



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/203327/>

Version: Published Version

Article:

Carnegy-Arbuthnott, Hannah Maeve (2023) Privacy, Publicity, and the Right to Be Forgotten. *Journal of Political Philosophy*. ISSN: 1467-9760

<https://doi.org/10.1111/jopp.12308>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Privacy, Publicity, and the Right to Be Forgotten

Hannah Carnegie-Arbuthnott 

Department of Philosophy, University of York, Heslington, York, UK

Correspondence: Hannah Carnegie-Arbuthnott, Department of Philosophy, University of York, Heslington, York YO1 5DD, UK.

Email: hannah.carnegie-arbuthnott@york.ac.uk

Many of us will remember the sense of foreboding induced by the simple threat, usually uttered by a teacher: ‘It will go on your permanent record’. This administrative bogeyman exploits our early awareness of the importance of being able to leave some things in the past. While some of the things we do may go on public record at various points in our lives, it used to be possible to comfort ourselves with the thought that these would soon be buried deep in the archives, where most people would not care to look. In the digital age, however, search engines do the digging and can serve up previously long-forgotten results on a simple search of a person's name.

This article takes up the question of what kinds of claims we have against information being dug up from our past. Specifically, it focuses on information that is legitimately a matter of public record. When some information has been stored in an archive and has, in principle, been accessible to anyone, what claims, if any, do we have against that information being brought back to light? Many people would find it intrusive for someone to dig through archives and publicize afresh the information they found out about you there, but it is not clear what the basis of such a complaint would be. Moreover, it is this kind of complaint that is provided protection through data-protection provisions colloquially referred to as the right to be forgotten. Such measures are typically framed as privacy protections. While some have argued that we can have privacy rights over information that has been made public,¹ I argue that appeals to privacy fail in cases concerning information that is legitimately a matter of public record.

Paying attention to the reasons we have to object to the dredging up of outdated information reveals a new category of claims that are distinct from claims to privacy, but serve the same general interest in self-presentation that privacy scholars have long been concerned with. I call these claims against distortion. We

¹Rumbold and Wilson 2019.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Author. *The Journal of Political Philosophy* published by John Wiley & Sons Ltd.

can understand such claims as falling under a general principle of reputational control embedded in the historic right of personality. That general principle provides the basis for claims to privacy, claims against defamation, and, as I will argue, claims against distortion.

The purpose of invoking the right of personality is to identify a general principle embedded in that concept, and to then draw out a taxonomy of claims that serve that principle in distinct ways. One advantage of separating these three categories of claims is that it allows us to retain clear boundaries around the concept of privacy, which has often come under fire for suffering from a plethora of meanings.²

On my account, claims to privacy are claims to prevent certain information from becoming publicly available beyond one's audience of choice. Privacy provides one aspect of reputational control by allowing us to decide who we share various aspects of oneself with. Claims against defamation are claims against having false statements made about oneself in a way that is detrimental to one's reputation. Claims against distortion are claims against having one's public profile distorted in ways that involve the dissemination of true information. Such distortion can happen 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.

Identifying the category of claims against distortion and analysing its place alongside claims to privacy and claims against defamation allows us to broaden the scope of the justificatory basis of data-protection provisions. In doing so, it widens the scope of debates about when and why we have reason to give individuals more control over their personal information. Data-protection provisions have tended to be framed and defended as measures which protect the privacy rights of individuals. There is an important place for data-protection regulations which serve that purpose. However, not all data-protection regulations can be understood or justified as privacy protections. Trying to analyse all data protections through the lens of privacy therefore unduly constrains our thinking about how such regulation may be justified on the basis of providing important protections to the individual.

The main argument proceeds by examining the data-protection provisions known as the 'right to be forgotten'. Specifically, it addresses the aspect of this regulation that gives individuals claims to the erasure of certain results that are displayed when their name is entered into a search engine. The aspect of the right to be forgotten that requires explaining is the rationale for erasing search-engine results that point to some content, while leaving the same content publicly available at the original source. Privacy-based approaches fail to adequately explain the purpose that is served by this aspect of the right to forgotten, as it has been implemented under the EU's General Data Protection Regulation (GDPR). This is because what is at stake in the case of the curation of a person's search profile is not always what information is publicly available, but how that information is framed or contextualized. When a

²For discussion of this line of criticism, see Solove 2006, pp. 479–82.

search engine serves outdated or irrelevant search results for a person's name, this can amount to a distortion of their public profile. We have reason to care about these practices, and give individuals claims against them, but such claims are not based in a right to privacy. They are better understood as claims against distortion.

While the right to be forgotten is a relatively narrow aspect of data-protection regulation, it shares a feature that is common to many concerns about our current situation with respect to the collection and use of personal data. It is a case in which the information in question has already been made public, and either cannot or should not be made entirely private again. And yet, concerns remain about the erosion of the individual's control over how their personal information is presented to others. Under these circumstances, we need to find a way to shut the stable door after the horse has already bolted. Claims against distortion provide a normative basis for regulations aimed at providing such control.

I | MOVING BEYOND PRIVACY

What is colloquially referred to as the right to be forgotten is not a general right as such, but a set of provisions laid out in the GDPR.³ Article 17 of the GDPR sets out a 'right to erasure', which gives data subjects the right to obtain from data controllers the erasure of personal data concerning the data subject. When relevant criteria are met, such as when the data is no longer necessary to the purposes for which it was collected, the data controller must respond to erasure requests and delete the data in question without undue delay.

One of the controversial aspects of this legislation is that search-engine providers such as Google are considered to be data controllers, and search-engine results are classed as personal data under certain circumstances.⁴ This means that an individual can request the removal of certain search results that are displayed when her name is typed into Google as a search term, and Google is obliged to remove them if the relevant conditions set out in Article 17 of the GDPR apply.

This aspect of the right to be forgotten has come under intense criticism from search-engine operators, media commentators, and scholars alike.⁵ Critics have argued that such measures go too far in curtailing freedom of information, especially when the search results in question pertain to information that was lawfully made public. It has been suggested that allowing individuals to make private information that has already been in the public domain, and indeed remains publicly accessible through other means, is futile at best, and at worst based on confused or incoherent premises.⁶ Finally, the way in which the right has been implemented has been criticized almost to the point of ridicule for being counterproductive on its own terms.

³Though a general right to be forgotten has at times been mooted as a possible new human right, no such general right has been enshrined in EU, UK, or US law. See e.g. Weber 2011.

⁴Case C-131/12, *Google Spain SL and Google Inc. v. Agencia Española de Protección de Datos (AEPD) and Mario Costeja González* (2014). Hereinafter referred to as *Google Spain*.

