"The Right to Be Forgotten": a Philosophical View

Luciano Floridi

I. Introduction

On May 13 2014, the Court of Justice of the European Union (CJEU) held that an Internet search engine operator (ISEO), such as Google, is responsible for processing personal information that appears on web pages published by third parties. As a consequence, the CJEU ruled that an ISEO must consider requests from individuals to remove links to (legally available) information online (web pages) that result from a search on their name. Justifications for removal include cases where the search results "appear to be inadequate, irrelevant or no longer relevant, or excessive in relation to those purposes and in the light of the time that has elapsed".[1] Limits to the delinking include "the nature of the information in question and its sensitivity for the data subject's private life and on the interest of the public in having that information, an interest which may vary, in particular, according to the role played by the data subject in public life".[2] If the ISEO rejects the request, the individual may ask relevant authorities to consider the case. In some cases, the ISEO may be ordered to remove the links from search results.

The CJEU's ruling just summarised concerned a Spanish citizen, Mario Costeja González. The events are by now very well known. Mr Costeja was displeased with the fact that Google searches on his name prominently featured an old foreclosure notice. This had been published, under legal requirement, by *La Vanguardia* – a daily newspaper with a large circulation, particularly in Catalonia – in 1998, when a property co-owned by Mr Costeja and Alicia Vargas Cots had been repossessed for debt (see Figure 1). The Spanish Data Protection Authority (the Agencia Española de Protección de Datos, AEPD) had rejected Mr Costeja's request to remove the original archived notice, but had asked Google to remove the referring links from its index. Google had appealed, and the Spanish court had requested guidance from the CJEU. The CJEU accepted Mr Costeja's claim that indexing the notice was irrelevant to Google's purposes as a search engine under the 1995 EU Data Protection Directive. Contrary to the advisory opinion of Advocate General Niilo Jääskinen,[3]

---

[1] (CJEU 13 May 2014), 93. The incompatibility is with the Directive 95/46/EC (EU Data Protection Directive).

[2] (CJEU 13 May 2014), 81.

[3] Opinion of 25 June 2013, ECLI:EU:C:2013:424.

the CJEU ruled that Google was a European data controller with associated responsibilities and so it ordered it to remove the links. Google complied.

The CJEU's ruling triggered an international discussion on the so-called "right to be forgotten".[4] As a direct consequence of the ruling, Google decided to offer a form online,[5] through which users can request the delinking of personal information; and it set up an Advisory Council[6] in order to analyse the implications of the ruling and provide guidelines about its implementation.

Regarding the requests, according to the Transparency Report published by Google,[7] by July 12 2015, the total number of URLs evaluated for removal was 1,027,207, corresponding to a total of 282,407 requests. So far, about 40% of the URLS have been removed. Note that these figures cover about a year of activities. They may seem huge but they are actually tiny when compared to the number of URLs removed by Google for copyright reasons, following the Digital Millennium Copyright Act. Just during the June-July 2015 period, for example, Google dealt with requests concerning almost 40m URLs.[8] The difference is not quantitative but qualitative. Requests about copyright issues are much less problematic and more easily processable, whereas the requests concerning the right to be forgotten are more complex, sensitive, and have to be managed by human experts, not by algorithms, by studying case by case.

As for the Advisory Council, this held several meetings and seven public consultations across Europe between September and November 2014, inviting experts and the public to join the discussion in Madrid, Rome, Paris, Warsaw, Berlin, London, and finally Brussels. It published its recommendations in February 2015.[9] As one of the eight independent members and the only philosopher and ethicist in the Council, I had the privilege of participating in such consultations and in the discussions surrounding them. That experience informs the following pages, in which I synthesise some reflections on what I believe are the most problematic aspects of the debate on "the right to be forgotten".

I wish I had a convincing answer for every question raised by the debate. I do not. But I do hope that some mistakes can be avoided and some strategies dismissed as fruitless, and that a clearer understanding of what exactly is at stake will help to find a more satisfactory solution. I firmly believe that it is crucial to abandon preconceived approaches for or against Google or the CJEU. Too much of the debate is

---

4 Since then, the literature on the topic has exploded. Julia Powles maintains an extensive bibliography online at http://www.cambridge-code.org/googlespain.html.

5 It is available here: https://support.google.com/legal/contact/lr_eudpa?product=websearch.

6 This is the official website of the Advisory Council: https://www.google.com/advisory council/.

7 http://www.google.com/transparencyreport/removals/europeprivacy/?hl=en-US.

