CHAPTER 9

Cleaning Cyber-Cesspools: *Google and Free Speech*

BRIAN LEITER

I shall use the term "cyber-cesspool" to refer to those places in cyberspace—chat rooms, websites, blogs, and often the comment sections of blogs[1]—which are devoted in whole or in part to demeaning, harassing, and humiliating individuals: in short, to violating their "dignity." Privacy is one component of dignity—thus its invasions represent an attack on dignity. But they are not the only such affront: implied threats of physical or sexual violence also violate dignity; so too non-defamatory lies and half-truths about someone's behavior and personality, so too especially demeaning and insulting language, so too tortious defamation and infliction of emotional distress. Cyber-cesspools are thus an amalgamation of what I will call "tortious harms" (harms giving rise to causes of action for torts such as defamation and infliction of emotional distress) and "dignitary harms," harms to individuals that are real enough to those affected and recognized by ordinary standards of decency, though not generally actionable.

The Internet is currently full of cyber-cesspools. For private individuals without substantial resources, current law provides almost no effective remedies for tortious harms, and none at all for dignitary harms. Dignitary harms are off-limits for legal remedy because U.S. constitutional law effectively subordinates the dignity of persons to a particular conception of liberty. Speech, however, causes real harms (dignitary and otherwise), so much so that the only reason to think government ought not protect against such harms is that government actors have too many obvious incentives to overreach in placing restrictions on speech.[2]

Since cyber-cesspools are in large part beyond the reach of regulation by the state in America because of constitutional protections, a number of commentators[3] have suggested enhancing private remedies by, for example, making intermediaries—those who host blogs or perhaps even service

providers—liable for tortious harms on their sites. This would require repeal of Section 230 of the Communications Decency Act (47 U.S.C. §230), which provides that "No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider." The effect of that simple provision has been to treat cyber-cesspools wholly differently from, for example, newspapers that decide to publish similar material. Whereas publishers of the latter are liable for the tortious letters or advertisements they publish, owners of cyber-cesspools are held legally unaccountable for even the most noxious material on their sites, even when put on notice as to its potentially tortious nature. But why should blogs, whose circulation sometimes dwarfs that of many newspapers, be insulated from liability for actionable material they permit on their site?[4] Although it is common for cyber libertarians to talk as if *all* speech is immune from legal regulation, even U.S. constitutional law permits the law to impose penalties for various kinds of "low-value" speech, such as defamation. So why should the law, via Section 230, treat cyberspace differently than the traditional media?

Defenders of Section 230 worry about what I shall refer to as "spillover effects": because website owners are more likely to err on the side of caution when facing legal liability, so the argument goes, if they do not have Section 230 immunity, they will be more likely to "censor" speech, including "valuable" speech. This is probably true, but it has a flip side: namely that insulation from liability via Section 230 will increase the prevalence of low-value speech, as well as speech that causes dignitary harms, as anyone familiar with cyberspace can attest. Why think the balance should be struck in one direction rather than the other? In all kinds of contexts—newspapers, classrooms, workplaces, and courtrooms—we restrict speech not only for the sake of legally protected interests but also for the sake of avoiding dignitary harms, no doubt at the cost of spillover effects. If no academic institution or newspaper would permit its classrooms or pages to turn into the analogue of cyber-cesspools, why should the law encourage that outcome in the virtual world?

The harm of speech in cyberspace is sufficiently serious that we should rethink the legal protections afforded cyber speech that causes dignitary harms. Thanks to Google (and similar search engines), cyber speech tends to be (1) permament, (2) divorced from context, and (3) available to anyone. If the law should not remedy this problem, it must be because the value of speech that inflicts dignitary harms or the value of the speech swept up in

the spillover effects is such that legal regulation is not justified. As I argue below, it is not clear whether either case can be made.

Let us first begin, however, with some case studies. I will quote *verbatim*, because too often academic discussion of this topic whitewashes what is really going on in the cyber-cesspools. Those easily offended—even those not so easily offended—are duly warned.