⁵For a particularly scathing critique, see Dans 2019.

⁶See Post 2018.

The case that formed the basis of the landmark ruling on the matter by the Court of Justice of the European Union (CJEU) is the prime target of those charges of inconsistency. The plaintiff in the case, Mario Costeja González, had requested the erasure of details of a foreclosure notice issued against him, which had been published on the web pages of the newspaper *La Vanguardia*. The court decided that the link to the original newspaper record that featured as a search result on Google when Costeja's name was entered as a search term should be erased. However, *La Vanguardia* was not required to remove the original article from its website. This inconsistency in the application of the rules of erasure between the original source and the Google search-results page is a principal bone of contention in analyses of the ruling.

Moreover, in an ironic twist, the case became subject to a phenomenon known as the Streisand effect—when the attempt to censor or protect private information results in it becoming much more widely publicized.⁷ The landmark ruling was widely reported. As a result, now when you enter Costeja's name as a search term on Google, the first page of search results is entirely populated with links to reports about the ruling, which contain reference to the facts of the case, including mention of the original foreclosure notice Costeja was seeking to have 'forgotten'. Because the facts of the legal case are a relevant and sufficiently weighty matter of public interest, it has not been deemed appropriate for these search results to be erased, but nor has the judgment about the erasure of the original link to the foreclosure notice been reversed.

Defending these aspects of the ruling through the lens of privacy presents a challenge, given that the information in question remains in the public domain via the original source, and has subsequently been amplified through the Streisand effect. I will argue that the ruling can nevertheless be interpreted as serving an important interest in self-presentation if we understand it as based on a claim against distortion, rather than a claim to privacy. This interpretation allows us to explain away the seeming inconsistency in treating the original source and the Google link differently with respect to erasure. While the Streisand effect is likely to be a rare side-effect particular to the landmark case, analysing this aspect of the case helps to illustrate the rationale for treating search-result links differently from source material. It sheds light on the importance of how information is framed in a given context, over and above the question of whether some piece of information is in the public domain. This idea is central to the case for claims against distortion.

Before making the positive case for a distortion-based justification for the right to be forgotten, further argument is required to motivate the move beyond privacy-based justifications. Several scholars have made the case in favour of the right to be

⁷The term was coined after a case in which Barbra Streisand sought to restrict the online publication of photos of her mansion. Streisand's attempts were widely publicized, resulting in increased traffic to the photos in question. For an analysis of the effect, see Jansen and Martin 2015.

forgotten by drawing on the idea that the forgetting of information over time has long been an important aspect of our cultural practices. Such gradual forgetting, it is argued, has been crucial to our ability as individuals to lead adequately autonomous lives, unencumbered by mistakes or embarrassing mishaps from our past. From this perspective, the widespread accessibility of old information made possible in the digital age represents a radical and destabilizing shift, which threatens our ability to live our lives on our own terms as the years go by.⁸

Proposed solutions to this problem have focused on the principles and pragmatics of how to strike the right balance between the protection of individual privacy on the one hand, and the public interest in information on the other, once we've recognized that the weighting of this balance might shift over time. Viktor Mayer-Schönberger, for example, argues that information should have an expiry date after which it automatically gets deleted.⁹ Less radically, Giovanni Sartor makes the case that regulators should direct data controllers as to when the balance has shifted such that information warrants deletion.¹⁰ On the basis of an argument from moral autonomy, Meg Leta Jones has argued that we should move towards a model of information stewardship for individual users, allowing them ultimate control over their online profiles.¹¹

I share the general concern that the widespread accessibility of outdated personal information may lead to a stifling social environment in which people are constantly reminded of and held back by events from their past. I am also sympathetic to the idea that individuals should have more control over their personal information profiles. However, to the extent that the focus in this domain has remained within the framework of assessing whether information in various cases should be 'public' (accessible) or 'private' (erased), this framing of the debate still implies a univocal approach to erasure across different information sources.

The univocal implications of the privacy framing are demonstrated by Robert Post in his critique of the *Google Spain* ruling.¹² Post argues that the decision conflated two conflicting senses of the right to be forgotten, one based on a bureaucratic principle of data protection, and the other on the notion of dignitary privacy. According to Post, the bureaucratic principle assumes that data are always collected for instrumentally specified purposes, and warrant erasure when no longer required for those purposes. Dignitary privacy rights, by contrast, seek to 'define and enforce social norms of respectful expression'.¹³ Protecting dignitary privacy requires a balance to be struck between protecting the dignity of the individual against the public interest in freedom of expression. Post points out that the public interest extends broadly to information that 'can or will become part of the agenda of public action',

⁸See Floridi 2015; Jones 2018; Mayer-Schönberger 2011; Pagallo and Durante 2014.

⁹Mayer-Schönberger 2011.

¹⁰Sartor 2014.

¹¹Jones 2012.

¹²Post 2018.

¹³*Ibid.*, pp. 991–2.

as well as to ‘maintaining the integrity of the structure of communication that makes public discourse possible’.¹⁴ Post argues that Google forms part of that structure of communication. The ruling therefore should have followed a principle of dignitary privacy. According to a principle of dignitary privacy, the standard by which to assess whether some information ought to be accessible to the public is the extent to which that information is deemed offensive or newsworthy. If deemed sufficiently offensive to cause harm to the dignity of the individual, the individual has a privacy claim against its publication.

In the context of dignitary privacy, according to Post, ‘what matters is not the content of information, but the context and meaning of particular communicative acts’.¹⁵ A key tenet of Post’s argument is that Google links carry no further meaning over and above the content of the underlying web pages to which they point. As a result, they cannot be deemed any more or less offensive or newsworthy than that original content. On Post’s account, the court should therefore have applied the same outcome to Google search results as to the underlying websites to which they point.

Post’s point about the contextual meaning of communicative acts is an important one. He interprets the significance of context in terms of the appropriateness of the audience to which the information is exposed, as evidenced by the following example: ‘It may be humiliating to reveal the CT scan of a private person to the general public, but unobjectionable to provide it to a doctor for evaluation’.¹⁶

I will argue in Section III that search-result links do carry meaning over and above the content to which they point, and that this gives us reason in some cases to treat them differently from that underlying content. However, the way in which they convey such contextual meaning is not by virtue of reaching a different or wider audience than the original source material. The contextual meaning of such communicative acts can thus fall foul of standards of offensiveness and give rise to concerns about dignitary harms in a way that is not adequately captured by the concept of privacy. The reason for this is that the concept of privacy is constrained to concerns about the appropriateness of the audience to which some information is exposed. Instead of expanding the concept of privacy, we have good reason to keep it so bounded, and instead expand the category of claims that function in the service of protecting individuals’ interests in self-presentation.

