



Privacy and Assurance: On the Right to be Forgotten

Scott Casleton, Philosophy, University of California, Berkeley, US, scott_casleton@berkeley.edu

The right to be forgotten enables individuals to remove certain links from search results that appear when their names are entered as search terms. Formulated as a distinct application of the general right to privacy, the right to be forgotten has proven highly controversial, for two reasons. First, it is difficult to see how the specific right to be forgotten can apply to the *withdrawal* of public information, since the general right to privacy typically covers the *disclosure* of private information. Second, as a putative right to withdraw information from public reach, the right to be forgotten poses a threat to freedom of speech, which depends on the accessibility of information. By responding to these two objections, this paper develops a novel account of the right to be forgotten, understood as a *claim of withdrawal* grounded in both privacy and free speech interests.



Privacy and Assurance: On the Right to be Forgotten

SCOTT CASLETON

Philosophy, University of California, Berkeley, US

Search engines like Google are extremely useful tools for finding information. Instead of spending time traveling to the local library or spending money on an encyclopedia, you can look up information on search. In this sense, search reduces the cost—in terms of time, effort, and money—of finding information. By reducing the cost of access to information, search engines provide us with an unequivocal benefit.

But easier access to information for one person may come at the expense of another. If I make it easier to peek over your fence while you sunbathe by setting up a stepladder, I may increase the incidence of violations of your privacy. You certainly have a privacy-based objection to being non-consensually observed while you sunbathe, and you also have an objection to my making it easier to observe you, so long, at least, as I have no good reason for setting up the stepladder. For me to increase the risk that your privacy will be violated for no good reason would itself count as a violation of your privacy right.

Search engines, like my step ladder, make it easier for people to access information about us, thereby increasing the risk that our privacy will be violated. This is what underpins the thesis that our right to privacy affords us protections against search engines—specifically, by including a right to be forgotten. Famously, a Spanish citizen, Mario Costeja González, sued Google in a European court in order to have a link removed from search results, which link led to a digitized newspaper announcement regarding a state-organized auction of Costeja's estate to pay off social security debts.¹ The European court sided with Costeja against Google, citing Costeja's claim to have his privacy protected by way of “de-listing” the link, which kind of privacy claim is now typically called the right to be forgotten.²

The concept of a right to be forgotten has generated a good deal of skeptical, even hostile, commentary. One line of skepticism notices a difference between search engines

¹ See Judgment of 13 May 2014, *Google Spain SL, Google Inc. v Agencia Española de Protección de Datos, Mario Costeja González*, C-131/12, EU:C:2014:317, paragraph 14 (hereafter *Google Spain*).

² For discussion, see Kulk and Borgesius 2018.

and my stepladder. Whereas my stepladder helps make publicly available information that is (still) private, search engines increase the accessibility of information that is (already) public. Why should the right to privacy—the right to be forgotten—give a person control over information if it has already entered the public sphere? A second form of criticism highlights the threat posed by the right to be forgotten to a culture of free speech. Affording everyone a right to withdraw links from search engine results jeopardizes vast swathes of information—information the public often has an interest in accessing, discussing, and circulating.

My goal in this paper is to respond to these two objections to the right to be forgotten and to thereby develop a positive account of this right. The bulk of my argument focuses on clarifying the normative structure of our claims to control some information once that information has in some sense become public—in particular, claims to *withdraw* information from search results. I make two key moves. First, I argue against the idea that just because someone’s personal information was legitimately disclosed—say, by a newspaper—that everyone is entitled to share or access this information in any way whatever. Such an approach, I claim, would wreak havoc on our interests in privacy. Second, I argue that withdrawal claims provide a form of insurance against future harm that might result from divulging personal information in the present. We often want to share personal information without knowing what the exact consequences of such sharing will be, and the right to be forgotten, by assuring us that we will not suffer gratuitously in the future for sharing personal information now, allows us to open ourselves up to others in a less inhibited way.

I refer to this second idea as the assurance value of the right to be forgotten. Crucially, by assuring us that we will not suffer unduly in the future by sharing information in the present, the right to be forgotten not only promotes our privacy interests but also our free speech interests. I thus argue that, contrary to the prevailing wisdom, the right to be forgotten does not conflict with free speech but in fact supports it. Interpreting the right to be forgotten as a form of withdrawal claim, based not only on the value of privacy but assurance as well, helps us to see this fact.

I. PRIVACY INTERESTS

It is common to interpret the right to be forgotten as a form of privacy claim.³ The appeal of this approach is easy enough to appreciate. Roughly put, a person’s right to privacy covers a variety of claims against people accessing information about her. The right

³ See *Google Spain*, paragraphs 3, 9, 38.

to be forgotten—as in cases of de-listing results from search—removes one pathway for people to access information about a person. That is to say, the right to privacy entitles a person to control how (or whether) her personal information is accessed by others, and the right to be forgotten seems to be one instance of a person exercising this control.

By conceptualizing the right to be forgotten as an instance of the privacy right, however, I encounter a problem. As Judith Jarvis Thomson famously said: “Perhaps the most striking thing about the right to privacy is that nobody seems to have any very clear idea what it is.”⁴ One may think that explaining the right to be forgotten as a privacy right cannot be a promising enterprise until the right to privacy itself has been demystified.

For our purposes, this problem is not insuperable. While an enormous amount has been written about the right to privacy, in this context all we need is an account of some of the *interests* that the right to privacy protects.⁵ This is different from giving a full account of the right to privacy, which would involve an exhaustive list of these privacy interests along with a systematic method for balancing these interests against competing moral values.⁶ Such a full account of the right to privacy would constitute an explanation of the *structure* of the right to privacy. My goal is to provide a partial account of the structure of the right to privacy by explaining how one application of this right covers claims to *withdraw* information from certain public fora or to *halt* the further spread of information by some person. This is the right to be forgotten.

Explaining this structural feature of the right to privacy, as I said, does not presuppose a complete account of the interests that are safeguarded by the right to privacy. Still, it is necessary to canvass a few of these interests, for two reasons. First, we need to understand the content of privacy claims so as to assess their *strength* and, ultimately, balance them against free speech considerations. Second, we need to bring into view the legal standard of harm that has been applied in right to be forgotten cases, as this is part of what needs to be explained in an account of the right to be forgotten.