A Tale of Two Cyber-Cesspools

In late 2004, I noticed that my blog[5] was getting hits from what purported to be a pre-law chat room called "AutoAdmit." I followed back some of the links, and so discovered a website that I have since described—generously, I might add—as a "cesspool of infantile morons, racists, and misogynistic freaks." Especially alarming was the fact that, while about half the "threads" in the chat room actually had something to do with law school or the practice of law—suggesting that there were actual law students utilizing this board—another half had as their primary purpose racist, misogynistic, and anti-Semitic abuse or simply vicious harassment, defamation, and implied threats against named individuals, usually other law students.

Oddly, there was no indication who was responsible for the AutoAdmit site, since it was devoid of contact information. In March 2005, after watching AutoAdmit for several months, I wrote a short note about AutoAdmit to Eugene Volokh, proprietor of a well-known right-wing law blog (prompted by a related post on his blog), knowing that many law students read his blog and thinking he might help "shame" the still-anonymous proprietors of the site into cleaning it up. To my surprise, Professor Volokh's posting about the site[6] led the "administrators" to "out" themselves the following day! Professor Volokh posted a response[7] signed by Anthony Ciolli, then a law student at the University of Pennsylvania, and Jarret Cohen, an insurance salesman in Allentown, Pennsylvania, defending the huge amount of racism, sexism, and anti-Semitism on the site on the grounds that,

> We are very strong believers in the freedom of expression and the marketplace of ideas. This is why we allow off-topic discussion and almost never censor content, no matter how abhorrent it may be.

We shall return to the contention that "the freedom of expression and the marketplace of ideas" are in any way relevant to the existence of cyber-cesspools. In any case, now that Mr. Ciolli had outed himself, I made my

any affiliation with it. You, instead, facilitated the expression and publication of such language. . . .

The increased attention to the plight of the two Yale women, alas, brought a proliferation of new abusive threads directed at them on AutoAdmit. The female victims eventually sued. They were represented pro bono by one of the nation's leading firms and a leading cyber-law expert. After two years, a handful of the harassers had been identified, some settled, and much of the abusive content had been removed from Google. In short, it took legal representation costing, one surmises, hundreds of thousands of dollars already, to achieve some modest success against only a handful of anonymous misogynists on one cyber-cesspool.

Our second cyber-cesspool is less colorful, and the harm it inflicted trivial by comparison. Yet, by way of contrast, it will prove useful for our analysis.

I run a blog that is widely read by philosophers. An obscure, and quite right-wing, philosophy professor at a university in the American South wrote to me asking that I link to his own blog. I had not been in the habit of simply linking to other blogs on request, but I did look at his blog, and it seemed a bit peculiar and not very interesting. I did not link to it. The obscure philosophy professor—I will call him K.[12]—had been writing nice things about my blog, but when I failed to link to him after a couple of months, he became angry. He started attacking me on his blog: he wrote a hundred different items attacking me as a "disgrace," a "buffoon," a "nut," and the like. The attacks became rather vicious and personal. He began making up incidents and hurling wild accusations. He linked to a story about a left-leaning professor in a small Texas town whose home was vandalized and who received other kinds of threats of violence. He suggested that this should be a warning to me (I lived in Texas at the time). My dean reported this incitement to violence to the university system in which K. worked. He calmed down for a while, but then resumed his irrational attacks.

I finally responded on my blog, documenting K.'s history of vicious and irrational attacks on me and others. I did not realize at the time that K. was probably mentally ill, and that my response would cause him to crack. Not quite two months later, K. spent his Christmas day creating a separate blog devoted to insulting, defaming, harassing, and threatening me—*and* my wife, my children, my parents, and anyone who reminded him of me! He wrote several hundred posts to this effect over the next two years, declaring,

more than once, that "By the time we're done with this sorry excuse for a human being, he'll be crying," and noting that my having responded at all to his attacks constituted "a terrible mistake, one that will haunt [Leiter] for the rest of his life. At 48 years of age and in great health, I expect to be around for another 25 to 30 years." This adult man, a tenured professor of philosophy, would often link to photographs of me in order to mock my appearance as "effeminate," also pointing out how "lucky" my children were that "President Bush protected them, for their father certainly wouldn't have." He also compared me to Stalin and Hitler, and denounced me as, variously, an "imbecile," a "monster," and a "cretin." He declared that he would humiliate me in front of not only my children, but my grandchildren![13] Soon I began to hear from other academics he had harassed. The whole display was sufficiently strange that my dean asked a psychiatrist to review the blog to try to shed some light on K.'s mental disturbance and whether he was dangerous.