In order to bring out the way in which the concept of privacy turns on the question of audience exposure, it will be instructive to consider key conceptions of privacy. Conceptions of privacy standardly invoke a distinction between public and private spheres.¹⁷ On this general view, it is assumed that a person’s privacy interest

¹⁴Ibid., p. 1015.

¹⁵Ibid., p. 1055.

¹⁶Ibid.

¹⁷Not all privacy scholars use the language of ‘spheres’ in this way. One reason for this is that such language risks oversimplifying the picture by implying that there are two clear and distinct spheres that apply universally, one ‘public’ and one ‘private’. However, as I demonstrate below, varying conceptions of privacy share the underlying structuring assumption that privacy is about ensuring that information is only shared in those spheres in which it is appropriate for it to be shared, while recognizing that the boundaries of those spheres will vary from one context to another.

is frustrated whenever some information that ought to be restricted to a specific audience crosses over into a more public domain. For example, in Judith Jarvis Thomson's influential view of privacy rights, once a person has made some information public, whether intentionally or unintentionally, they are considered to have thereby waived their privacy right over that information.¹⁸ It is this feature of approaches to privacy that makes it difficult to make sense of the outcome in *Google Spain*, if we interpret the data-protection measures in question as aimed at protecting a privacy right.

This general structuring idea underpins a range of approaches to privacy: privacy as control, privacy as accessibility, and privacy as contextual integrity. The conception of privacy as control paints a picture of the individual as sovereign over her personal domain, with her consent required to let others in, at her own discretion.¹⁹ In doing so, it most directly invokes a sharp distinction between the private sphere of the individual as a domain that requires protection against unwanted encroachment from the public.

The conception of privacy as accessibility departs from the notion that privacy is best served when the individual has ultimate control over the flow of information between the private and public spheres. Arguments in this camp point to the fact that the value of privacy can be undermined when individuals willingly open the floodgates, especially if the resulting dissolution of their private sphere is irreversible.²⁰ This represents a more paternalistic stance towards protecting privacy, but one nevertheless based on the premise that some things should be kept within a private sphere protected from public scrutiny.

The third approach, privacy as contextual integrity, more directly challenges the simple picture of private and public as two distinct spheres. Instead, it emphasizes privacy as a right to an appropriate flow of information, with different information being appropriately known by different parties in different social contexts.²¹ This approach rejects the idea that the dividing line of privacy is between activities that are carried out in private spheres such as the home or within the family, and those carried out in public. It does so to make space for the idea that certain ways of collecting personal information can constitute violations of privacy even where the individuals in question are acting in public. However, in positing the idea of an appropriate flow of information, it nevertheless draws on the same general structuring idea of different spheres of accessibility of information. Each distinct social context in which some information is appropriately known can thus be thought of as its

¹⁸Thomson 1975, pp. 301–2.

¹⁹See e.g. Brandeis and Warren (1890), who argue that the idea of privacy was already part of the common-law principles behind the protection of one's home as one's castle. They paint a picture of privacy rights by analogy to the idea of an individual in his castle, exerting control over the drawbridge to let others in or keep them out. This has, of course, led to debates about whether privacy is reducible to other rights such as property. On that point, see Marmor 2014; Scanlon 1975; Thomson 1975. For a critique of the privacy-as-control approach, see Allen 1999.

²⁰Allen 2011.

²¹See Nissenbaum 1998, 2004, 2015, 2020.

own restricted sphere as against the rest of the public who are not included in that particular context.

While each of these approaches to privacy offers its own take on how best to think about the divide between public and private in any given context, they all share the idea that some privacy interest is frustrated when information becomes publicized in a context in which it ought to be restricted to a smaller or different group of people. One response to this might be to seek to expand the concept of privacy beyond this framing, in order to capture concerns about the ways in which different sources present information to us, over and above the question of the audiences to which they make it available. Such concerns about the contextual presentation of information ultimately explain what is at stake in cases like *Google Spain*.

However, expanding the concept of privacy to attempt to capture these concerns risks undermining its coherence, leaving it open to the critique that it is an ill-defined term used to cover a range of disparate interests.²² This gives us good reason to keep the concept of privacy bounded and specific, instead of expanding it to capture all the situations that give rise to concerns about reputational control. In keeping the concept of privacy bounded to the question of audience exposure, I follow Carissa Véliz's description of privacy as

the quality of having one's personal information and 'sensorial space' unaccessed. You have privacy with respect to a certain person to the extent that that person does not know anything personal about you, and to the extent they cannot see, hear, or touch you in contexts in which people do not commonly want to be the object of others' attention.²³

Where we identify concerns about protecting an interest in self-presentation which stretch beyond what the concept of privacy can capture, these are better captured by distinct categories of claims.²⁴

This general structuring framework of the concept of privacy underpins the contention that the ruling in *Google Spain* was inconsistent, if based on a privacy claim. If privacy is about keeping information restricted to appropriately restricted spheres, then it is hard to see how one could justify treating Google search results differently from the original source material they point to, when that original source material is in principle publicly accessible to the same audience. It would be an even harder task to justify the differential treatment of Google links pointing to different sources, when those sources contain the same informational content, as happened with respect to the links to newspaper articles

²²For discussion of such critiques, see Solove 2015, p. 479.

²³Véliz 2022, pp. 35–6.

²⁴Miranda Fricker (2017, p. 53) makes a similar argument with respect to keeping the concept of epistemic injustice bounded and specific.

reporting the details of the *Google Spain* case compared to the links to the original foreclosure notice in *La Vanguardia*. If we cannot justify these elements of the right to be forgotten by appeal to the importance of privacy, we need to find another way to justify the measures implemented in *Google Spain* or accept that the case was dealt with inconsistently.

A separate question that arises in response to understanding privacy as functioning to keep information restricted to appropriate spheres is whether we can have claims to privacy over information that has already been made public. Benedict Rumbold and James Wilson have argued that people can retain a right to privacy over information which has been made public.²⁵ They argue that one infringes another person's privacy right when one deduces some piece of information about that person that they did not intend to reveal, from information which they did intentionally make public. While it may be possible to apply privacy rights in this way to information that has been made public in the specific way identified by Rumbold and Wilson, the cases relevant to the right to be forgotten do not fit that mould. Instead, they are cases where some piece of information has intentionally been made public, either willingly by the data subject in the first instance, or by virtue of it being a legitimate matter of public record. The particular challenge this article takes up is whether there is any justification for treating different sources for reporting that information differently with respect to erasure. Rumbold and Wilson's argument thus does not apply to the features of the right to be forgotten that are of central concern to this article.