Of the interests protected by the right to privacy, not all are uniquely *privacy* interests. For instance, I have a generic interest in not being blackmailed, and this interest just happens to be protected by my claim against you reading my diary uninvited. The primary focus of theories of privacy, though, is to provide an account of certain interests that are uniquely suited to explaining our intuitive objections to privacy violations.

⁴ Thomson 1975, p. 295.

⁵ Compare Scanlon 1975.

⁶ See, for example: Marmor 2015; Nissenbaum 2004; 2009; and Allen 2000.

At least three such interests are widely acknowledged. These include an interest in maintaining distinct kinds of relationships; an interest in avoiding unwanted intimacy; and an interest in avoiding psychological harm, such as shame or embarrassment.⁷

This list is not exhaustive. For we agree that there are some plain cases of privacy violations that do not frustrate any of these interests. For example, you might read my diary without my permission and yet not interfere with my ability to maintain valuable relationships, foist on me unwanted intimacy, or cause me to feel psychological discomfort. Figuring out what interest is harmed in such cases is a difficult philosophical problem; solving it is not my goal, here. I bring attention to this issue simply to indicate that some cases of the right to be forgotten may involve intuitive violations of privacy without us having a ready explanation as to which interest of the right-holder has been set back. This is, in fact, reflected in the *Google Spain* ruling, as the court said of the right to be forgotten, somewhat cryptically, that “it is not necessary... to find such a right that the inclusion of the information in question in the list of [search] results causes prejudice to the data subject.”⁸ Strictly speaking, I do not agree with this, since a violation of the right to privacy—of the right to be forgotten—must involve *some* interest being set back. Identifying which interest, though, may depend on one’s theory of the right to privacy. To give an account of the right to be forgotten, all we need to assume is that in certain cases, as with the case of you reading my diary, a person’s interests are indeed harmed, and thus the right to privacy violated, even if the interest in question eludes us, or is controversial.

As I have said, what connects the right to be forgotten with the right to privacy is that both concern an individual’s control over her personal information and how (or whether) it is accessed by some public. In some cases, the right to privacy grants a person full discretion over the disclosure of a particular piece of information. For instance, it is up to me and only me whether I reveal to other people whom I voted for in the 2016 US Presidential election. In other cases, though, the right to privacy only grants a person partial discretion over the disclosure of some piece of information. My discretion to choose who gets to look in the trunk of my car, for example, is limited by the discretion of the police to search it when public safety justifies their looking. The right to be forgotten, I will show, covers a claim of partial discretion insofar as it applies to links in search results but not the original publisher’s website to which the link on search directs.

⁷ On relationships, see Fried 1968 and Rachels 1975. On intimacy, see Nagel 1998. On psychological harm, see Marmor 2015.

⁸ *Google Spain*, paragraph 96.

What makes the right to be forgotten distinct, then, is not that it applies in a limited way but that it applies to information that has already entered public view, or at least become available for public viewing. My argument is that a claim to withdraw information from public view is very often justified because it serves privacy interests, and in some cases free speech interests. The key question is why these interests are not undercut or defused by the fact of a prior, legitimate disclosure of the information in question. I will address this question in Sections III and IV; it is worthwhile to first consider an argument that an alternative framework for the right to be forgotten is available that does not rely on the admittedly unusual idea of a privacy-based right of withdrawal.

II. REPUTATION?

Hannah Carnegie-Arbuthnott argues that the right to be forgotten should be understood in terms of *reputational* interests rather than privacy interests.⁹ More specifically, Carnegie-Arbuthnott argues that the right to be forgotten is violated when someone, such as a search engine operator, distorts another person's reputation. The concept of *reputational distortion* is a cousin of the more familiar idea of *reputational defamation*. We all agree that it is objectionable for Steve to defame Sally by circulating *false* information about her. Carnegie-Arbuthnott's strategy is to appeal to this consensus on objections to defamation and argue that we also have objections to distortion—specifically, when someone circulates *true* information about us that is *outdated or irrelevant*.¹⁰ This would seem to mirror Costeja's own claim, in the *Google Spain* proceedings, that the issue of his social security debts “had been fully resolved for a number of years and that reference to them was now entirely irrelevant.”¹¹

My aim in this section is to criticize this approach so as to highlight the strengths of the privacy view. While I do not think the reputation-based view is the correct model for the right to be forgotten, I suggest it may still point to certain objections we have to the way search engines present our personal information to the public.

Carnegie-Arbuthnott begins with a specific account of the wrong of defamation, with the aim of extending the underlying rationale to reputational distortion. She says that the right against defamation “arises from the principle that a person should

⁹ Carnegie-Arbuthnott 2023.

¹⁰ See Carnegie-Arbuthnott 2023, pp. 3, 13.

¹¹ *Google Spain*, paragraph 15.

only be held accountable for actions which are properly attributable to them.”¹² When Arnold dismisses Delilah’s statements because he has been told, falsely, that she is an inveterate liar, he holds her accountable for actions—telling lies—which are not properly attributable to her. Delilah’s right against defamation was violated by whomever spread this lie.

It is not perfectly clear, however, how this model is supposed to be extended to the case of reputational *distortion*. Defamation involves one person causing another to hold a false belief—but, as Carnegy-Arbuthnott emphasizes, distortion occurs “when true information from someone’s past is presented in a way that suggests it would be appropriate to hold them accountable for it, when it is no longer appropriate to do so.”¹³ Perhaps the idea is this: by learning true information that is outdated or irrelevant, a person may draw a mistaken inference and form a further, *false* belief. If I truly tell you that Sally was fired from her first two jobs because she kept showing up to work late—without my adding that Sally has since changed her tardy ways—you might wrongly infer that Sally is not a dependable person. This false belief could cause you to treat Sally worse than you otherwise would, to hold her accountable for character traits that are not, in fact, applicable to her.

Though intelligible as a way of wronging someone, this notion of reputational distortion does not map on to the right to be forgotten. The right to be forgotten seems to ground claims against the availability of information on search independently of any worries about others forming false beliefs about us. Suppose, for example, that a college newspaper truly and legitimately reports about a female student that she had an abortion. Later in life, this woman may want this article removed from search, and the removal seems *prima facie* justified. Yet the justification for her claim need not depend on the worry that others will draw mistaken inferences about her character. It seems more natural to say that this piece of information is just not anyone else’s business—that it belongs to her zone of privacy—rather than that the information may lead to the formation of false beliefs on account of its being outdated or irrelevant.