Economists say, correctly, that the "barriers to entry" are low in cyberspace. They are thinking mainly of financial cost, but the barriers are "low" in a more significant way as well. Prior to cyberspace, if you wanted to reach more than your immediate circle of acquaintances, you usually had to have some kind of competence, education, status, intelligence, and ability: otherwise no one would listen to or publish you. Indeed, in the old days, you generally had to be moderately sane to get an audience! That is no more. The K.s of yesteryear were confined to writing letters, or ranting to their friends (if they had any), or handing out leaflets, or doing other things that involve personal contact, the kind of contact that would, in most instances, reveal the profound level of emotional disturbance afflicting the speaker. Now the K.s and the AutoAdmit sociopaths need only a computer in order to abuse their targets, and to do so in a way that permits their defamation and harassment to be visited and revisited again and again by countless people anywhere on the planet, visitors who are often deprived of almost all relevant information about the speaker or his targets.

Google Is Part of the Problem

This brings us to the final villain in our story of cyber-cesspools: Google, the dominant search engine in the market today. (I shall refer in what follows just to "Google," but the same points apply mutatis mutandis to other search engines.) For without Google, every K. and every chat-room sociopath stews

in obscurity. It is Google that retrieves the rantings of a friendless madman typing away on his hate blog, or the anonymous smears directed at a female law student by a vicious misogynist in a chat room, and associates those rantings and smears with the victim's name for any Google user to find. How exactly Google decides what search results to return is shrouded in a bit of mystery, though they do reveal[14] using a "PageRank algorithm" that "considers the importance of each page that casts a vote [by linking], as votes from some pages are considered to have greater value, thus giving the linked page greater value" and a "hypertext-matching analysis," which looks at how a particular search term (e.g., "Brian Leiter") figures in a web page's content.

The idea that the "most relevant and reliable results" about a female student at Yale Law School consist of the anonymous rantings of misogynistic sociopaths would be amusing if real people were not involved. Why Google searches give such prominence to blogs and Internet chat rooms—which, as a class, may be among the least reliable sources of information in human history—is puzzling. As the historian and blogger Juan Cole delicately puts the point: Google does not necessarily "put[the most relevant and reliable results first," though it most certainly facilitates what Professor Cole aptly dubs "the Google smear": the discrediting of an opponent by abuse that has a high web profile and is indexed through Google.[15] When the AutoAdmit posters attacked the Does, they were attacking private individuals with hardly any Internet presence at all. But when posted on a highly trafficked site, AutoAdmit, with the names of the victims in the thread titles, the attacks very quickly became the top search results for anyone—a friend, a family member, a prospective employer, a new acquaintance—Googling their names. Google "smearing" someone like me, with a substantial Internet presence (in the form of blogs and university home pages) is a bit more challenging. K., however, had a cyber friend, another far-right racist blogger, also of dubious sanity. This blogger emailed his entire circle of far-right blogging friends to advertise K.'s hate blog, telling them in the process a series of bizarre falsehoods about me and concluding: "[P]lease blogroll the new blog so that it rises in Google's rankings, so that when people type 'Brian Leiter' into Google, the new blog comes up." His network of extremists obliged, and a new Google smear was briefly born.

What, if anything, should Google do about its clear complicity in the viability of cyber-cesspools? What, if anything, should the law do about it? Is the moral value of "free speech" an obstacle? To these issues we now turn.