An alternative path to justifying the right to be forgotten can be drawn from the principles underpinning the more general concept of the right of personality. To the extent that this alternative justification stands up to scrutiny, it extends the theoretical tools we can draw on to defend various aspects of data-protection provisions beyond the constraints of what privacy-based arguments alone can offer.

The argument is not meant to assume that the measures applied in *Google Spain* were intuitively correct, and that there must be an explanation for their consistency. Rather, the aim is to show that there is a plausible path to justifying data-protection measures which need not draw on a right to privacy, and provides a more coherent way to make sense of the measures applied in *Google Spain*.

II | DISTINCT THREATS TO REPUTATIONAL CONTROL

The origins of a right to be forgotten pertaining to personal data can be traced back at least as far as a legislative project developed in France in 2010.²⁶ The French legal concept of a *droit à l'oubli*, sometimes translated rather dramatically into English as a right to oblivion, has historical precedent in setting standards for the way in which spent criminal convictions are expunged from the public record after a certain

²⁵Rumbold and Wilson 2019.

²⁶Weber 2011 The project in question was the Charte du droit à l'oubli dans les sites collaboratifs et les moteurs de recherche, 13 Oct. 2010.

amount of time has elapsed. Within the legal tradition of continental Europe, the right to be forgotten stems from the right of personality, which has historically been conceived of as protecting the dignity, honour, and right to private life of individuals.²⁷

As Rolf Weber points out, the right of personality is a right that incorporates a number of different concepts and terminologies, and is mainly intended to protect the moral and legal integrity of the person. It delineates a sphere of privacy as part of that goal. It is in this context that the right to be forgotten and the way it has been enshrined in the provisions of the GDPR has been interpreted and implemented as a privacy right.

However, as argued above, the framing of the right to be forgotten as a privacy right constrains the scope of data-protection provisions that can be included in such a right. In order to provide a better justificatory basis for those data-protection provisions, we need to shift focus away from a narrow framing around privacy, and instead lean on some of the other concepts invoked by the right of personality to motivate and justify the claims in question.

In doing so, the aim is not to provide a historical account of the legal origins of these existing data-protection provisions. Rather, it is to provide a philosophical analysis of the principles already embedded in the right of personality which can provide foundations for the right to be forgotten, in a way that provides a path to justifying a broader range of data-protection provisions under that umbrella than a focus on privacy alone. More specifically, the task will be to outline three different categories of claims which serve to protect against different threats to reputational control. These are claims to privacy, claims against defamation, and claims against distortion.

The concepts of dignity and honour embedded in the right of personality are broad ones, as is the goal of protecting the 'moral and legal integrity of the person'. When it comes to regulating the sharing of information, what is at stake is the integrity of a person's public persona. The key battleground over the integrity of one's public persona is located in the gap between how other people see, understand, and talk about us, compared to how we would like to present ourselves to others. The question this raises is how much control over the shaping of one's reputation to leave in the hands of individuals, and how much to allow to be shaped by others.

This idea of reputational control embedded in the more general right of personality has been invoked in support of privacy protections. Consider the following arguments that have been offered in defence of the importance of privacy. As we saw, Warren and Brandeis argued that the right to privacy was based on a principle of inviolate personality, and was required to secure peace of mind for the individual.²⁸ Contemporary accounts have argued that privacy rights are required to enable valuable social relationships, allowing us to regulate the boundaries between intimate relationships and more distanced ones. Relatedly, appeals are made to our interest in

²⁷Weber 2011, p. 121.

²⁸Brandeis and Warren 1890.

having a reasonable degree of control over how one presents to others, or to make decisions about one's own life, protected from the interference and scrutiny of others, with the suggestion that these are key aspects of living an autonomous life.²⁹

A common thread in the literature on privacy is thus the idea that privacy protects an interest in self-presentation. Marmor, for example, argues that the distinctive thing about privacy rights is that they safeguard our interest in 'having a reasonable measure of control over ways in which we present ourselves to others and the ability to present different aspects of ourselves, and what is ours, to different people'.³⁰ The second part of Marmor's formulation of the interest reveals again the general framing of privacy in terms of the appropriateness of the audience to which some information is exposed. Marmor explains the distinctiveness of privacy by reference to the importance of being able to conduct intimate relationships, and the formulation of the second part of the interest is key to that case.

However, if we focus on the more general interest in having control over how we present to others (as expressed in the first half of Marmor's formulation), we can notice that privacy plays a relatively narrow role in safeguarding that interest. The distinct way in which it does so is by giving us a reasonable measure of control over which audiences we share different aspects of our lives with. That more general interest in self-presentation, however, can be undermined by ways of sharing information which do not infringe any privacy claims. Some ways of presenting information can constitute a claim-infringing distortion of one's reputation, even where the content of the information in question is already legitimately in the public domain.

Given the focus on the interest in self-presentation, it is no surprise that scholars have been concerned with different ways of distorting or manipulating a person's public presentation, and have sought to capture these under the umbrella of privacy rights. Daniel Solove, for example, includes distortion in his taxonomy of privacy.³¹ Solove distinguishes distortion from disclosure as follows:

Distortion, like disclosure, involves the spreading of information that affects the way society views a person. Both distortion and disclosure can result in embarrassment, humiliation, stigma, and reputational harm. They both involve the ability to control information about oneself and to have some limited dominion over the way one is viewed by society. Distortion differs from disclosure, however, because with distortion, the information revealed is false and misleading.³²

Distortion, on Solove's account, includes cases of defamation. Both are included in his taxonomy of privacy on the basis that privacy is ultimately grounded in the importance of protecting the dignity and honour of individuals.

²⁹See Gerstein 1978; Marmor 2014; Nissenbaum 2015; Rachels 1975; Roessler 2006; Roessler and Mokrosinska 2013.

³⁰Marmor 2014, p. 7.

³¹Solove 2006.

³²*Ibid.*, p. 550.

However, as argued above, expanding the concept of privacy to include distortion and defamation risks playing into the hands of critics who complain that privacy suffers from an embarrassment of meanings. Having identified a general concern to protect the dignity and honour of the individual, we are better off identifying the distinct categories of claims which serve that aim by protecting against different types of threats to one's reputation. We can do this by separating claims to privacy, claims against defamation, and claims against distortion. These claims that come under the general right of personality can be thought of as institutional provisions which help to construct the juridical concept of the individual. They play a role in delineating the boundaries of the domain of individual authority as against the claims of the collective with respect to the shaping of our public personas. They do so by providing distinct elements of reputational control to individuals.