Carnegy-Arbuthnott might reply by saying that the wrongfulness of distortion does not depend on the generation of false beliefs but instead on the fact that reputational distortion involves people concentrating on some facts about us rather than others, which may reduce a person’s power to decide how she presents herself to others.¹⁴ For instance, if people can google Sam’s name and discover that he is the son of a

¹² Carnegy-Arbuthnott 2023, p. 13.

¹³ Carnegy-Arbuthnott 2023, p. 2, my emphasis.

¹⁴ See Carnegy-Arbuthnott 2023, pp. 11, 14, 18.

reviled politician, they may be inclined to make jokes about Sam or simply hold him in contempt. When people have Google at their disposal in this case, Sam has less control over how he presents his public-facing personality, potentially to his detriment.

But this interpretation of the concept of distortion will not move us closer to explaining the right to be forgotten. As a general matter, I am skeptical that we have generic claims against people causing some audience to lower their opinion of us, or to judge us negatively. Suppose a French comedian makes jokes about Americans. This may cause one of my students, who is a French foreign exchange student, to regard me with contempt. I am not inclined to think I am harmed by being thought of this way, and thus I do not seem to have even a *pro tanto* objection to the comedian's joke-telling. In any event, the range of right to be forgotten claims is much narrower than the range of cases where our standing in the eyes of others might be negatively affected. There are lots of cases where people may be concerned about their reputations—say, a former college quarterback who now does not want to be associated with football because it has become uncool—without having even a *pro tanto* claim against the availability of some digitized news article on Google.

Legitimate right to be forgotten claims seem to track cases in which a person's privacy interests are harmed, as in cases where a person wants to avoid uninvited intimacy or the pain of being exposed to shameful exposure. It seems to me that the reputation view is too capacious in what it counts as a harm, such that it cannot discriminate between legitimate right to be forgotten claims and cases in which a person is worried others will, for whatever reason, think less of her.

The privacy model of the right to be forgotten thus seems more appealing than the reputational model. However, this is not to say that we should dispose of the concept of reputational distortion altogether. Sometimes it seems quite true that we have objections to people indirectly causing others to hold false beliefs about us by telling them true but irrelevant or outdated information. As we noted, Sally may have an objection to my blithely sharing information about her employment history with others, because this might cause them to hold false, damaging beliefs about her. I thus leave open the possibility that search engine operators could be guilty of distorting our reputations in this way. Search engines may provide information that lacks sufficient context to prevent the drawing of false inferences.¹⁵ Whether objections to this kind of harm ground a right to be forgotten claim is a separate matter.¹⁶ I hope to show that the

¹⁵ Compare Carnegy-Arbuthnott 2023, pp. 2–3.

¹⁶ It might only ground a claim to provide context, for example.

privacy model of the right to be forgotten can accommodate the kinds of cases I have discussed, which seem not to fit into the reputational framework.

III. DISCLOSURE AND WITHDRAWAL

Paradigmatically, the right to privacy covers a person's claim to decide whether or how to *disclose* information. Yet the right to be forgotten concerns the *withdrawal* of information, say from search results. It is therefore conceptually puzzling how the right to be forgotten could be an instance of the privacy right.¹⁷ Worse, the normative basis of a withdrawal claim is obscure. If information about a person is legitimately disclosed, either by a third-party such as a newspaper or by that person herself, what could ground that person's claim to control how others use that information? My goal in this section is to address these two issues.

Begin with idea of information withdrawal. While at first sight it may seem odd to say that the right to privacy can include claims to withdraw information, there are uncontroversial cases where we recognize such claims. For example, if you steal an intimate photo from my wall-safe and make photocopies of it, I can demand that all these photos be returned.¹⁸ Or, consider a case in which one person voluntarily gives another person a nude photo, only for the recipient to multiply it—say, digitally—and share it with others. The mere fact that the initial act of sharing in this latter case was voluntary—that the initial disclosure was legitimate—does not entail that the recipient has an unrestricted right to circulate the photo.¹⁹ My broader case for claims of withdrawal takes this observation as its point of departure.

I will consider two kinds of legitimate disclosure, third-party and first-person. Third-party legitimate disclosure occurs when someone, such as a journalist, publishes a person's personal information without violating that person's privacy right. Very often this occurs when someone *forfeits* a privacy claim by acting in such a way as to lower her objection to the publication of the personal information, as when someone commits a crime and forfeits her claims over the publication of personal information

¹⁷ Compare Leta Jones 2016, p. 81: “issues raised by the right to be forgotten are difficult to understand as *privacy* issues because they are about information that has been properly disclosed but has become or remained problematic.”

¹⁸ Basu (2022) argues for the unusually strong thesis that a person can even have a duty to *forget* some information when that information is garnered through a privacy violation.

¹⁹ Compare Citron (2015), discussing the phenomenon of “revenge porn.”

related to the crime.²⁰ First-person legitimate disclosure occurs when someone *waives* a privacy right, as when someone intentionally reveals personal information on a blog.

The idea that first-person waiver depends on intent to disclose is important. It is usually argued that in the absence of intent to disclose, a person cannot be said to have waived a privacy right. Benedict Rumbold and James Wilson argue that the intentional disclosure of a data set does not necessarily imply that the person who disclosed this set waives her right over all information that can be *inferred from* this data set.²¹ Suppose, for example, a woman intentionally discloses a portion of her DNA to a researcher, but does not intend to disclose that she has a rare genetic disease, and she is unaware that this fact can be inferred from an analysis of her DNA.²² In cases such as this, I agree with Rumbold and Wilson that the woman retains a claim on others not to publish—or, say, sell—this further information about her rare disease.²³ This line of argument suggests a tight link between intent to disclose and waiver.

We must be careful, though. For speaking of intent to disclose *simpliciter* obscures the fact that we typically intend to disclose information *with respect to specific audiences*. My argument is that, in cases of first-person disclosure, we can intentionally disclose information to one group of people without licensing the transfer of that information to a wider audience. What is more, as I argue in the next section, we can sometimes retain claims over personal information *even when* we intentionally disclose information to the widest possible audience. For now, the claim I want to focus on is that the intentional disclosure of information, insofar as it is audience relative, may carry normative constraints on how others may use the information we disclose to them. Compare, in this regard, the transfer of private property. Whereas when I give you my old bike you are at liberty to do with it whatever you please, when I tell you that I take depression medication, you are intuitively not at liberty to do what you like with this information.