The Value of "Free Speech"

All young children are advised at some point to remember that "Sticks and stones can break your bones, but names can never hurt you." Like many things told to young children, this isn't true. Indeed, on its face, the advice is a non sequitur: there are harms other than broken bones, and there is no reason at all to think that "names," that is, words, are not capable of causing them. To be sure, "names" do not break bones, but humans are creatures whose lives are suffused in meaning, and these meanings constitute their sense of self and large parts of their well-being. Words may not be the unmediated cause of a fracture, but they can certainly cause humiliation, depression, debilitating anxiety, incapacitating self-doubt, and devastating fear about loss of safety, respect, and privacy. There are three standard rationales offered for permitting speech, even when it causes some harm: individual autonomy, democratic self-governance, and the discovery of the truth ("the marketplace of ideas"). I will assume that something like John Stuart Mill's "Harm Principle" should be a limitation on individual liberty, and that certain degrees of harm can override the value of speech. I will also assume (contra, perhaps, Mill) that "harms" can include psychological ones—such as dignitary harms to reputation and privacy interests, as well as tortious harms that our law does recognize.[16] Since no one contests the propriety of regulating tortious harms, I concentrate on speech that causes dignitary harms, as well as speech that is included in the spillover effects of more effective regulation of tortious harms through the abolition of Section 230 immunity for website owners. The question, in short, is what value the speech on cyber-cesspools can be said to have. If there is any legally significant difference between the virtual and actual worlds, it is that speech in the virtual world may be more likely to cause harms because of its ability to reach a wide audience stripped of relevant context thanks, in large part, to Google.

Notice, to start, that cyber-cesspools, at least insofar as they target private individuals, will get no help from considerations of democratic self-governance:[17] the viability of informed democratic decision making is not at stake when an anonymous poster on AutoAdmit reports that Jane Doe has herpes or that he would like to sodomize her forcibly. That means that if there is a reason not to regulate the kind of abusive speech that is the hallmark of cyber-cesspools it must come from the other two considerations: individual autonomy and/or the discovery of the truth ("the marketplace of ideas"). Let us consider the "marketplace of ideas" rationale first.

Mill believed that discovering the truth (or believing what is true *in the right kind of way*) contributes to overall utility, and that an unregulated "marketplace of ideas" was most likely to secure the discovery of truth (or believing what is true in the right kind of way). Mill's commitment to the so-called "marketplace" is based on three claims about truth and our knowledge of it. First, Mill thinks we are not justified in assuming that we are infallible: we may be wrong, and that is a reason to permit dissident opinions, which may well be true. Second, even to the extent our beliefs are partially true, we are more likely to appreciate the whole truth to the extent we are exposed to different beliefs that, themselves, may capture other parts of the truth. Third, and finally, even to the extent our present beliefs are *wholly* true, we are more likely to hold them *for the right kinds of reasons*, and thus more reliably, to the extent we must confront other opinions, even those that are false.

For this line of argument to justify a type of speech, the speech in question must be related to the truth or our knowledge of it, and discovering this kind of truth must actually help us maximize utility. Now one might wonder whether some of the purported "truths" that cyber-cesspools proffer—for example, the purported truth that Jane Doe has herpes—are actually truths that contribute to maximizing utility. But, from the utililtarian perspective, that is not even the right way of framing the question: for the real question is whether claims about Jane Doe's alleged herpes on Internet sites by anonymous individuals with unknown motives (it is even unknown whether they have any interest in the truth!) are likely to maximize utility. It would seem not unreasonable, I venture, to be, at most, agnostic about an affirmative answer to this question, especially once we factor in the likely harms in the event that the claim is false.

But Mill, it is important to recall, did not actually accept the thesis about our fallibility in its strongest form. For Mill held that there is no reason to have a "free market" of ideas and arguments in the case of mathematics (geometry in particular) since "there is nothing at all to be said on the wrong side of the question [in the case of geometry]. The peculiarity of the evidence of mathematical truths is that all of the argument is on one side."[18] This is all the more striking a posture in light of the fact that Mill is a radical empiricist, and so denies that there is any a priori knowledge: even logical and mathematical truths are a posteriori, vindicated by inductive generalizations based on past experience. On Mill's view, then, there simply would not be any epistemic case for making room for the expression

of opinions on which there is no contrary point of view that could make any contribution to the truth. This point is particularly important to bear in mind when it comes to material on cyber-cesspools aimed at private individuals.