While Solove defines distortion as involving the dissemination of false or misleading information, I reserve the term distortion to refer to cases that involve the dissemination of true information in ways that nevertheless distort a person's public information profile. To make the case for claims against distortion, I provide an analysis of the way in which various ways of distorting a person's public information profile are distinct from defamation, but function to undermine the same general interest which claims against defamation serve to protect. An analysis of the function and significance of search-engine results sheds light on the importance of claims against distortion and the types of context in which they arise.

We can think of the search results that are served when someone types your name into a search engine as being an aspect of your public information profile. A person's online search profile forms an increasingly significant part of their public profile, in a way that can affect important parts of one's life. Employers and universities use search engines to research and assess prospective candidates, to an extent that has led to a burgeoning industry in online reputation-management services aimed at keeping clients' top search results appropriate to the professional image they wish to project.³³ The same goes for assessing the suitability of a prospective romantic partner.³⁴ 'Googling' other people has become an accepted social practice in contexts where we seek to make an assessment of someone's competence or character, and we take search engines such as Google to be useful tools in helping us to make some of those initial assessments.

Not only do our social practices reflect this, but Google's own strategy for delivering search results is explicitly to deliver a ranking by relevance to the search term in question. The results are also intended to be tailored to the purposes of the searcher, based on the profile that Google has built up about them through their other internet activity.³⁵ The results that are served by a search engine operator when

³³See Jones 2012.

³⁴According to one recent global survey of 1500 people, 50% of users of dating apps or websites had searched for someone they had met on a dating app; McGowan 2023. A 2021 survey of online dating users in the US found that 40% had searched a prospective date's name online; Dixon 2022. Another survey from 2018 found that 77% of single Americans research prospective dates on a regular basis; JDP 2018.

³⁵See [Google.com](https://www.google.com) for its public-facing description of its services. For a more detailed explanation of how Google produces a ranking by relevance, see Stuart 2013, pp. 473–6.

a person's name is entered as a search term are therefore not a neutral reference index. They form a curated profile of the search subject deemed by the search engine operator's algorithm to be most relevant to the purposes of the searcher. As such, this profile carries communicative content over and above the content of the information it points to.

In light of this, I suggest that when outdated or irrelevant results are featured prominently in the list of content served by a search engine when a person's name is used as a search term (that is, featured as highly relevant results), this can amount to a distortion of that person's public information profile. It is a distortion insofar as outdated or irrelevant results are presented as relevant to the searcher's assessment of the search subject. We have reason to object to this in the same way we would have reason to object to someone digging up some information from an old archive and re-publicizing it. However, further explanation is required as to what the nature of that objection is.

Claims against distortion contrast importantly with claims against defamation. Much as in the case of privacy, however, paying attention to the normative underpinnings of claims against defamation proves instructive in carving out a role for claims against distortion. In contrasting distortion and defamation, the specific conception of defamation I draw on is Arthur Ripstein's. Ripstein argues that the law against defamation serves to protect each person's entitlement that no other person should determine his or her standing in the eyes of others.³⁶ The right is an inherently relational one, setting out the limits of what others may say about you. It arises from the principle that a person should only be held accountable for actions which are properly attributable to them.

The interest in reputation which a right against defamation protects is fundamentally juridical. What Ripstein means by this is that it is based on a normative idea that is inherent in private law as a system of individual responsibility and relational rights. The interest in reputation is not one that is contingently taken up by the law against defamation and weighed against other competing interests. It is baked into the very structure of private law itself. As such, Ripstein argues that the right to one's reputation is 'at the root' of a system of ordered liberty.

From this perspective, the notion of an interest in reputation is not merely about protecting the vanity of individuals or their sense of honour. Rights which protect the interest in reputation serve a more fundamental purpose in carving out the balance between the control that individuals have over their own lives as against the encroachments of other members of society. They play a part in ensuring that the control that individuals have over their own lives is never unduly outweighed by more powerful parties. Claims against distortion of one's public information profile have a part to play in protecting that same interest in reputation. They can do so by setting out the boundaries of reputational control with respect to practices of accountability.

³⁶Ripstein 2016, pp. 185–232.

A clear distinction between the forms of distortion mentioned above and defamation is that the information presented by the search engine in the cases in question is, strictly speaking, attributable to the data subject. However, if we draw on one of the principles underpinning defamation, that people should only be held accountable for things that are properly attributable to them, we can make headway by focusing on how practices of accountability might bear on a person's right to reputation.³⁷

A closer analysis of normative practices for holding people accountable for what they've done will help to flesh out the claim that when outdated or irrelevant search results are featured prominently, that can amount to a distortion of a person's online public profile. It will do so by providing an explanation of why the framing of the information in the context of search results matters, beyond the mere question of whether or not the information contained therein is publicly available.

III | ACCOUNTABILITY

A key feature of the *Google Spain* ruling was that the link to the original foreclosure notice was removed from the search results associated with Costeja's name, while the original source material remained in place. Additionally, as a result of the Streisand effect, other search results which contained the same information presented in a different context were allowed to remain. A key point at issue is therefore whether individuals have claims to control how information about them that is not private, but a matter of public record, is presented.

While privacy pertains to the question of the audience to which some information is made available, and defamation hangs in part on whether some claim about a person is truly attributable to them, claims against distortion, I suggest, hang on the relevance of the contextual presentation of some information to our practices of accountability.

The case for allowing individuals some measure of control over their public information profile builds from the idea that the principles embedded in the right of personality are about ensuring adequate protection of individuals' control over how to present themselves to others. Our systems of reputation are not just about what is known about a person, but also how we are held accountable for what is known about us. Drawing on Nagel, we can invoke the importance that 'no more should be subjected to the demands of public response than is necessary for the requirements of collective life'.³⁸

The question of what is subjected to the demands of public response, and how, goes beyond what is true or false about a person, or whether certain information is private or made public. Our practices of accountability form an important part of the context that individuals must navigate in presenting themselves to others. And the question of what things they can be appropriately held accountable for bears on a person's ability

³⁷For a discussion of the distinction between accountability and attributability, and how practices of accountability involve imposing certain demands or sanctions on others, see Watson 1996.

³⁸Nagel 1998, p. 14.

to navigate valuable social relationships. We therefore have reason to interrogate how presenting outdated information about a person as salient in a given context could undermine or disrupt our practices of accountability in concerning ways.

For the purposes of this article, I'll focus on a particular aspect of our norms around accountability. This is the idea that there exist various social norms around what information about a person it is acceptable to publicly acknowledge in various social contexts, even if that information is common knowledge. There are some situations where everyone knows some piece of information about a person, and everyone knows that everyone knows this, and yet it would not be polite to mention the information in question, nor to expect the other person to answer for it. Our social norms are such that we would not deem it appropriate in such circumstances to hold someone accountable for the thing they are known to have done.