To be sure, this is not *always* the case. Sometimes our personal information is disclosed such that *all* of our privacy claims are lowered. For instance, when I publish an op-ed in the *New York Times*, intending to reach as wide an audience as possible, I seem to lower all my privacy claims over this information. Those who read the op-ed are free to pass on this information to whomever they like. Or, to take a case of third-party disclosure, I may run for President and thereby forfeit all my claims against journalists

²⁰ See Hanin (2022) for a detailed discussion of privacy-rights forfeiture.

²¹ Rumbold and Wilson 2019.

²² I take this example from Rumbold and Wilson 2019, pp. 14–15.

²³ But see Hanin (2022, especially pp. 260–1), arguing that unintentional revelation can sometimes amount to forfeiture.

digging up and publishing information about specific areas of my past, such as my tax filings. Considering such cases of forfeiture will allow us to refine the idea of claims on the further sharing of information that has been legitimately disclosed by third parties, most notably in the *Google Spain* case.

Suppose I am arrested for drunk driving, and this fact is reported in the police blotter of the *San Francisco Chronicle*. I have no objection to the *Chronicle* reporting this information; it is an instance of legitimate information disclosure. However, we may ask: does this mean that I forfeit all standing to object to other people's sharing of this information? Is my colleague, who reads the *Chronicle* every day, at liberty to share this information with our boss?

I do not think so, at least not necessarily. The fact that I have forfeited a claim over the disclosure of my personal information in the *Chronicle* does not entail that I have forfeited all claims over this information. If telling my boss about the incident would harm my privacy interests without any foreseeable benefits, I cannot see how my colleague could be at liberty, normatively speaking, to share this information. The point of the right to privacy is precisely to protect us from gratuitous harm to our privacy interests. Of course, if some comparably serious benefit were in the offing, my colleague might have a justification for passing along the information. But passing along the information is not justified *just because* the *Chronicle's* initial disclosure was legitimate.

It may be asked, here, why I have a claim against my colleague relaying the information to my boss given that my boss could just as well read about the incident in the *Chronicle* herself.²⁴ The same interest would be harmed in either case, after all. In reply, I note that the right to privacy does not entitle us to others not knowing things about us, but that they not come to know things about us in specific ways.²⁵ Remember my stepladder. You, sunbathing in your backyard, are not entitled to not being looked at *tout court*. For a youth climbing a tree might accidentally spot you from a leafy bough without violating your right to privacy. I, however, am prohibited from peeking over your fence by standing on a stepladder. The right to privacy blocks certain pathways, so to speak, that others might take to access personal information about us; but it does not necessarily block all pathways. Which pathways the right to privacy blocks depends on the costs that are involved in blocking them. I take it that the costs of a general prohibition on setting up stepladders next to the fences of private homes are not unreasonably high, whereas the costs of a general prohibition on just-for-fun

²⁴ I thank a reviewer for raising this question.

²⁵ Compare Thomson (1975, p. 307), discussed by Marmor (2015, p. 4).

tree-climbing would inordinately burden the interests of the neighborhood kids in having fun. Analogously, in the drunk driving case, I am entitled to my privacy being protected compatible with not unduly burdening other people's interests. The *Chronicle* promotes the public interest when it publishes my personal information, and therefore does not violate my right to privacy. By contrast, the restriction on my colleague's liberty to tell other people in the workplace about my drunk driving incident protects my privacy interest without substantially burdening anyone else's interests.

We can now provide an initial account of the *Google Spain* case. Recall that, in this case, a Spanish newspaper, *La Vanguardia*, published an announcement for an auction of Costeja's estate that was held to pay off his debt. Like my drunk driving, Costeja's failure to handle his debt justified a newspaper announcement that broadcast a fact about his personal life to the public. Costeja has no objection to the public having access to his personal information by reading it in *La Vanguardia*. However, his right to privacy could very well apply to *other* pathways that the public might use to access this information—such as search engines. This is precisely how I understand his right to be forgotten claim, as a claim against the accessibility of his personal information through one particular pathway, the blocking of which would not impose an undue cost on the public interest. Of course, it must still be argued that the burden on the public interest is not inordinately high. I take this up in Section V. For now, the point is just that the right to be forgotten, as a specific application of the right to privacy, is aimed at blocking particular ways that people might use to come to learn things about us. And, moreover, the raising of a right to be forgotten claim is not undermined by the fact that the person's personal information was legitimately disclosed in the first place, in this case by *La Vanguardia*.

Now, right to be forgotten claims are not limited to cases of third-party disclosure; they apply in some cases to first-person disclosure as well. First-person disclosure is legitimate in most cases because of the privacy subject's *waiver*, as opposed to forfeiture. However, as I have noted, waiver is audience-relative. Just consider, for example, a man who tells his friends that he takes hair loss medicine. He would be justified in feeling indignant if he were to discover that one of his so-called friends turned around and started telling others.

Note that the man's grievance does not depend on some prior *agreement* to respect his intentions vis-à-vis the disclosed information. While we sometimes solemnly ask a confidant to keep information "between us" or in some contexts sign non-disclosure forms, these are outliers rather than representative cases. Very often we intentionally communicate information without securing agreement from our audience that they will keep the information to themselves. The normative demand of respecting someone's

privacy interests in such cases remains in force because the information discloser has not intentionally waived the disclosure of the information *for everybody*.

I want to emphasize, though, that I am *not* saying that in the absence of *knowledge* of someone's intentions we should not share information. In describing the way an information-sharer's intentions limit the further-sharing of information, I have given examples in which either personal knowledge of the speaker or general knowledge of cultural norms, such as what tends to cause embarrassment, allows an audience to infer the speaker's intentions. As a rule, we take a person sharing information about his hair loss medicine or sharing a nude photo as not intended for further sharing. However, most cases of information sharing do not occur between people who can reliably infer an information sharer's intentions, if the original information sharer is even present. To treat every case of information disclosure as, by default, not intended for further sharing in the absence of knowledge of someone's intentions would be too costly. The free flow of information is too valuable to assume that unless we know a person is OK with us passing on some information, we ought not to pass it on. Indeed, in most cases of information circulation, we are passing on information regarding people whom we do not know and whose intentions we cannot divine.