Permit me to take what I hope is not a very controversial position, namely, that there actually are *not* two sides to the question of whether Jane Doe ought to be forcibly sodomized. If there are any moral truths, surely all the epistemic bona fides are on just one side of this issue. In other words, the explicit and implied threats of sexual violence central to cyber-cesspools like AutoAdmit simply have no moral standing based on the "marketplace of ideas": they are in the same boat, for any Millian, as a website devoted to establishing that the square of the hypotenuse of a right triangle is equal to the product, rather than the sum, of the squares on the other two sides.

But what of dignitary harms more generally, and what of the spillover effects attendant upon a legal regime in which website owners face intermediary liability? Surely some speech that causes dignitary harms actually *does* facilitate the discovery of the truth, and surely much of the speech that falls within the scope of spillover effects from more effective regulation of tortious harms in cyberspace would do so as well (and some of it might even affect democratic self-governance). If we are to be genuine utilitarians, we must weigh the competing utilities and disutilities of different schemes of regulation of speech. I shall advance two claims: first, dignitary harms are much more harmful in the age of Google; and, second, spillover effects of more effective regulation of tortious harms in cyberspace will have little effect on the discovery of truth or democratic self-government.

The AutoAdmit sociopath no doubt had his analogue in an earlier era: call him the Luddite Sociopath. The Luddite Sociopath could indeed tell his friends and acquaintances that Jane Doe is a "slut" with herpes, but there is little reason to think the law ought to provide redress, except in extreme circumstances. The reasons are worth emphasizing. The Luddite Sociopath, in the first instance, reaches hardly anyone with his hateful message. We cannot control, and would not in any case want the law to control, the *thoughts* of others. People may think whatever they want, however false, foolish, disgusting, or demeaning. Even when the Luddite Sociopath articulates his thoughts, the impact is minimal: a small circle of acquaintances, perhaps, hear it, and some of them, thanks to their familiarity with

the Luddite Sociopath, may appropriately discount them. The harm to Jane Doe is still almost nonexistent: she is insulated both by the size of the audience *and* the availability to the audience of their *experience* with the Luddite Sociopath. Jane Doe may prefer, understandably, that no one think these thoughts or express them, but that is not a preference the law can satisfy.

Suppose, now, that the Luddite Sociopath is dissatisfied with his limited audience, and with the fact that his audience generally knows a fair bit about him—for example, his propensity to rant and rave, or his misogyny, or his inability to interact normally with other people, or his membership in fringe political groups, and the like. The Luddite Sociopath wants the *world at large* to "know" about Jane Doe, he wants to *harm* Jane Doe with his words. Our Luddite Sociopath needs an intermediary who can broadcast his words far beyond any audience he can reach, and who can detach his words, and their meaning, *from him* so that they are free-standing meanings that supply no context for interpretation that might defuse their force.[19] The Luddite Sociopath thus sends letters to the editors of newspapers, tries to place ads in magazines, and tries to weasel his way on to radio and cable television programs that will give him a potent forum for his message about Jane Doe.

But now, of course, the law steps in and places some obstacles in his path. For the law declares that any one of these intermediaries who picks up the Luddite Sociopath's "message" about Jane Doe can be liable for defamation and infliction of emotional distress. None of these intermediaries can say, "We did not say those nasty things, the Luddite Sociopath did!" Thus, the law gives every intermediary a significant incentive to be cautious, to investigate what the Luddite Sociopath says before broadcasting it, and to look into the Luddite Sociopath's background and motivations. Notice, too, that even in the absence of intermediary liability, most of the traditional media also give weight to dignitary harms in deciding what ought to be published about private persons.

In the age of blogs, Internet chat rooms, and Google, our formerly Luddite Sociopath has new intermediaries who have no current incentive to place *any* obstacles in his way. With the help of a chat room or blog, he can disseminate his message about Jane Doe to those who know nothing about him, and with the help of Google, the Sociopath's message can now be widely disseminated well beyond the blog or chat room to anyone with any interest in Jane Doe. Because the law, through Section 230, insulates the

intermediaries from any liability, the law no longer puts any obstacles in the way of the Sociopath: no blog owner, or chat room administrator, or search engine operator, has any legal reason to make it harder for the Sociopath to express his thoughts about Jane Doe, to express them with no contextual information about the Sociopath or his target, and to do so in ways that are no longer ephemeral, but etched into the Internet's permanent memory, thanks to Google, for anyone, anywhere to discover. Both Tortious harms and dignitary harms are, in consequence, *more harmful* than ever before.