Nagel gives a vivid example of this phenomenon. Two academics, one of whom has just written a cutting review of the other's book, meet at a party. What Nagel calls 'conventions of reticence' in this case require both to avoid the contentious subject, and to engage in small talk on other topics rather than enter into acrimonious debate over the review. Nagel discusses the importance of such conventions of reticence in terms of the public/private divide:

What is allowed to become public and what is kept private in any given transaction will depend on what needs to be taken into collective consideration for the purposes of the transaction and what would, on the contrary, disrupt it if introduced to the public space.³⁹

Nagel presents these as norms to do with the privacy/publicity divide. However, given that the information in such situations is public knowledge, a more fitting reading of conventions of reticence in cases such as this is that they bear on practices of accountability, rather than privacy. That is, they indicate to us what information it is appropriate to demand a response from someone for in a given social context.

When a search-engine operator presents old information as a prominent search result, that implies the relevance of the result to the purposes for which the search operator knows we use its function (including, for example, assessing a prospective date or employee). To the extent that such rankings signal a degree of relevance that cuts against norms of accountability, we can judge that they amount to a distortion of a person's public profile. They do so because they deliver an implicit message that it would be appropriate to hold someone accountable for the information in question. The distortion consists in calling attention to something that it is no longer reasonable to expect a person to answer for, even if that information is a matter of public record.

³⁹Ibid., p. 12.

An important caveat is required here, which is that claims against distortion will only apply in cases where the information in question is of a kind that would in principle be appropriate to hold someone accountable for, and the appropriateness of holding them so accountable has elapsed. This caveat is required to account for a principle which holds that people cannot be morally assessed for things that are due to factors beyond their control. Without taking account of this principle, the category of claims against distortion would extend too far, to cover search results detailing information about a person that is outside of their control, such as eye colour or place of birth.⁴⁰ We can exclude such results from being subject to claims against distortion because such information falls outside the scope of norms of accountability. Because it falls outside the norms of accountability, the inclusion of such information within the search results associated with a person's name will not convey the implicit meaning that this is relevant information by which to hold a person accountable or judge their character, though it might imply relevance to the searcher's personal preferences in selecting a romantic partner, for example.

If there is an important interest in reputational control, we have reason to be wary about practices that increase the social pressure for people to be held to account for things which it is no longer reasonable to hold them accountable for. This gives us some reason to defend the idea that individuals have claims against such distortion of their online public profiles. Claims against distortion serve a similar interest as claims against defamation, but do so in a distinct way, by focusing on questions of accountability, rather than attributability.

Focusing on distortion from the perspective of accountability, as I've suggested here, also provides an alternative explanation of the claims and duties involved in the kinds of cases discussed by Rumbold and Wilson.⁴¹ Their argument about privacy rights over public information has the consequence that there may be many cases in which one infringes a person's privacy right by unintentionally deducing something that a person did not intend to reveal from information which they willingly made public. They note that this is a bullet they are willing to bite. Shifting the focus from claims to privacy to claims against distortion offers a way to avoid that bullet in some of the cases that Rumbold and Wilson address. One would not infringe another person's claim against distortion simply by coming to know the information in question, but only by holding the person to account for it, or by disclosing it to others in certain ways.

Of course, much of the determination of whether a particular ranking of a search result amounts to distortion will come down to an assessment of whether or not it transgresses the social norms of accountability in the relevant context.⁴² To illustrate

⁴⁰Thanks to an anonymous referee for bringing this point to my attention.

⁴¹Rumbold and Wilson 2019.

⁴²This is similar to the structure of claims against defamation. For a successful claim against defamation, it must be the case both that the statement made about the claimant be untrue, and that it be such as to lower the claimant's standing in the eyes of others.

by way of example, I provide a sketch of how this assessment might be applied in the *Google Spain* case.

Take, for example, a person who had Googled Costeja ahead of a first date. It has become completely socially acceptable to conduct a Google search of someone before a date, and perhaps even to admit to the other person that one had done so. It may even be deemed acceptable within current norms of accountability to ask one's date about some of the results found in the course of that search.

On the other hand, if one had gone to the effort of trawling through old newspaper archives to dig up information, our conventions of reticence would probably make it gauche at best to mention that to one's date, or to expect them to answer questions about what had been dug up. From that perspective, we can think of Google's prominent ranking of the outdated information from the newspaper archive as cutting against the social norms by which we take it to be appropriate to hold people accountable in various contexts.

An important complication arises here. Given the fact that it is now acceptable to Google someone before meeting them, we might think that this has already caused a shift in our social norms of accountability, such that it is deemed generally appropriate to bring up even very old information, which is now readily available through search engines. We therefore cannot make a straightforward assessment of whether any given search result is irrelevant by reference to existing social norms, precisely because the set of information presented in search results pages is part of the shifting landscape of those very social norms.

There are two considerations to put forward in response to this point. The first is that there is already a wider cultural conversation over the extent to which it is appropriate for people to be publicly held to account for bad behaviour from their past, in an age where it is increasingly easy to find the skeletons in people's digital wardrobes.⁴³ I take this to be part of a public reckoning around our collective norms of accountability, and the extent to which we have reason to resist too radical a shift in those norms.

While it is overly simplistic to think that we can assess standards of relevance or irrelevance of information by reference to existing social norms of accountability as though those are static, we can nevertheless recognize the importance of making such assessments in light of the standards of accountability we ought to strive for or preserve. Once we recognize the importance of that, it's a short step to proposing that some of our institutional structures and regulatory practices might have a role in seeking to stabilize the norms we have reason to preserve by setting out the boundaries of claims against distortion. One candidate principle to be guided by in the determination of those boundaries is the general liberal principle highlighted by Nagel of ensuring that the demands of collective life do not become too stifling on the individual.

⁴³For discussion on this point, see Jones 2018.

That principle provides a general guide, rather than a precise recipe for determining the boundaries in every case. One aspect to assess is the severity of transgressions of existing norms of accountability, and the extent to which those transgressions encroach on an individual's reputational control. Another aspect will be to assess the effects of existing norms, relative to how much they encroach in general on the ability of individuals to present themselves to others on their own terms. There is a parallel here to the question of how to strike the balance between privacy claims and the public interest in the assessment of what is deemed 'newsworthy'. The determination of newsworthiness relies on a combination of considerations, including the factual question of what is of interest to the public, as well as the normative question of what information the public ought to have access to as a matter of public good.⁴⁴ In assessing the extent to which information-sharing practices risk becoming too stifling to the individual, we'll need to rely on a sociologically informed assessment of the effects of existing accountability norms combined with a normative assessment of how best to balance the public benefit of those information-sharing practices against the importance of securing sufficient reputational control for individuals.