8 http://www.google.com/transparencyreport/removals/copyright/?hl=en.

9 All the recordings and the final report are available from https://www.google.com/advisorycouncil/.

still blindly ideological, with socio-legal, ethical or political prejudices being couched as critical and objective evaluations. The fact that the debate overlapped with the European elections did not help to moderate its terms and generated some strumentalization. That journalists with vested interests on either side of the debate or merely in search for some headlines joined the conversation further constrained the space for reasonable exchanges of ideas. All these factors remain a major hurdle towards any constructive and fruitful dialogue.

II. Privacy vs. Freedom of Speech

By ordering the deletion of links to personal information legally published online, the CJEU generated a challenging debate about privacy, freedom of speech, and complex questions arising from the implementation of the ruling. What might have begun as a small and irrelevant legal episode soon became the spark that ignited a lively international debate on how to regulate the availability and accessibility of information online. It is easy to see now that it was a ticking bomb waiting to explode.

The outcome of the current debate is likely to have long-lasting consequences and represent a watershed in the evolution of the infosphere, the informational environment represented by our increasingly hyperconnected world (Floridi 2014a, b). When is it appropriate for a search engine to provide a link to truthful information about a person that a third-party has legally published online? And what is the best way of dealing with each request to remove from search results a link that refers to such personal information? Similar questions may seem merely technical and boring. But they are actually points of collision of two immensely important spheres of interest in our everyday life.

On the one hand, there is the right to privacy (European Convention on Human Rights, art. 8; and Declaration on Human Rights, art. 12). In our case, this means allowing legally available, truthful information about a person to sediment, without being constantly rehearsed, and let bygones be bygones, so that the past might become a springboard for the future, not a stone around its neck. We should "let us not burden our remembrances with a heaviness that's gone" to quote Prospero from Shakespeare's *The Tempest* (Act V, sc. 1), hence the imprecise but very sticky and by now unavoidable label "right to be forgotten", which is at most a metaphorical "right to be delinked" (in more Latin-based languages the so-called "right" sounds even more dramatic, as in French: *droit à l'oubli*; Italian: *diritto all'oblio*; or Spanish: *derecho al olvido*).

On the other hand, there is the right to freedom of speech – also referred to as the right to freedom of opinion and expression – which includes "the right to receive and impart information and ideas without interference by public authority and regardless of frontiers" (European Convention on Human Rights, art. 10). Note that the Declaration on Human Rights, art. 19 is slightly but significantly different for it

also includes the right "to *seek*, receive and impart" information, quite a crucial verb given the context. In our case, this means the freedom to refer to, and access truthful information that is legally available online, hence the so-called "right to information".

Neither the "right to be forgotten" nor "the right to information" is actually a legal right, so both are misnamed and give rise to a conflict by proxy. Both proxies, and the actual rights behind them, are subject to legitimate interpretations, and when such interpretations are in conflict, Europe and the United States dissent on how to reconcile them. All this makes the debate even more "complex" and in need of a "balance", two rather apt terms.

The debate is complex because there are many elements interacting with each other. The actual ruling, with its pro and contra; the role of ISEOs as intermediaries or data controllers; the difference between availability and accessibility of information online; as we just saw, the so-called "rights" (to be forgotten, to information), the real rights behind them (privacy and freedom of speech), and the ways in which they are interpreted on the two sides of the Atlantic; the concepts of relevance and of public interest, both very slippery; the procedural uncertainty about who should decide which links are rightfully removed and who should be informed about it. I shall return to these points below. Here, suffice to note that, because the debate is complex, almost all parties involved seem to think that its resolution requires a balancing act. We saw above that the ruling itself relies on such a crucial concept. However, the superficial agreement on the need for balance probably hides a deeper disagreement on exactly what kind of balance may be needed. Let me use a very simple analogy. Imagine you like French food but I like Italian food. We may find a balance by alternating between going to a French restaurant one day and to an Italian restaurant next time; or by going to a third restaurant that we both like, say Spanish; or perhaps to another restaurant that serves both French and Italian food. We could let a third friend choose… any combination of the former solutions and more complicated ones are possible. The same holds true about the debate on the right to be forgotten. While most parties involved seem keen on finding a balance, few may have the same arrangement in mind.