As Internet sources gradually displace or replicate the functions of other media, the reasons for thinking that they, unlike their old media counterparts, should be exempt from familiar forms of legal regulation will seem increasingly bizarre. Let us assume, then, that Section 230 will be repealed or significantly modified. Hopefully we shall then see the application of ordinary tort law not to Internet service providers, but to the intermediaries more proximate to the harmful words: for example, blog proprietors and chat room administrators/owners. The result would unquestionably be a significant reduction in the freedom with which individuals, especially anonymous individuals, are able to speak on the Internet. That effect would be enhanced if the law were also to provide remedies for some dignitary harms in cyberspace.

There would, however, be no reduction at all in the ability of individuals to speak freely, just in their ability to exercise that purported right to speak freely in cyberspace. It is important to emphasize *purported,* since, as with Ciolli and Cohen's defense of AutoAdmit, appeals to "free speech" are invoked on behalf of speech that in fact enjoys no special legal or moral standing (e.g., defamation of private individuals). Repeal of Section 230 together with causes of action for some dignitary harms will undoubtedly reduce, dramatically, the number of comments sections on blogs, since most blog proprietors fail to monitor the content on their sites. Why that would be a greater loss in cyberspace than it is in the traditional media, which do not permit nearly as much unregulated anonymous self-expression, is a question I have not seen addressed. Certainly anyone who has spent much time reading anonymous comments on blogs would not conclude that they are an especially notable repository of human wisdom, rational insight, or moral acuity. Indeed, if the entire Internet vanished tomorrow, we would still have all the traditional media and the traditional fora of communication: not just the so-called "mainstream media," but the

alternative newspapers and presses, the foreign newspapers, the libraries, the scholarly periodicals, the satellite radio and the cable television, and on and on.

The issue, though, is not the Internet, but only certain sites on the Internet, like blogs and chat rooms, which are the primary loci of cyber-cesspools. The world is not obviously better because of these parts of the Internet, and in many ways it is obviously worse. Prior to blogs and chat rooms and Google, female law students were not subjected to campaigns of anonymous vicious harassment accessible to thousands of other students and lawyers around the country. Prior to blogs and chat rooms and Google, it was rather harder to irresponsibly invade privacy, circulate defamatory statements, or threaten sexual and criminal violence with seeming impunity. What *precisely* are the contributions to human knowledge and well-being that are attributable *solely* to these aspects of the Internet, that would have been impossible without its existence in its current unregulated form?[20] It is far from obvious that there are any, at least in otherwise democratic societies.

The preceding considerations leave us, it seems, with only one free speech argument for not regulating cyber-cesspools: namely, the value of permitting individuals to express themselves freely. But what exactly is *valuable* about such expressive freedom or autonomy? Consider the idea that the value of autonomy resides not in free choice per se but in choosing wisely or valuably.[21] If *autonomy* or *freedom* per se has value, then we should think it better that Hitler chooses *freely* to kill the Jews of Europe than that he does so because of a chemical imbalance in his brain. But most of us think the opposite: freedom of choice, exercised poorly, has even less value than the same action performed unfreely![22]

The line of thought I am criticizing here trades on an ambiguity about the "value" of an action: between, that is, its blameworthiness (which is increased when one *autonomously* chooses badly) and its utility for the agent. What is really at stake is the idea that an individual is better off when he can "express" himself than if he has to "bottle up" who he is, what he feels, and so on.