Second, we have reason to raise concerns about differential abilities of individuals to exert control over their public information profiles, depending on the resources at their disposal. While the configuration of results yielded for any given search term is determined by the search-engine operators' algorithms, it is possible for companies or individuals to exert some degree of control over their search results through a process of search-engine optimization. For those who have enough time, resources, and expertise, this allows them to push inconvenient or embarrassing search results down the search-results ranking for a given search term (for example, their name), by ensuring that other, more positive content is ranked more highly. The fact that some people can engage online reputation-management services to manipulate their own search pages in this way, while others lack the resources to do so, gives us further reason to support data-protection regulation that would give everyone more equitable protection for claims against distortion. The effectiveness of such practices, however, does indicate that claims against distortion may be met by requiring search engines to display certain search results less prominently in the list of results, rather than requiring the removal of the search result in all cases.⁴⁵

A broader concern is that search-engine operators have acquired the power not only to distort the public profiles of individuals, but also to shift the very conditions in which we shape our practices of accountability. Giving individuals more control over their public information profiles represents one way to resist such power becoming too concentrated in the hands of a relatively small number of large corporations.

⁴⁴For a discussion of this, see Post 2018, pp. 1057–60.

⁴⁵For a similar argument in favour of demoting, see Stuart 2013.

IV | CONCLUSION

I have argued for a category of claims against distortion that are distinct from claims against privacy and claims against defamation, but serve the same general interest in reputational control. Claims against distortion apply in cases where some true information that is legitimately a matter of public record is communicated in a way that implies that it is relevant information by which to hold a person accountable, in contexts where that information is outdated or irrelevant to existing norms of accountability.

The approach I've outlined provides a case for giving individuals protection for claims against the distortion of their online public profiles, and provides some guidance on how to make the assessment of which types of information presentation amount to distortion. A question remains over how this is supposed to help to resolve the seeming inconsistency in *Google Spain*.

The three things that were supposed to be in tension in that case were (1) that Google was required to remove the link to the original newspaper article from their search results page for Costeja's name; (2) that the original article was allowed to remain on the newspaper's website; and (3) that subsequent search result links to newspaper articles reporting the CJEU ruling, including all the details of the original debt foreclosure, were not required to be removed. The case in favour of extending claims against the distortion of one's public information profile by reference to normative standards of accountability helps us to explain away the tension here.

With respect to the original newspaper article, allowing it to remain on the internet, but not directly accessible via a simple Google search of Costeja's name would ensure that anybody who wanted to seek it out would have to dig quite deep. While that level of investigation may be appropriate or even necessary in certain restricted contexts, it would not be deemed appropriate in many of the social contexts in which it might be normal to Google someone's name (such as dating or deciding whether to invite someone to interview after assessing their job application). Ensuring that such information remains accessible, but relatively hard to find, is compatible with preserving social norms around conventions of reticence and accountability, as outlined above.

As regards the links to newspaper articles reporting the CJEU case being allowed to remain prominently listed, in light of the argument I've proposed about accountability, we can notice that the framing of the information in question will matter to our assessment of its relevance in the search rankings. It matters, from this perspective, that these articles embed the information about the foreclosure in the context of reporting the facts of the *Google Spain* case. That is, they explicitly mention that the original foreclosure notice was subject to a court ruling in which it was deemed to be sufficiently irrelevant or outdated to warrant erasure from the search results associated with Costeja's name.

While this result still allows the information about the original foreclosure notice to be readily accessible in the public sphere, it ensures that it is framed in a way that simultaneously signals the fact that it is not the kind of information that is

appropriate to hold a person accountable for in the many contexts in which we might conduct a quick Google search of that person's name. In that way, this outcome can also be read as a coherent way to protect individuals from having their public information profiles distorted, by taking measures which seek to reinforce or preserve norms of accountability.

A possible objection is that while such measures may indeed project some public message about the relevance of information to what we hold people accountable for, this provides no guarantee that people will refrain from holding us accountable for the things in question. This objection could be pressed most strongly by those who would advocate an analysis of accountability and blame in terms of the emotional responses or reactive attitudes that are warranted in response to a person's wrongdoing.⁴⁶ Even if people recognize that it would not be appropriate to mention a person's particular misdemeanour in a given context, they may nevertheless be holding the relevant attitudes towards them, and in that sense be holding them accountable in their judgements of the person's character. And of course, there would be nothing to stop them from modifying their own behaviour accordingly, for example by cancelling a date.

It is certainly true that we cannot legislate for the attitudes of others. As Ripstein points out in his discussion of the right against defamation, a right to reputation is not a right that others think well of you. It is a right that nobody else should determine your reputation in the eyes of others. Widening our perspective on the right to reputation to take into consideration the importance of accountability over time, alongside the narrower question of attributability, introduces a grey area on the question of what constitutes someone else interfering with or distorting one's reputation. However, engaging in a careful assessment of existing social practices around what we publicly deem it acceptable for people to be held accountable for, and which of these practices ought to be preserved or protected, provides at least some indication of what data-protection legislation might do in service of that goal.

What I've suggested is that it is possible to justify data-protection regulations that give people claims against distortions of their public information profiles, and in doing so aim to stabilize or protect norms of accountability. Doing so goes at least some way to ensuring that individuals retain a reasonable degree of control over their reputation as against entities like search-engine operators. In the face of the increasing influence that tech companies have over how we present to others, this provides an institutional counterbalance to put some of that reputational control back into the hands of individuals.

I hope to have shown that moving beyond the constraints of a focus on privacy rights alone allows us to make a case for a broader range of data-protection provisions. We can do so by appeal to the importance of claims against distortion alongside claims to privacy and claims against defamation. Thinking more carefully about

⁴⁶E.g. those following a broadly Strawsonian approach to accountability. See Strawson [1962] 2008.

how norms of accountability bear on our assessment of what counts as a distortion of someone's public profile allows us to make a case for extending the right to reputation beyond the narrower standard of attributability on which claims against defamation are based. It does so in a way that addresses concerns about the increasing influence that search-engine operators have over our online public profiles. A key advantage of this approach is that it provides a path to defending and implementing measures that go some way to give individuals control over their personal information even in cases where from the perspective of the privacy/publicity divide, the cat is already out of the bag.

ACKNOWLEDGEMENTS

Early versions of this article were presented at the Centre de recherche en éthique in Montreal and the Departments of Philosophy at Queen's University Belfast and the University of York. I am grateful to the audiences at those talks for their invaluable feedback.

FUNDING INFORMATION

None relevant.

