# THE UNMASKING OPTION

## JAMES GRIMMELMANN[†]

I'd like to tell a story about online harassment and extract a surprising proposal from it. I'm going to argue that we should consider selectively unmasking anonymous online speakers, not as an aid to litigation, but as a substitute for it. Identifying harassers can be an effective way of holding them accountable, while causing less of a chilling effect on socially valuable speech than liability would.

In the end, I'll conclude that this proposal is unworkable due to the danger of pretextual uses of an unmasking remedy by plaintiffs looking to engage in extra-legal retaliation. Even this conclusion, though, has something valuable to teach us about the uses and abuses of online anonymity. Decoupling anonymity from liability enables us to understand more clearly what's at stake with each.

## I. SKANKS IN NYC

To set the stage, let's talk about Skanks in NYC.[1] That's the name of an anonymous blog someone created on Google's Blogspot service. Actually, calling it a "blog" may be something of an overstatement. It consisted of five entries, all posted the same day, in which the anonymous author called a model named Liskula Cohen, a "psychotic, lying, whoring . . . skank," "Skankiest in NYC," a "ho" and so on.

Cohen filed for a "pre-action disclosure" order against Google to disclose the anonymous blogger's name so she could sue for defamation. The blogger objected, saying the posts were just hyperbole and "trash talk," not anything actionable. The judge, however, agreed with Cohen, looking to the American Heritage Dictionary definition of "skank" to conclude that calling someone "disgustingly foul or filthy and often considered sexually promiscuous" is defamatory. Thus, since Cohen had a "meritorious cause of action," the judge ordered Google to disclose the blogger's identity.[2]

In an O'Henry-esque plot twist, the anonymous blogger turned out to be one Rosemary Port—if not quite a friend of Cohen's, then certainly

---

　　1. 　Wendy Davis, *Judge Rules That Model Has the Right to Learn 'Skank' Blogger's Identity*, MEDIAPOST, Aug. 17, 2009 http://www.mediapost.com/publications/?fa=Articles.showArticle&art_aid=111783.
　　2. 　Cohen v. Google, Inc., 887 N.Y.S.2d 424, 428-30 (N.Y. Sup. Ct. 2009), *available at* http://m.mediapost.com/pdf/Cohen_doc.pdf.

a frenemy. According to an (anonymous) source who spoke to the New York Post,[3] the source of Port's anger was that Cohen had criticized the company Port kept to Port's boyfriend. After learning who her antagonist was, Cohen filed a $3 million defamation suit, but quickly dropped it, saying, "It adds nothing to my life to hurt hers. I wish her happiness."

*A. Right and Wrong*

Port's conduct may have been unfortunate, but what should we make of Cohen's? Although they vary in the threshold they require the plaintiff to meet, courts across the country agree that a "John Doe subpoena" of this sort should issue only where the plaintiff appears to have a winnable lawsuit against the (as-yet unknown) defendant. Cohen represented to the court that she had an urgent legal need for Port's identity—to file her defamation lawsuit—and that was the basis for the court's ruling. But almost as soon Cohen had Port's name in hand, the lawsuit went by the wayside. So much for urgent legal need. Was this a hypocritical abuse of the legal system?

Dan Solove thought so. He's written, "The law must restrict bad-faith lawsuits designed solely to unmask anonymous speakers."[4] He saw Cohen's suit in precisely those terms, saying it appeared "she was using the lawsuit only to unmask the blogger."[5] For him, the Skanks in NYC case is an abuse of the justice system.

I think Solove has things exactly backwards. *Cohen v. Google* wasn't an abuse of the justice system, it *was* justice. Rosemary Port got exactly what she deserved. She tried to shame Cohen; the result was that she herself was shamed. That seems about right. There's something beautifully Kantian about it. Lawrence Becker would say that it was a "fitting" and "proportionate" "return."[6]

It strikes me as a good thing that Cohen dropped her lawsuit. For one thing, lawsuits are shockingly expensive. Cohen resolved her beef against Port for a small fraction of what litigation through final judgment would have cost. If the only response to online harassment is willingness to litigate, then only the rich will have any protection against it at all. For another, what more would Cohen have achieved by carrying her lawsuit through to the bitter end? Port was apparently close to judgment-proof, which is another way of saying that a verdict for Cohen would have

---

3.    Lachlan Cartwright et al., *Secret Grudge of NY 'Skankies'*, N.Y. POST, August 21, 2009, at 9, *available at* http://www.nypost.com/p/news/regional/secret_grudge_of_ny_skankies_f6c4ttnK4 zchSR51tDJoYJ.
    4.    DANIEL J. SOLOVE, THE FUTURE OF REPUTATION 149 (2007), *available at* http://docs.law.gwu.edu/facweb/dsolove/Future-of-Reputation/text/futureofreputation-ch6.pdf.
    5.    Posting of Daniel J. Solove to CONCURRING OPINIONS, http://www.concurringopinions.com/archives/2009/08/can-you-be-sued-for-unmasking-an-anonymous-blogger.html (Aug. 25, 2009, 7:04 EDT).
    6.    LAWRENCE C. BECKER, RECIPROCITY (1990).

bankrupted Port without actually achieving anything for Cohen. And for yet another, it's not self-evident that Cohen would have won a defamation suit. I'm more confident that calling someone a "ho" and a "whoring skank" online is morally wrong than I am that it's legally actionable. In many cases, the convoluted doctrines of defamation and privacy law will deny recovery for reasons that have little to do with the blameworthiness of the defendant's conduct.