There may well be a type of *value* for the agent in his being able to express himself: Hitler feels better, one suspects, if given the opportunity to rant and rave about the Jews. But that fact leaves unanswered key questions. Is Hitler's "feeling better" a relevant criterion of utility? Can his "feeling better" be outweighed by the disutility to others? We should not conceive of utility in terms of preference-satisfaction alone, so that if Hitler's preference is to spew his venom about the Jews, then it creates utility to let him do so. Satisfying many kinds of preferences makes people worse off: the heroin addict's ability to satisfy his preference for more heroin does not add to his well-being

Even if self-expression has utility for the self that gets to express itself—however depraved or ignorant or foolish—we still need to weigh the utility of others. Let us assume the AutoAdmit sociopath gets utility, in the sense of preference-satisfaction, from his ability to express his desire to sodomize Jane Doe. It surely is not plausible that this utility outweighs the harm to Jane Doe of having that message broadcast, repeatedly and widely. But in that case, we no longer have a justification for permitting such speech.

I conclude that there is no clear reason to think that speech about private individuals on cyber-cesspools has any moral standing as free speech that should be protected, and there is no reason to think spillover effects of better regulation of cyber-cesspools will not be offset, many times over, by all the other avenues by which knowledge is shared and opinions expressed, both on the Internet and in the other media of communication. Legal defenses already exist against abuse of legal process, in the form of SLAPP (strategic lawsuit against public participation) suits against meritless defamation actions whose intent is to suppress protected speech. Yet the main prophylactic against such abuse is to restrict remedies against cyber-cesspools to "private" individuals, as understood in American libel law.[23] "Private" individuals, unlike public figures, are less likely to have the resources to mount frivolous assaults on cyber-cesspools and, by the same token, speech about them is less likely to implicate democratic values or truths that really maximize utility. This is, after all, the solution we have preferred in the rest of American law. The real question is why cyberspace should be treated *more protectively* when it comes to tortious harms and why it should not, in fact, be treated more restrictively when it comes to dignitary harms, given how much more harmful they are in cyberspace.

Regulating Google to Reduce Tortious and Dignitary Harms

What *ought* Google to do about its role in facilitating cyber-cesspools? Here are some simple steps an ethical search engine company might take in response to the harms caused by cyber-cesspools:

First, Google could set up a panel of neutral arbitrators who would evaluate claims by *private* individuals that Google is returning search results that might constitute tortious or dignitary harms. I would limit the right of appeal to *private* individuals precisely because speech about public figures is far more likely to implicate actual free speech values such as democratic self-governance. Google might impose a modest fee for this right of appeal, in order to reduce the number of frivolous complaints filed with the panel. But even a modest fee (say, $500) would be miniscule by comparison to the cost to victims of filing a legal action. The Google panel would receive and evaluate whatever materials the complainant deems relevant.

Second, the Google panel would have authority to provide several possible remedies in the event it concurs with the complainant that the material in question is more likely than not to constitute actionable material or a dignitary harm (the panel would apply something akin to a "preponderance of the evidence" standard). Possible remedies might include (1) delisting the material in question from the search engine results or demoting the results so that they turn up after the first page of results (the first page, or the "top five," being the only ones that most search engine users peruse); (2) awarding to the complainant, per a proposal of Frank Pasquale,[24] a "right of reply" in the form of an asterisk attached to the search result that links the searcher to the complainant's response; or (3) requiring the proprietor of the site on which the material in question appears to provide evidence to the Google panel that the material in question is neither actionable nor a dignitary harm; in the event the proprietor fails to do that, either one of the first two remedies would be available.

Imagine how this system would have worked in the case of the AutoAdmit Does. They are private individuals. They could have provided URLs to the offending postings at AutoAdmit, the majority of which, on their face, were actionable. For those not constituting per se libel or obvious inflictions of emotional distress, the Does might have needed to submit some additional information: for example, evidence of actual LSAT scores (which were alleged on AutoAdmit to be extremely low). The Google panel would have, presumably, awarded the delisting remedy, and the whole matter would likely have ended. The anonymous AutoAdmit sociopaths could rant and rave about the Does, but their ranting and raving would be far less likely to reach employers, friends, and relatives, and so would be far less harmful.