It is at this juncture that we can understand how the right to be forgotten emerges as an instance of the general right to privacy. As I have said, our normative default with regard to information sharing, in the absence of knowledge of a privacy subject's intentions, need not be to keep mum: instead, our normative default should be being at liberty to share. This default of being at liberty to share will keep the information economy running like a well-oiled machine. However—and this is the key point—this default of being at liberty to share should be complemented by individuals having a defeasible claim to halt further sharing. This is to say that we should be at liberty to share information when we do not know the intentions of the privacy subject, but that this subject retains a claim to ask us to stop sharing the information, which includes a claim to remove artifacts containing the information from public view.

Let me illustrate. Suppose I go to my local café to order an oat milk latte, and while my barista is preparing the drink, I flippantly reveal that I voted for Trump in 2016. After leaving the café, the barista, who unbeknownst to me is a social media influencer, posts an Instagram reel describing his encounter with a well-known philosopher who voted for Trump.²⁶ When I see this reel on my Instagram feed, I rush back to the café and ask—maybe I demand—that he take the reel down. Ought he to comply?

²⁶ Indulge my fantasy that I am well enough known for the public to identify me as the person who voted for Trump.

The answer, it seems to me, is “Yes.” As I see it, my barista did not act wrongly by sharing the information in the first place, but he would act wrongly by denying my request for withdrawal. This is just to say that he would act wrongly by violating my right to privacy—my right to be forgotten, in this particular case. It is important to point out that by saying the barista did not act wrongly in the original posting, I assume that he did not know what my intentions were; more exactly, I assume that he did not have good grounds to infer my intentions. This is what allows me to say that the initial disclosure of the information to the public of Instagram was not impermissible. I thus find perfectly justified the conclusion that the permissible, or legitimate, disclosure of information about an individual does not entail that that individual thereafter loses all privacy-based claims over the circulation of the information. Of course, the barista, or whoever else, may have an interest in sharing the information that contends with my privacy interest in withdrawal. But the point here is just that, in virtue of my privacy interests, I have a *pro tanto* claim to the withdrawal of the information, which is not undercut by the legitimacy of the initial disclosure.

We can now see how search engines fit into the picture. Unlike human beings, search engines cannot use context or conversational cues to determine whether it is appropriate to pass on information. They just gobble up information, feed it through an algorithm, and display it in response to queries. On the one hand, it certainly seems like a mistake to impose a default of not-sharing on search engines, since this would drastically stymie the free flow of information. Search does not wrong us by disclosing information that was, in the first instance, legitimately disclosed. But, on the other hand, the indiscriminate nature of search in collecting and presenting information seems to render all the clearer a person’s claim to restrict the further sharing of personal information when it extends beyond her intended audience. Or, in the case of third-party disclosure, when it harms her privacy interests gratuitously.

Let’s conclude this section with an example. Suppose a college student gives an interview with a campus magazine describing her evolving attitude toward her own gender expression. By giving the interview, the student consents to the disclosure of the information, although it seems reasonable to say the size of her intended audience is not anyone and everyone but instead a smaller group of people, those who are likely to pick up the magazine. (Suppose, by comparison, that she would *not* consent to describing her experiences in an op-ed for the *Times*.) In this case, we should first note that if a digitized version of the article were to become accessible through Google, it would increase the potential size of her audience far beyond what she had originally imagined. It is because of this fact that right to be forgotten claims can be directed

against search engines but not the original publisher of the personal information in question. The search engine, but not the original publication, expands the audience beyond what the privacy subject initially intended.

And, crucially, we must add that the search engine has a *weaker claim against the withdrawal request*. The magazine may want to maintain the integrity of its archive, entice new readers, boast about past content, and so on. By contrast, a search engine like Google will have a small pecuniary interest in keeping the link on its search results pages; and supposing my claim applies to all search engines equally, it will not result in competitive disadvantages for any one search engine. At most, the search engine's parent corporation might appeal to the public's interest, that is, the interests of the search engine users. But there is no reason to assume the public's interest will ground an objection to my withdrawal claim in the same way as the magazine's interests.

It therefore seems that the right to be forgotten, understood as an application of the privacy right, can discriminate between claims against search engines and claims against original publishers. The claim against the search engine protects a person from gratuitous harm to her privacy interests and does not impose an inordinate burden on the public's interest in access to information, at least in the kinds of cases I have considered so far. I now turn to strengthening the normative basis for making withdrawal claims, to round out my account of the right to be forgotten.

IV. ASSURANCE AND WITHDRAWAL

If we know that we will, or can, enjoy a claim to withdraw information in the future, we may be more likely to share information *now*. This suggests that our interests in sharing information *now* can support structuring the privacy right to include claims to withdraw information *later*. I will refer to this as the *assurance value* of the right to privacy. The assurance value of the right to privacy helps ground the specific right to be forgotten.

Suppose that you want to share a secret with me, but you are afraid of how I will treat you once I have learned the secret. Now suppose that I tell you that I have a pill that I can take, and this pill will have the following effect: one week from now, I will forget all the new facts I learned on the day on which I took the pill. So, I offer to take the pill so that you can tell me your secret, with the assurance that I will forget the secret in a week. This will allow you to test out our relationship with me knowing your secret. If the test goes well, you can simply re-share the secret with me; if not, not.

The value of this pill lies in the assurance that it gives you that I will forget a week from now. This is its assurance value. This value is constituted by your being better able to satisfy your privacy interests—in this case, your interest in having a special kind of relationship with me, founded in part on the sharing of private information. One way to think about the value of assurance is that it reduces our inhibitions about sharing private information just insofar as it reduces our fear of some future bad consequence. Normally we must balance our desire to satisfy an interest now against a fear of some later consequence. For instance, Victor may want to come out of the closet to his parents so that they will better understand him, but he fears the possibility of their rejection. He is thus inhibited in disclosing his sexual orientation.

In order to better satisfy our privacy interests, we should reduce such inhibitions to the extent possible. One way to reduce inhibitions on sharing is by affording people a claim to withdraw personal information from public view. If a person knows that she will be able to remove certain links from Google in the future, she will have less to fear about selectively disclosing information now. Clearly, this will allow her to disclose information about herself in testing out the boundaries of relationships and intimacy at a particular time, which, remember, are two of the important interests underlying the right to privacy. The right to be forgotten is thus valuable insofar as it provides insurance against certain bad outcomes of a person satisfying her privacy interests.

