensive speech in the context of friendship is illustrative of the way in which our emotional dispositions towards each other are crucial in determining whether and to what extent offensive speech has its characteristic negative consequences. The case of mere mutual intellectual trust – i.e. mutual intellectual trust without friendship – similarly involves emotional dispositions. Empathy, in particular, should be given pride of place in our picture of what mutual intellectual trust looks like, as it typically requires exercises of empathy to trust non-friends to be intellectually virtuous. To be reasonably confident that someone will be more-or-less or relatively open-minded, fair-minded, intellectually honest, and impartial, you have to able to see things from their point of view and to acknowledge their perspective as valuable.

Now, I want to concede: mutual intellectual trust, like friendship, involves emotional vulnerability. Imagine that your prize pupil, whom you have been mentoring all semester, whom you trusted to be open-minded and careful in preparing their final paper, turns in some sloppily-reasoned nonsense in defense of the proposition that “all lives matter.” There is a kind of disappointment, a feeling of being let down, here that would not be present if the author of the bad essay were merely some anonymous pundit. Trust, of any kind, essentially includes the possibility of being betrayed. However, it seems to me that the possibility of engagement that comes with mutual intellectual trust mitigates these risks. Even when our intellectual trust has been betrayed, as in the case of the sloppy essay, so long as it has not been extinguished, there is potential for dialogue and debate – and that provides a benefit to weigh against the cost of betrayal. There are two points that are worth making here. The first is that the negative emotions made possible by trust (e.g. disappointment) can exist at the same time as (relatively) positive emotions made possible by the presence of trust (e.g. surprise, as opposed to outrage). The second is that the possibility of engagement, in contexts of mutual intellectual trust, means that it is possible for a betrayal to be redeemed through conversation, as when you rationally persuade a trusted interlocutor to abandon some offensive view or are rationally persuaded by them that their view, properly understood, was not offensive in the first place. No such redemption seems possible in the case of the aimless outrage we experience when we confront offensive speech in the absence of mutual intellectual trust. For these reasons, I think we can still say that offensive speech is less unpleasant, annoying, upsetting, and painful when it comes from someone with whom you enjoy mutual intellectual trust

6 Conclusion

Where does all of this leave the liberal vis-à-vis the problem of offensive speech (§2)? They are in a position to articulate some modest but substantial progress. The consequentialist articulation of the idea of a marketplace of ideas (§2) leads us to the insight that intellectual trust is conductive to the functioning of a marketplace of ideas (§3), in the same way that certain kinds of trust are conducive to the functioning of a market (§2). The internal logic of liberalism, as it were, leads us to posit the value of intellectual trust. However, mutual intellectual trust among its participants mitigates the bad consequences of offensive speech. Thus, there is a good that liberals are led to embrace, for the sake of the functioning of the marketplace of ideas, that at the same time mitigates the bad consequences of offensive speech. Offensive speech is less of a problem when this distinctively liberal good – mutual intellectual trust – is present.

I offered this argument both as an alternative to ignoring or dismissing the problem of offensive speech and as an alternative to societal regulation of offensive speech (§2). However, I have said noting about *policies* that liberals might pursue in response to the problem of offensive speech. In this respect, my discussion is incomplete, as I have not yet articulated a *practical* alternative to societal regulation of offensive speech. It is easy to say that, if I am right, liberals ought to *promote* intellectual trust among interlocutors for whom offensive speech has become an issue. But what exactly does that mean? There are at least two kinds of cases to consider in which we might say that intellectual trust is lacking and in need of promotion. First, there are cases in which the trustworthy are not trusted, as when someone is wrongly perceived to be insincere or bigoted. In this kind of case, the liberal would like to correct these mistaken perceptions. Second, there are cases in which people are untrustworthy, as when someone really is insincere or bigoted. In this kind of case, the liberal would like to cure these intellectual vices, or perhaps to recommend so-called "therapeutic" trust, aimed at inculcating trustworthiness in the person trusted. In neither kind of case is it clear what concrete social, legal, or institutional *policy* might do the job.[12]

Bibliography

Baier, A.C. (1986), "Trust and Antitrust," *Ethics* 96, pp. 231-60.

Foley, R. (2001), *Intellectual Trust in Oneself and Others* (Cambridge University Press).

Holton, R. (1994), "Deciding to Trust, Coming to Believe," *Australasian Journal of Philosophy*, 72(1), pp. 63-76.

Jones, K. (1996), "Trust as an Affective Attitude," *Ethics* 107, pp. 4-25.

Feinberg, J. (1973), ""Harmless Immoralities" and Offensive Nuisances," in N.S. Care and T.K. Trelogan (eds.), *Issues in Law and Morality: Proceedings of the 1971 Oberlin Colloquium in Philosophy* (Case Western Reserve University Press).

Feinberg, J. (1984), *Harm to Others: The Moral Limits of the Criminal Law, Volume One* (Oxford University Press).

Feinberg, J. (1985), *Offense to Others: The Moral Limits of the Criminal Law, Volume Two* (Oxford University Press).

Fricker, M. (2007), *Epistemic Injustice: Power and the Ethics of Knowing* (Oxford University Press).

Langton, R. (1993), "Speech Acts and Unspeakable Acts," *Philosophy and Public Affairs* 22(4), pp. 293-330.

[12] I presented versions of this paper at a workshop at the Collège de France in 2017 and at the World Congress of Philosophy, in Beijing, in 2018. For valuable feedback, I would like to thank my audiences at those talks, the students in my course on free speech, at Washington University in St. Louis, in 2019, and Alessandra Tanesini.

Langton, R., and West, C. (1999), “Scorekeeping in a Pornographic Language Game,” *Australasian Journal of Philosophy* 77(3), pp. 303-19.

Lukianoff, G., and Haidt, J. (2015), “The Coddling of the American Mind,” *The Atlantic*, September 2015.

MacKinnon, C.A. (1987), “Francis Biddle’s Sister,” in *Feminism Unmodified: Discourses on Life and Law* (Harvard University Press), pp. 163-97.

MacKinnon, C.A. (1993), *Only Words* (Harvard University Press).

Maitra, I. (2009), “Silencing Speech,” *Canadian Journal of Philosophy* 39(2), pp. 309-38.

Maitra, I. (2012), “Subordinating Speech,” in I. Maitra and M.K. McGowan (eds.), *Speech and Harm: Controversies over Free Speech* (Oxford University Press), pp. 94-120.

Medina, J. (2013), *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and the Social Imagination* (Oxford University Press).

Mill, J.S. (1859/1989), *On Liberty and Other Writings* (Cambridge University Press).

Waldron, J. (2012), *The Harm in Hate Speech* (Harvard University Press).

Zagzebski, L.T. (2012), *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief* (Oxford University Press).