Finding such an arrangement requires acknowledging the fact that there is a problem to begin with. And this means clearing the ground from a mistaken view that was repeatedly aired during the Council's meetings. Just because privacy and freedom of speech are two fundamental ethical principles, this does not mean that they are always compatible. Sometimes they can be in serious and genuine conflict. The naïve view to the contrary seems to hold that the two principles are like two pillars of a democratic and liberal edifice. Pillars complement and sustain each other; they do not and cannot clash. The trouble is that this is merely a nice metaphor that can at most illustrate the view, but does not justify it. The truth is rather that, like all other ethical principles, sometimes privacy and freedom of speech need to be prioritized or reconciled in different ways. Think of the textbook example of being honest vs. being altruistic. Both are fundamental principles

and yet, if a Nazi asks where an innocent hides, you better lie, *pace* Kant, but when you swear to tell the truth during a fair and legal trial, you better do so, even if this may hurt an innocent. If there were no conflicts between fundamental ethical principles most of the ethical debate would be superfluous. At the same time, there is no fixed rule or mere algorithm to resolve the conflict, and only human intelligence can strike the right balance. The right to be forgotten is precisely a case in which individual privacy and the social right to freedom of speech pull the debate in opposite directions.

Now, a fruitful way of approaching this tension, and hence the need for a balancing act, is to consider what kind of information is in question and, for this, a distinction between availability and accessibility is in order.

III. Availability vs. Accessibility of Information

The CJEU's ruling and the issuing debate do not concern illegal information, such as pedophilic images, which is regulated by other parts of the legal system. The debate in the popular press has often been confused about this point. The ruling concerns information that is legally published and made available online but should not be linked and hence become easily accessible through a non-local search engine. This means that it is perfectly legal to publish here, or on a t-shirt, the text of the foreclosure notice that appeared in *La Vanguardia* (see *Figure 1*).

> ▪ Les dues meitats indivises d'un habitatge al carrer Montseny, 8, propietat de MARIO COSTEJA GONZÁLEZ i ALICIA VARGAS COTS, respectivament. Superfície: 90 m². Càrregues: 8,5 milions de ptes. Tipus de subhasta: 2 milions de ptes. cadascuna de les meitats.

Source: http://hemeroteca.lavanguardia.com/preview/1998/01/19/pagina-23/33842001/pdf.html?search=Costeja%201998.

Figure 1: The Original Text of the Foreclosure Notice Published by *La Vanguardia*

Yet it is no longer legal for Google to provide a link to it. If this seems counterintuitive consider further that the original information remains accessible through Google using other search terms (at the moment of writing, if you search "Alicia Vargas Cots + La Vanguardia" on a European search engine you can still find the link to the information in *Figure 1*). And the information is also available by direct access to the source: it is legal for *La Vanguardia* to maintain a *local* search engine that works only on its own archive and provides a search through which that information can be found. It is only a *non-local* search engine that ought not to make it accessible, and only insofar as the name of the person requesting the delinking is

concerned.[10] This raises further questions about territoriality. Only non-local search engines are affected by the ruling, such as Google or Bing, and yet by its very nature any non-local search engine is such that it can provide access to the available information through versions of itself that are not regional, in this case neither Spanish nor European but American, e.g. google.com. I shall return to this point in section 4. Here, I wish to stress a different aspect that should counterbalance the counterintuitiveness stressed above. In an analogue world, the *availability* ("the information is there") of some information is usually coupled to its *accessibility* ("the information that is there is known to be there")[11] at least in principle. The best way to block the accessibility to a printed book is to burn it, as all dictators have always known in the past, because as long as it is available it is also potentially accessible. In the infosphere, this is no longer the case. Since we obtain information by going first through the gate of its online accessibility, today a new, two-tier approach to information is normal, with the availability of information online (content) being completely detached from its accessibility (link). The decoupling and re-prioritising of links first and contents next mean that removing the former is tantamount to hiding the latter. More and more frequently, if something is inaccessible (not linked or delinked) online in theory it simply becomes unavailable in practice. In the infosphere, the map becomes more important than the territory: erase something from the map and it does not matter whether the territory remains unchanged. This shift has at least one very significant consequence. It is disingenuous both to argue – as many participants in the debate have – that the CJEU's ruling does not affect the available information *and* to argue that an ISEO like Google, which has a virtual monopoly on search in Europe, provides just the equivalent of a library catalogue. For in both cases one is mistaken in reasoning in analogue terms ("the CJEU's ruling is not burning books" vs. "Google is only indexing books") about digital documents. We must realize that linking and delinking are today ontological operations about the territory. We should all accept and agree that the new battleground for the control of information is the map, and that it is there that power is exercised.