Importantly, the assurance value of the right to be forgotten is not limited to privacy interests. Assuring someone that she can remove information from Google at a later time will also promote her interests in freedom of expression.²⁷ A college student's willingness to air her opinions in the college newspaper will be less inhibited if she knows that she will be able to remove from Google a link directing users to a digitized version of one of her more embarrassing op-eds. This illustrates how the right to be forgotten is not exclusively in *conflict* with freedom of speech, but in fact can, in some cases, promote freedom of speech, by allowing otherwise inhibited individuals to publicly articulate their ideas. As Seana Shiffrin has argued, the freedom to make public one's thoughts is one of the fundamental values of freedom of speech, since it affords us the experience of freely articulating an idea, receiving feedback, and perhaps revising our beliefs accordingly.²⁸ The value of the public articulation of one's beliefs is

²⁷ Citron (2022) also connects privacy protections with freedom of speech, though Citron focuses on legal protections for privacy violations rather than on what I call the assurance value of the right to privacy.

²⁸ Shiffrin 2011.

promoted by the assurance value of the right to be forgotten, to the benefit of both the speaker and the public at large.

A major caveat regarding the assurance value of the right to be forgotten must be registered, though. Earlier I distinguished between an individual's intentional disclosure of information and third-party disclosures. Right to be forgotten claims are only supported by reasons of assurance in cases of individual intentional disclosure. Consider the original *Google Spain* case in which Mario Costeja González wanted to remove from search a link leading to an article announcing an auction of his estate. There is obviously no sense in which Costeja might have been inhibited when it came to the release of this information—because he was not the one that chose to release it. Evidently, then, we need to restrict right to be forgotten claims on grounds of assurance to those claims that are plausibly tied to the possibility of an individual having felt inhibited in disclosing information in the past. This rules out the possibility that right to be forgotten claims can be applied by individuals who were not, in fact, exploring the values of privacy and free speech. However, an assurance claim can still apply in cases where a third-party reports on a person sharing information in a way that she may have been inhibited from doing. A report, on a third-party platform, about someone's potentially-inhibited disclosure does not bar the person from raising a claim to have a link to the third party platform de-listed.

With this account of the assurance value of the right to be forgotten in place, we arrive at the following picture:

The Right to be Forgotten: An individual has a *pro tanto* claim to the withdrawal of information from some public forum when at least one of the following conditions is satisfied.

- First, when an individual has not, through waiver or forfeiture, licensed the spread of information to a certain audience and the illicit spread frustrates at least one of her privacy interests.
- Second, when an individual could plausibly have been inhibited in sharing the information in the past due to fear of consequences following from its having been made public.

Some comments are in order. The first is that a claim to withdrawal from *a* public forum is not a claim to withdrawal from *all* public fora. This is clear in cases in which my claim to remove a link from Google search does not apply equally well to the magazine webpage that the link leads to. The plausibility of this point is confirmed if we consider

analogous cases of claims to withdrawal. Having donated a box of photographs to the local library, I might request that one of the photos not be displayed in a public exhibition but lack grounds to have the library destroy, or even return, the photo. The library might keep it in its archives for researchers to look at. There is thus a balance at work between my withdrawal claim—based on my privacy and free speech interests—and the public's interest in access to information.

A second comment is that a right to be forgotten claim is stronger when both conditions are satisfied. For example, a young man in college might tell a small social group that he once experimented with homosexual sex. Suppose he runs for political office ten years later in a conservative political district, only to discover someone has blogged about his sexual experimentation. His claim of withdrawal will satisfy both conditions. On my model of the right to be forgotten, his claim is as strong as it could be—and this fits with my intuition about the case, which is that he has powerful grounds to demand the removal of the blog page from search. Of course, this may not be strong enough to outweigh the blogger's interests in keeping the information up, but it may very well outweigh any objection from Google (or, Google on the public's behalf).

Notice that not every case will satisfy both conditions. I might explicitly license the wide circulation of some nude photographs of myself by consenting to an international magazine's invitation for a photo shoot. However, later in life, I may cite the second ground for withdrawal, namely that I had very good reason to be inhibited in consenting to this photo shoot. I had good grounds, that is, to expect these photographs to cause embarrassment to my future spouse, children, and so on. But the weight of my reasons for experimenting with my modeling career and my body earlier in life seemed strong enough to justify doing the shoot—and I was comforted by the fact that I would have grounds for removal later in life.

An objection may be raised, here. If a person is morally responsible for the release of her own personal information, why shouldn't the burden of responsibility include living with the public accessibility of this information? There are plenty of decisions that we make—making promises, accepting a request to help a friend, and so on—which we are bound to live with even if we come to regret making them. Perhaps the disclosure of personal information should function this way, with the information-sharer being asked to bear the burden of public access rather than depriving the public of some good once it has it.

This objection misunderstands the notion of assurance value that supports the right to be forgotten. The value of assurance lies precisely in the fact that it exempts us from certain kinds of responsibility. When I agree to take my forgetting-pill so

that you can tell me your secret, you are freed from the possibility of suffering certain consequences—which is just to say that you do not have to bear the responsibilities that would otherwise come with the information disclosure. If I were to insist, a week after having taken the forgetting-pill, that I have a right to take the forgetting-pill-antidote in order to not forget what you have told me, I would frustrate your interest in maintaining a certain kind of relationship with me. That is, I would harm your privacy interest. Moreover, the availability of this antidote would threaten the general possibility of assurance, and with it the privacy interests that assurance promotes.

I believe it is a mistake to insist in an unqualified way on the importance of a person taking responsibility for her actions. In the context of the disclosure of private information, this may simply mask a different concern, which is that a person, by seemingly seeking to evade responsibility for her decisions, will deprive the public of valuable information. Saying that “actions have consequences” should not justify unnecessarily imposing a burden on someone who disclosed personal information, but it might justify the public’s claim to have access to information when it overrides the interests supporting a withdrawal claim.

V. THE FREE FLOW OF INFORMATION

The right to privacy—like all rights—has limits. As I noted, my right to maintain privacy about the contents of my car’s trunk is limited by the public’s interest in having a police force that can sometimes search cars. The right to privacy is limited by the public’s interest in access to information in other ways. It is, for instance, permissible for reporters to ask prying questions about the lives of politicians, which questions might be inappropriate when directed at non-public individuals. Various critics of the right to be forgotten argue that the doctrine, as currently understood by European courts, sweeps too broadly, transgressing the limits set by the public’s interest in access to information.