Perhaps this lawsuit was pretextual. But if so, then bring on the pretextual lawsuits! It's better to have pretextual lawsuits that are resolved quickly and lead to appropriate embarrassment than protracted lawsuits that cause serious additional harm to the defendant. And once we put it this way, why not cut out the middleman? If there's nothing wrong with a pretextual lawsuit brought to unmask the defendant, we might as well drop the fiction of the lawsuit as the basis for unmasking. I'm proposing, in other words, that the legal system *prefer* unmasking to the standard remedies at law. Without dwelling on the details, what if we had a system that routinely unmasked defendants, one that channeled plaintiffs into unmasking and away from damage suits?

## II. A THOUGHT EXPERIMENT

Thus, here's a proposal for a kind of minimally invasive surgery to deal with online harassment. Suppose that we were to give the victims of online harassment an expedited procedure to unmask their harassers. Specifically, following a quick judicial proceeding with an easier required showing, a relevant intermediary would be required to turn over whatever it knew about the harasser (typically an IP address or subscriber information). In return, the plaintiff would be required to give up all remedies at law. These two rules, taken together, would channel many cases into unmasking rather than into litigation.

My intent is not to endorse complete reciprocal transparency in all things, along the lines of David Brin's *The Transparent Society*.[7] That's a recipe for madness; privacy is a basic element of the human condition. Most people who choose to go online without identifying themselves have a good reason for it, and we should ordinarily respect that decision. I'm also not suggesting any new data-retention requirements. At least for now, the Internet's ad hoc balance—it's easy to keep your identity superficially private and hard to keep it truly private—is about right. The harassers we really think we can reach—the AutoAdmit posters, the lulz-mobs, the Rosemary Ports—aren't using advanced techniques to hide their identities.

There are many things to like about unmasking. In the first place, it's particularly effective at dealing with harassment. Many of the worst

---

7.    DAVID BRIN, THE TRANSPARENT SOCIETY (1998).

cases involve online mobs: crowds of mutually anonymous people who spur each other on to increasingly nasty behavior. One of the best ways to bust up a mob is to call out its members by name, like Atticus Finch in front of the jailhouse. It rehumanizes them, activating feelings of empathy and shame, removing the dangerous psychological condition in which they fear no reprisal. In this respect, visible acts of unmasking—which make members of the crowd more aware that their actions have consequences—may be a more effective deterrent than actually punishing them.

Unmasking also has some major advantages over other possible responses to anonymous online harassment. The First Amendment puts significant limits on the use of tort law. This leads to cases in which harmful, wrongful speech can't be redressed through a suit for damages. In response, we've seen equally dangerous calls to pare back the First Amendment's protections. Unmasking sidesteps that dilemma. Not all the speech that we'd like to protect under the First Amendment needs to be protected as *anonymous* speech.

Similarly, unmasking is a better option in many cases than holding intermediaries liable. The typical poster to a web site is more morally responsible, and better able to control her own speech, than the web site operator, its hosting provider, or the ISP. Making any of these intermediaries liable is likely to lead to substantial chilling effects, as they take down any potentially problematic material at the drop of a hat. Our experience with the DMCA in this regard hasn't been particularly cheerful. In contrast, requiring these intermediaries only to turn over what information they have on the identity of the poster is a smaller burden, and one that doesn't give them bad incentives to take down too much material.

On balance, an identification requirement is likely to be more speech-friendly than most of the alternatives on the table. It avoids the excessively censorious effects of direct and intermediary liability—but it also helps protect the speech interests of the *victims* of anonymous online harassment, who in many cases today are forced off the web in fear.

## A. Shame, Good and Bad

Let us be clear. An argument for regular unmasking is, in effect, an argument for vigilantism. One of the reasons unmasking works is that it exposes anonymous harassers to mass shaming. Solove has argued[8] that online shaming can be "the scarlet letter in digital form," a point he illustrates with the story of Dog Poop Girl, who was vilified by millions on the Internet after failing to clean up after her dog on the subway.

From that perspective, to unmask posters is to open up Pandora's Box. Rosemary Port could become the next Dog Poop Girl, her face plas-

---

8.      SOLOVE, *supra* note 4, at 1–11.

tered everywhere online, as millions of people mock her, exposing her to shame and retaliation that far exceeds anything she deserved. Aren't we unleashing exactly the same forces of hate and innuendo that we're supposed to be tamping down, leading to a never-ending shame spiral? Compared with legal process and societal oversight, isn't this illiberalism, pure and simple?

Perhaps. But if so, it's a surprisingly tolerable kind of illiberalism. The legal system does violence, too; it uses the full power of society and the state against its victims in a very real and direct way. Dog Poop Girl-level abuse will be rare, but damage lawsuits in run-of-the-mill harassment cases will routinely all but wipe out defendants. If the alternative is being sued into bankruptcy, online shaming isn't the worst option out there.