CONFLICT OF INTEREST STATEMENT

There are no potential conflicts of interest relevant to this article.

DATA AVAILABILITY STATEMENT

All relevant data are included in the article.

ETHICS STATEMENT

The author declares human ethics approval was not needed for this study.

ORCID

Hannah Carnegie-Arbuthnott  <https://orcid.org/0000-0002-9414-2469>

REFERENCES

- Allen, Anita. 1999. Privacy-as-data control: conceptual, practical, and moral limits of the paradigm. *Connecticut Law Review*, 32, 861–76.
- Allen, Anita. 2011. *Unpopular Privacy: What Must We Hide?* Oxford: Oxford University Press.
- Brandeis, Louis and Samuel Warren. 1890. The right to privacy. *Harvard Law Review*, 4, 193–220.
- Dans, Enrique. 2019. It's time to forget the right to be forgotten. *Forbes*. <<https://www.forbes.com/sites/enriquedans/2019/09/25/its-time-to-forget-the-right-to-be-forgotten/>>.
- Dixon, Stacey Jo. 2022. US user actions before meeting somebody via online dating 2021. Statista. <<https://www.statista.com/statistics/976154/following-actions-completed-adults-prior-meeting-someone-met-dating-app-usa/>>.
- Floridi, Luciano. 2015. The right to be forgotten: a philosophical view. *Jahrbuch für Recht und Ethik/Annual Review of Law and Ethics*, 23, 163–79.
- Fricker, Miranda. 2017. Evolving concepts of epistemic injustice. Pp. 53–60 in I. J. Kidd, J. Medina, and G. Pohlhaus Jr (eds), *The Routledge Handbook of Epistemic Injustice*. London: Routledge.

- Gerstein, Robert S. 1978. Intimacy and privacy. *Ethics*, 89, 76–81.
- Google.com. n.d. Google search: discover how Google search works. <https://www.google.com/intl/en_uk/search/howsearchworks/how-search-works/>.
- Jansen, Sue Curry and Brian Martin. 2015. The Streisand effect and censorship backfire. *International Journal of Communication*, 9, 656–71.
- JDP. 2018. New data reveals trends in modern dating. 11 Dec. <<https://www.jdp.com/blog/study-online-dating-statistics/>>.
- Jones, Meg Leta. 2012. You are what Google says you are: the right to be forgotten and information stewardship. Rochester: Social Science Research Network scholarly paper 2154353. <<https://papers.ssrn.com/abstract=2154353>>.
- Jones, Meg Leta. 2018. *Ctrl+ Z: The Right to Be Forgotten*. New York: NYU Press.
- Marmor, Andrei. 2014. What is the right to privacy? *USC Law Legal Studies Paper*, no. 14-13.
- Mayer-Schönberger, Viktor. 2011. *Delete: The Virtue of Forgetting in the Digital Age*. Princeton: Princeton University Press.
- McGowan, Emma. 2023. Is your online date googling you? <<https://blog.avast.com/internet-research-in-online-dating-avast>>.
- Nagel, Thomas. 1998. Concealment and exposure. *Philosophy and Public Affairs*, 27, 3–30.
- Nissenbaum, Helen. 1998. Protecting privacy in an information age: the problem of privacy in public. *Law and Philosophy*, 17, 559–96.
- Nissenbaum, Helen. 2004. Privacy as contextual integrity. *Washington Law Review*, 79, 119.
- Nissenbaum, Helen. 2015. Respect for context as a benchmark for privacy online: what it is and isn't. Pp. 278–302 in Beate Roessler and Dorota Mokrosinska (eds), *Social Dimensions of Privacy: Interdisciplinary Perspectives*. Cambridge: Cambridge University Press.
- Nissenbaum, Helen. 2020. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Palo Alto: Stanford University Press.
- Pagallo, Ugo and Massimo Durante. 2014. Legal memories and the right to be forgotten. Pp. 17–30 in Luciano Floridi (ed.), *Protection of Information and the Right to Privacy: A New Equilibrium?* Cham: Springer International Publishing.
- Post, Robert C. 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.
- Rachels, James. 1975. Why privacy is important. *Philosophy and Public Affairs*, 4, 323–33.
- Ripstein, Arthur. 2016. *Private Wrongs*. Cambridge, MA: Harvard University Press.
- Roessler, Beate. 2008. New ways of thinking about privacy. Pp. 694–713 in J. S. Dryzek, B. Honig, and A. Phillips (eds), *The Oxford Handbook of Political Theory*. Oxford: Oxford University Press.
- Roessler, Beate and Dorota Mokrosinska. 2013. Privacy and social interaction. *Philosophy and Social Criticism*, 39, 771–91.
- Rumbold, Benedict and James Wilson. 2019. Privacy rights and public information. *Journal of Political Philosophy*, 27, 3–25.
- Sartor, Giovanni. 2014. The right to be forgotten: dynamics of privacy and publicity. Pp. 1–15 in Luciano Floridi (ed.), *Protection of Information and the Right to Privacy: A New Equilibrium?* Cham: Springer International Publishing.
- Scanlon, Thomas. 1975. Thomson on privacy. *Philosophy and Public Affairs*, 4, 315–22.
- Solove, Daniel J. 2006. A taxonomy of privacy. *University of Pennsylvania Law Review*, 154, 477–564.
- Solove, Daniel. 2015. The meaning and value of privacy. Pp. 71–82 in Beate Roessler and Dorota Mokrosinska (eds), *Social Dimensions of Privacy: Interdisciplinary Perspectives*. Cambridge: Cambridge University Press.
- Strawson, Peter. [1962] 2008. Freedom and resentment. *Proceedings of the British Academy*, 48, 187–211.
- Stuart, Allyson Haynes. 2013. Google search results: buried if not forgotten. *North Carolina Journal of Law and Technology*, 15, 463–518.
- Thomson, Judith Jarvis. 1975. The right to privacy. *Philosophy and Public Affairs*, 4, 295–314.
- Véliz, Carissa. 2022. Self-presentation and privacy online. *Journal of Practical Ethics*, 9, <<https://doi.org/10.3998/jpe.2379>>.

Watson, Gary. 1996. Two faces of responsibility. *Philosophical Topics*, 24, 227–48.

Weber, Rolf H. 2011. The right to be forgotten: more than a Pandora's box. *Journal of Intellectual Property, Information Technology and Electronic Commerce Law*, 2, 120–30.

How to cite this article: Carnegie-Arbuthnott, Hannah. 2023. "Privacy, Publicity, and the Right to Be Forgotten." *Journal of Political Philosophy* 00, 1–23. <https://doi.org/10.1111/jopp.12308>.