No one has made this case more trenchantly than Robert Post, who argues that the demands of a free speech culture place strict limits on right to be forgotten claims.²⁹ He places such strict limits on the right to be forgotten because of the way he understands the countervailing free speech values. The public sphere is a place for dialogue: where speakers speak, audiences listen, and in response members of the audience share their

²⁹ Post 2018. Post allows some right to be forgotten claims, when the material in question is extraordinarily offensive or humiliating (Post 2018, pp. 1007–8, 1054–60.) As my argument indicates, I think this standard is unduly restrictive.

own views.³⁰ This dialogue requires a background of shared information, and shared access to information. To remove information from this background would be to imperil the dialogue itself by undermining the possibility of shared understanding—thus threatening two distinct kinds of value. Open public dialogue is intrinsically valuable when and because it constitutes (a) political deliberation or (b) the development of culture. Assuming that these intrinsically valuable activities are preeminent in any open society, Post concludes that only the most serious of privacy concerns can justify interfering with the information backdrop of social communication.

Let me start by noting that all advocates of the right to be forgotten allow for free speech limits on withdrawal claims.³¹ As I have made clear, my formulation of the right to be forgotten rests on *pro tanto* claims that can be overridden by competing considerations. This prompts the question as to the proper way to conceive of the competing interests at play—more specifically the interests of speakers and their audiences, that is, the public. Post suggests, mistakenly I think, that a huge number of seemingly benign requests for withdrawal will have, as a knock-on effect, a deleterious impact on political and cultural dialogue. I turn now to scrutinizing this claim.

There is an important ambiguity in Post's argument. Like other commentators, Post emphasizes the sheer number of right to be forgotten requests that have been processed by Google in Europe.³² But this raises the question as to whether the threat to the informational fabric of the digital public sphere is an issue of the *volume* of information lost or, instead, an issue of losing specific bits of valuable information. Do these distinct issues represent equally problematic threats to a free speech culture?

While the number of links removed from Google under the European right to be forgotten regime is typically reported in an apocalyptic register, I believe it is a mistake to say members of the public have a claim against the removal of large amounts of information from the internet *per se*. Suppose there is a large group of people who appeared in televised advertisements for a company that, it turns out, relied on slave labor to produce its products. These people may all, individually, want to remove links from search results that appear in response to their name that indicate they appeared in these advertisements, simply because they do not want to risk being associated with a company that, unbeknownst to them, relied on slave labor. If each of these people submits a right to be forgotten request, a large number of links will disappear from search. Of course, people will still be able to find out about the company's use of slave

³⁰ Post 2019.

³¹ As Post (2018, p. 987, note 18) himself notes.

³² See, for example, Post 2018, p. 988.

labor by googling a number of other things—just not by googling these individuals' names. In this case, a large volume of information is retracted from public access, but it does not seem particularly worrisome. It would be overblown to label this a dire threat to the fabric of our intersubjective social world.

This leaves us with the second interpretation of Post's argument, which is that right to be forgotten claims threaten to remove specific, high value links from the public realm. On the one hand, this argument is not particularly interesting—because even defenders of the right to be forgotten concede this! No one believes that the right to be forgotten should empower Donald Trump to remove links from Google that direct to stories about his illicit dealings with a porn star. But, on the other hand, there may still be bite to Post's argument. For it could be argued that seemingly unimportant stories on the internet have an outsized value for cultural and political dialogue. Indeed, in the *Google Spain* case, it might be said that the link Costeja wanted removed from Google concerned a political subject matter, namely the way the Spanish state handles, or handled, social security auctions. Approached this way, it may seem that a huge amount of information that at first blush seemed uninteresting in fact falls under a description that has at least some political or cultural importance. So long as a webpage contains *something* of cultural or political significance, it may be argued, the public has a strong interest in having easy access to a link to that webpage on search and the webpage author has a strong interest in the public's having said access.

This argument fails, though, because the criterion of whether something falls under a politically or culturally relevant description is too weak to justify the strong claims we usually attach to free speech interests. There will be many cases in which a webpage may contain information that is in some sense politically or culturally relevant but still not valuable enough to justify providing a link on search results that could override a person's right to be forgotten claim. Consider an example. A college student might consent to having a blogger write about his participation in an LGBTQ summer sports league, perhaps hoping that this will do something to normalize gay men in sports. Ten years later, the topic of gay men in professional sports might be a national topic of conversation. As a result, the webpage hosting the blog post will be clearly relevant to an ongoing cultural issue. Nevertheless, it strikes me as unreasonable to deny the erstwhile college student's claim to have the link removed from search, given that he may have developed a strong preference for privacy about his sex life.

I believe this argument holds even for some people who become, in some sense, public figures. Suppose that a woman who is now 40 told a group of friends, when she was 20, that she is a child of incest. Now the woman is running to be the governor of Wisconsin, and one of her former friends posts on social media that the woman is the

child of incest. By virtue of running for public office, the woman is, by all accounts, a public political figure, and as a result, information relating to her takes on a political cast. We thus ask: does the woman have a claim against the social media post appearing in search when her name is googled? I believe the woman does have a claim against Google. Indeed, I would go even further and say that the woman has a claim against her former friend, who has the power to delete the social media post in question. The strength of the woman's claim is explained by the fact that both the criteria for a right to be forgotten claim are fulfilled. First, the information is now being circulated among an audience she did not intend, violating her interests in privacy. Second, she plausibly would have felt an inhibition about sharing this information with her friends when she was younger.

These cases illustrate my general conclusion, which is that the mere fact of some information's having political or cultural relevance is not sufficient to justify the claim that the public has a strong interest in the information. More specifically, it is not sufficient to justify a strong enough public interest to override the values of privacy and assurance. Whether a case involves free speech values strong enough to override these interests will, in each case, be a rather complex question. For not every case will involve *both* the values of privacy and assurance, and the strength of the public's interest in knowing about public figures, for instance, will depend on the person in question and the information at stake. Suppose, for example, that a former politician wants to remove from search results a link to a story saying that he attended a strip club while in office. Assuming he was photographed unaware, he will not have an assurance claim—but he will certainly have a privacy interest in this story not being available on Google to, for instance, his grandchildren. However, the public clearly has a very strong interest in accessing stories about how politicians act when in office. In this case, the politician lacks a right to be forgotten claim to the de-listing of the link.