Perhaps even more tellingly, look who started the hate. As between the innocent plaintiff and the defendant who originally posted mean things about her, it seems clear which of these two ought to bear the risk of a disproportionate response. There's still a plausible fit between the harm the shamer caused and the consequences she must endure. And if massive online shame for the shamer *is* a potential outcome, this seems like a singularly appropriate form of deterrence, one that might actually be psychologically effective with would-be harassers.

## B. Retaliation

And now for my own O'Henry-esque twist. I've just argued that an unmasking option is superior on most theoretical dimensions to traditional lawsuits. But I don't see a way of making it work in practice.

Sometimes a lawsuit, with a good old-fashioned damage remedy, really is the best outcome. If harassment leads you to lose your job, that's a real, economic harm, and compensatory damages make sense. Forcing a plaintiff to give up any hope of that remedy is making matters worse.

In theory, we could design the unmasking option so that the plaintiff gets to choose between unmasking (with a lowered threshold) or a lawsuit (with the usual John Doe subpoena standard). But that's an awful choice to put the plaintiff to, because of Arrow's Information Paradox. Until she finds out who her harasser is, she's not in a good position to choose: she can't tell whether the harasser is embarrassment-proof or judgment-proof. What if she chooses the identification, only to learn that her nemesis is a rich recluse who enjoys victimizing women and doesn't care about his own reputation?

If the unmasking option is unfair to plaintiffs, it's also unfair to defendants. You can bet that a corporate CEO would love to characterize some salty criticism of his leadership as "harassment," trace it back to an employee, and take a little revenge. Here, even if we require the plaintiff to give up legal remedies, identification itself imposes serious harms. A

company that can retaliate in ways other than filing a lawsuit would be delighted with the unmasking option's lowered threshold.

Thus, it turns out that the trade at the core of the unmasking option—get an identity in exchange for giving up the right to sue—is poorly matched. Sometimes plaintiffs get far too little; sometimes they get far too much.

## C. Pretext

This conclusion, however, tells us something important about online privacy. Many anonymous posters justifiably fear the pretextual plaintiff. As soon as we lower the standard to unmask people online, we open the door to all sorts of disquieting uses. Companies want to unmask whistleblowers, and perhaps some stalkers might find a way to use it to learn more about their victims.

This is a classic problem of privacy as a second-best solution. I said earlier that people have legitimate reasons to go online anonymously. Our belief that those reasons are legitimate stems from the idea that it would be wrong for these people to have to suffer being fired, being stalked, being personally embarrassed, and so on. But in many cases, these wrongs are harms the law has principled reasons not to redress directly, or simply has practical difficulties in dealing with. Free speech rights, freedom of contract, and the difficulties of proving causation will mean that many people who suffer retaliation will have no legal redress for it. Anonymity is the best we can practically do, and so, unless we're prepared to make much bigger changes to the legal landscape, we'll have to protect people from pretextual unmasking.

But if the fear of pretext is legitimate, the strength of the plaintiff's cause of action isn't always a very good proxy for it. Some plaintiffs will have a valid lawsuit, but bring it for totally pretextual reasons—a few stray comments about a mid-level corporate executive could blow a whistleblower's anonymity. Contrariwise as I've been arguing, there are plenty of people who ought to be unmasked, but who haven't done anything actionably tortious, given the labyrinthine folds of defamation and privacy law. Pretextual lawsuits need not be baseless, and vice versa.

## CONCLUSIONS

Thus, I take two lessons from this thought experiment. The first is that we need to decouple unmasking and litigation. The precise inversion I proposed—give up your lawsuit to make unmasking easier—doesn't work. But we should be more creatively exploring unmasking standards that aren't directly tied to the strength of the plaintiff's case in chief. We should consider the pros and cons of unmasking directly, on their own merits, without always referring back to the lawsuit.

So, on the one hand, in order to better protect the victims' interests in these lawsuits, we should find ways of dropping elements from a typical John Doe subpoena. Thus, for example, a plaintiff typically needs to show necessity: that she's exhausted other options to learn the harasser's identity. Chuck that one out the window; if the plaintiff thinks that asking the intermediary for the identifying information is the best way to learn who the harasser is, that ought to be good enough for us.

On the other hand, to protect defendants, we should be more explicit about pretextual unmasking. Right now, we're protecting defendants by testing the strength of the plaintiff's case. We should acknowledge explicitly that the true threat is retaliation, and develop doctrines that directly ask whether the defendant legitimately fears retaliation from being unmasked. Those doctrines could then usefully be applied in any case where unmasking is at stake, regardless of the area of law in which it arises.

This is a Legal Realist argument. It's concerned with the social goals the law is trying to achieve—and with what the law on the ground is actually doing, regardless of what the law says it's doing. A John Doe subpoena standard that sees only the strength of the plaintiff's case is ultimately both unjust and unstable, because it's asking the wrong question. Unmasking is the very best kind of wrong answer: it helps us understand the question we meant to ask.