Given this premise about our new information cartography, those who publish legal contents should probably have a saying on whether and how it is linked. This is what I shall argue in the next section.

IV. Transparency

We saw that today the indexing by search engines influences the life of the corresponding information: those who control the map control the territory. At what

---

[10] Add to this the technical fact that it is actually the whole pdf that is no longer linked, not just the specific relevant text. If you search for something else available on that pdf, this is not delinked just because the personal information has been delinked.

[11] *Full* (as opposed to *partial*) accessibility means that the information known to be there can also be obtained, but I shall not use this finer distinction here.

stage, if any, and how (from mere notification to co-decision) could publishers – the owners of the territory, so to speak be involved when their legal information about an individual may be delinked? During the consultations of the Advisory Board it emerged quite clearly that publishers would like to be involved in the decisional process about delinking requests as early as possible. For example, before Google decides to delink some information published by *Le Monde*, it should consult the newspaper to check whether the removal of the link clashes against the freedom of the press or the public interest. To put it more technically: notification should be *ex ante*.

Publishers should be fully involved in the evaluation of a delinking request. They should have the right to know about whether someone has requested a search engine to delink some information that they legally published; to be informed about what decision has been taken by the search engine with regard to such a request; and to appeal, if they disagree with the delinking decision. All this applies even more strongly if a global approach to delinking were to be adopted (see section 4). Of course, the risk is that, by informing the publishers, one may enable them to re-publish the same contents in ways that can by-pass the ruling and the delinking decision itself, both of which concern only personal information and hence "name and surname" searches. Yet this is a case in which I would recommend a principled approach. One could certainly implement disincentives, but the fact that publishers may misuse the meta-information about a delinking request is not an argument against their right to know and hence being able to appeal. This will be even more obvious the further we move towards a situation in which not being indexed by a search engine simply means "not being", full stop. In this case too, the report by the Advisory Council has found a fair balance, by recommending Google to follow the good practice of notifying the publishers "to the extent allowed by the law". It is a bit vague, and I would have liked to see an even more incisive position in favour of a full involvement of the publishers throughout the process, but it is a satisfactory compromise.

Getting the publishers on board seems a sensible idea. After all we are talking about personal information that has been legally published. However, even if one crosses this particular bridge, one stumbles upon a new problem. Suppose that the notification has been discussed and the link has been removed and not re-pubished, should the public be informed about such removal? Once again, this may seem a mere technicality, but unpack it and a significant difficulty surfaces. Currently, if one searches for "Mario Costeja" on a European version of Google, e.g. google.co.uk, one is informed that "Some results may have been removed under data protection law in Europe." There follows a link to further information about the ruling. This public notification appears perhaps due to a similar procedure adopted when a search engine removes content that infringes copyright laws. Yet publicising the fact that some information about someone may have been delinked is problematic. After all, being told that an individual may have gone through the effort of asking a search engine to remove a link to some personal information may raise some serious

suspicions, which may damage that individual as much as, and perhaps even more than, the original information itself. So a current solution by Google is to add the disclaimer seen above at the end of the first page of a search in any European-based version of the engine, containing links to personal information, even if that person has made no request and all and all the links are actually available. This is the crucial value of the "may" in "may have been removed" clause. Perhaps they have not, but the user is supposed to remain uncertain. This seems fine: users are no longer able to infer whether someone has actually asked to have some links removed. But there is an ironic twist. Part of the problem is that, if information about a removed link is made public, it can be exploited to create services that can re-link the de-linked material (see for example http://hiddenfromgoogle.afaqtariq.com/). And if no detailed information is provided, the disclaimer may still work as a reminder and a tempting invitation to look for the same personal information using a search engine that is not based in Europe, such as google.com or duckduckgo.com.

Despite its shortcomings, the blank disclaimer seems to be the best compromise. Those who dislike the solution do not enjoy a more palatable alternative. Suppose one is not told that a search page one is viewing has had some links removed. This appears to be more in line with the spirit of the CJEU's ruling and the protection of privacy. The problem is that now one is truly kept in the dark. Not only one does not get the link to some information one may be looking for, one is not even been told that one is not getting it. In mathematical logic this is called negation by failure: if you do not obtain some information, you assume that the information in question is not there to be obtained in the first place. A bit like a fisherman assuming that if he fails to catch any fish in the lake then there is no fish in the lake. This starts smelling very badly, in terms of freedom of speech. Using our previous analogy, now the map is missing some details and it is also hiding the fact that it is missing them. The partial representation of the territory becomes indistinguishable from a more complete one, and it is by default assumed to be equal to it. Well-informed users will be chronically suspicious about what links they are obtaining any time they run a search about someone. So they may end up using a non-European search engine such as google.com by default anyway.