On this proposed voluntary scheme, it is less clear whether a quasi-public figure like me would be helped. The likelihood that highly critical speech about a public or quasi-public figure actually has some *value* is prima facie higher; one cannot craft legal rules around freak cases, like mentally ill individuals with delusional obsessions. By the same token, the ability of what is really actionable speech to do damage to a public or quasi-public figure is significantly less, precisely because there is so much other information available.

Google is unlikely to adopt this voluntary scheme, so the law will probably have to create incentives for Google to address its role in the proliferation of cyber-cesspools. The most promising analogue from existing law would make Google liable for its negligence in disseminating tortious material. (A more radical proposal would make Google liable for disseminating material constituting dignitary harms as well; I remain agnostic on whether that would be advisable.) To be sure, under existing law, neutral disseminators of even actionable material are rarely deemed liable, except in cases where they are aware of the tortious nature of the material, the harms are serious and highly probable, and the burdens on the disseminator to deflect the harm are not too great. The Dobbs torts treatise takes the view,[25] for example, that libraries are not likely to be held liable for maintaining on their shelves material that they know to be defamatory. Yet one of the state statutes Dobbs cites, Cal. Civ. Code §48.5, actually does impose liability on a radio broadcaster who fails to exercise "due care" in disseminating tortious material.[26] This is, I think, the right paradigm for the treatment of Google.

Here is the proposal: if Google is put on notice by a private party complainant that material returned in its searches is tortious, it has an obligation to evaluate the claim in accordance with something like the procedure sketched above; if it fails to evaluate the soundness of the claim, then it is liable for negligent dissemination of tortious material, and can be sued by the complainant. If it undertakes a fair review of the complainant's claim about the material, and deems it nonactionable, then it is not liable, unless the complainant can establish negligence in the review process. Google is, after all, a formidable defendant to sue, and one can imagine that will deter many private individuals. On the other hand, if Google, even after being put on notice, continues to disseminate actionable material, why should it not face liability? We already impose similar obligations on Internet intermediaries with respect to copyright infringement,[27] why not accord as much

protection for private individuals from tortious harms (and maybe even dignitary ones)?

Conclusion

The rhetoric about "free speech" in cyberspace usually obscures more than it illuminates, even in scholarly discussions. Daniel Solove observes that, with respect to the regulation of speech on the Internet, "We are witnessing a clash between privacy and free speech, a conflict between two important values that are essential to our autonomy, self-development, freedom and democracy."[28] Yet when it comes to cyber-cesspools, most of these values are not implicated at all, except on the side of the victims. Only the incredible view that all expression, regardless of its subject or character, has value could sustain the idea that there is a significant clash here. James Grimmelman writes that "[R]emoving content from a search engine's index at the demand of a third party . . . is offensive to free-speech values."[29] His example is China's demand that search engines block users from finding information about the Falun Gong, a quasi-religious movement banned by the Chinese government. That might, indeed, implicate some free speech values—such as democratic self-governance and the discovery of the truth—but removing AutoAdmit content threatening sexual violence from a search engine implicates no free speech values I can discern. Mark Lemley—who represented pro bono the women suing AutoAdmit—claims that "The amazing diversity of the Internet, with its abundance of user-generated content, would be impossible" without some safe harbors for intermediaries from liability.[30] But is there really an amazing diversity of *valuable speech* in and around cyber-cesspools such that we should give them safe harbors? We do not protect safe harbors in the traditional media for "cesspool speech"; why is cyberspace different?

It would, of course, be a cost not worth bearing if measures like those described here chilled rough-and-tumble political debate and scathing social criticism. There is already too much faux civility in our public discourse, which permits charlatans and villains to claim the patina of legitimacy because no one dares, for fear of being rude, to call them out for what they are. We do not want to regulate speech, on the Internet or elsewhere, in a way that would make it impossible for a modern-day H. L. Mencken to excoriate his targets. But surely it is not hard to draw the line between a Menckensque scathing critique of public figures and calls to sexually assault female law students or to make someone "cry" and humiliate her in front of her children. Recognizing how little moral standing cyber-cesspools have as bastions of "valuable" speech about private individuals ought to encourage us to rethink Section 230 of the Communications Decency Act and rethink tort liability for search engines like Google. The Internet, and the real world, would both be better places if we did so.[31]