While there are, therefore, limits to the right to be forgotten, we should not assume these limits will be as restrictive as Post's argument implies. Especially in the digitized world, we should be wary of giving too much credence to claims of public interest in access to the private lives of individuals. A claim to withdraw information from search can help shield our privacy interests, and in some cases promote the individual's own free speech interests, without necessarily impeding the free flow of information.

VI. CONCLUSION

The right to privacy protects individuals from the prying eyes of the public. The right to be forgotten is an instance of the right to privacy, and a crucial one in the age of the

internet. Most of us find it unsettling that, because of the internet, our bad decisions or silly mistakes can follow us for our whole lives. What is worse, even seemingly benign decisions can come back to haunt us. A picture or video that a man uploads to a social media website might go viral, launching him into the public eye and possibly garnering him a massive, mean, and relentless audience.³³ If we think that people should have some insulation from such unfortunate eventualities, the right to be forgotten has powerful explanatory value. Without the assurance value supplied by the right to be forgotten, an unpleasant fame may deter many people from sharing information online. This would not only inhibit our experimentation with the boundaries of privacy and free speech, it would also threaten the public good of an open and diverse internet.

ACKNOWLEDGEMENTS

I am extremely grateful to R. Jay Wallace and Niko Kolodny for their extensive, constructive criticism of successive drafts of this paper. I would also like to thank Elek Lane and Milan Mossé for giving written comments on an early draft. I presented that early draft at the Richard Wollheim Society at U.C. Berkeley, where I received much helpful feedback. Finally, I would like to thank two anonymous reviewers for *Political Philosophy*. Responding to their careful remarks greatly improved the paper.

COMPETING INTERESTS

The author declares that he has no competing interests

REFERENCES

- Allen, Anita. 2000. Privacy-as-data control: conceptual, practical, and moral limits of the paradigm. *Connecticut Law Review*, 32: 861–75. https://scholarship.law.upenn.edu/faculty_scholarship/790
- Basu, Rima. 2022. The importance of forgetting. *Episteme*, 19: 471–90. DOI: <https://doi.org/10.1017/epi.2022.36>
- Carnegy-Arbuthnott, Hannah. 2023. Privacy, publicity, and the right to be forgotten. *Journal of Political Philosophy*, 31: 494–516. DOI: <https://doi.org/10.1111/jopp.12308>
- Citron, Danielle Keats. 2015. Protecting sexual privacy in the information age. Pp. 46–54 in *Privacy in the Modern Age: The Search for Solutions*, ed. Marc Rotenberg, Julia Horwitz, and Jeramie Scott. New York: The New Press.

³³ Or consider the man whose photo became the template for the meme titled “the worst person you know” (Kassam 2022). I thank a reviewer for this example.

- Citron, Danielle Keats. 2022. Intimate privacy's protection enables free speech. *Journal of Free Speech Law*, 2: 3–16. <https://www.journaloffreespeechlaw.org/citron.pdf>
- Fried, Charles. 1968. Privacy. *Yale Law Journal*, 77: 475–93. <https://openyls.law.yale.edu/handle/20.500.13051/15184?show=full>
- Hanin, Mark. 2022. Privacy rights forfeiture. *Journal of Ethics and Social Philosophy*, 22: 239–67. DOI: <https://doi.org/10.26556/jesp.v22i2.1633>
- Jones, Meg Leta. 2016. *Ctrl + Z: The Right to Be Forgotten*. New York: New York University Press. DOI: <https://doi.org/10.18574/nyu/9781479801510.001.0001>
- Kassam, Ashifa. 2022. “The worst person you know”: the man who unwittingly became a meme. *Guardian*, 19 June. <https://www.theguardian.com/technology/2022/jun/19/the-worst-person-you-know-the-man-who-unwittingly-became-a-meme>
- Kulk, Stefan and Frederik Zuiderveen Borgesius. 2018. Privacy, freedom of expression, and the right to be forgotten in Europe. Pp. 301–20 in *The Cambridge Handbook of Consumer Privacy*, ed. Evan Selinger, Jules Polonetsky, and Omer Tene. Cambridge: Cambridge University Press. DOI: <https://doi.org/10.1017/9781316831960.018>
- Marmor, Andrei. 2015. What is the right to privacy? *Philosophy & Public Affairs*, 43: 3–26. DOI: <https://doi.org/10.1111/papa.12040>
- Nagel, Thomas. 1998. Concealment and exposure. *Philosophy & Public Affairs*, 27: 3–30. DOI: <https://doi.org/10.1111/j.1088-4963.1998.tb00057.x>
- Nissenbaum, Helen. 2004. Privacy as contextual integrity. *Washington Law Review*, 79: 119–58. <https://digitalcommons.law.uw.edu/wlr/vol79/iss1/10>
- Nissenbaum, Helen. 2009. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Redwood City, CA: Stanford University Press. DOI: <https://doi.org/10.1515/9780804772891>
- Post, Robert. 2018. Data privacy and dignitary privacy: *Google Spain*, the right to be forgotten, and the construction of the public sphere. *Duke Law Journal*, 67: 981–1072. <https://scholarship.law.duke.edu/dlj/vol67/iss5/2>
- Post, Robert. 2019. Privacy, speech, and digital imagination. Pp. 104–21 in *Free Speech and the Digital Age*, ed. Susan J. Brison and Katharine Gelber. Oxford: Oxford University Press. DOI: <https://doi.org/10.1093/oso/9780190883591.001.0001>
- Rachels, James. 1975. Why privacy is important. *Philosophy & Public Affairs*, 4: 323–33. <https://www.jstor.org/stable/2265077>
- Rumbold, Benedict and James Wilson. 2019. Privacy rights and public information. *Journal of Political Philosophy*, 27: 3–25. DOI: <https://doi.org/10.1111/jopp.12158>
- Scanlon, Thomas. 1975. Thomson on privacy. *Philosophy & Public Affairs*, 4: 315–22. <https://www.jstor.org/stable/2265076>
- Shiffrin, Seana. 2011. A thinker-based approach to freedom of speech. *Constitutional Commentary*, 27: 283–307. <https://hdl.handle.net/11299/163435>
- Thomson, Judith Jarvis. 1975. The right to privacy. *Philosophy & Public Affairs*, 4: 295–314. <https://www.jstor.org/stable/2265075>