There is no easy solution because the CJEU's decision currently applies only to ISEOs operating in Europe. No coordination among search engines is currently envisaged, either within Europe or internationally. At the moment of writing it still seems that what happens in Europe stays in Europe, despite a ruling in France,[12] the guidelines issued by Article 29 Working Party,[13] and a request by Isabelle Falque-

---

12 See http://www.legalis.net/spip.php?page=jurisprudence-decision&id_article=4291.

13 See http://ec.europa.eu/justice/data-protection/article-29/press-material/press-release/art29_press_material/20141126_wp29_press_release_ecj_de-listing.pdf In particular, on the first page one can read: "The WP29 considers that in order to give full effect to the data subject's rights as defined in the Court's ruling, de-listing decisions must be implemented in such a way that they guarantee the effective and complete protection of data subjects' rights and that EU law cannot be circumvented. In that sense, limiting de-listing to EU domains on the grounds

Pierrotin of the French data protection authority that any de-listing is applied worldwide.[14] This is known as the territoriality issue. It is a thorny problem that deserves its own section.

V. Territoriality

As specified above, the CJEU's ruling currently affects only ISEO insofar as they operate in Europe. The same personal information no longer accessible through European search engines, such as google.co.uk, is still accessible through non-European search engines, such as google.com or DuckDuckGo. The ruling may seem as ineffective as closing the door of the house while leaving all the windows open. Yet the reality is more complicated.

For centuries, roughly since the Peace of Westphalia (1648), political geography has provided jurisprudence with an easy answer to the question of how far a ruling should apply, and that is as far as the national borders within which the legal authority operates. A bit like "my place my rules, your place your rules". It may now seem obvious, but it took a long time and immense suffering to reach such a simple approach. Which is still perfectly fine today, as long as one operates within a physical space. However, the Internet is a logical not a physical space (more on this distinction presently), and the territoriality problem is due to an ontological misalignment between these two spaces, in the following sense.

The CJEU's ruling is based on an offline physical space: the Westphalian system of sovereign states with controllable borders, which has served us so well for so long. In contrast, search engines operate within an online logical space of nodes, links, protocols, resources, services, URLs, interfaces and so forth. Which means that any place is only a click away. This is a new space that we are still learning to manage. You may compare the difference between the two spaces to the difference between the physical size of a chessboard, which may vary considerably, and its logical space, which is always the same structure of 64 squares in eight rows and eight columns arranged in two alternating colours. A physical space can be *constrained* by rules: think about the problem of immigration across physical borders. But a logical space is *constituted* by rules. Think of a pawn, which can move in a given way on the chessboard because the rules say so, not because of some physical possibilities or impediments. The *accessibility* of some personal information online belongs to a logical space. If there are problems in the logical space, this is probably where they should be solved (Floridi 2009, 2013). A ruling that concerns the Inter-

---

that users tend to access search engines via their national domains cannot be considered a sufficient means to satisfactorily guarantee the rights of data subjects according to the ruling. In practice, this means that in any case de-listing should also be effective on all relevant domains, including .com".

[14] See http://www.wsj.com/articles/french-privacy-watchdog-orders-google-to-expand-right-to-be-forgotten-1434098033.

net cannot rely on the old, Westphalian solution. Note that this misalignment of spaces does not cause only problems, it also provides solutions. The non-territoriality of the Internet works wonders for the unobstructed circulation of information. In China, for example, the government has to make a constant and sustained effort to control information online. But the same feature proves awkward if you are trying to implement the right to be forgotten. The report issued by the Advisory Council strikes a fair balance, recommending to implement the delinking policy at the European level: if a request is approved, links are removed from all European version of Google's search engine. Personally, I accepted this decision but I also argued within the Council in favour of a more restricted, nation-based delinking, for the following reasons.

Although users can easily by-pass the current restriction, according to Google very few do, preferring to rely on their local search engines in their languages. Spaniards use google.es, Italians google.it, Germans google.de, and so forth. The so-called "power of default" is enormous. It follows that if Alice, who is French and lives in Paris, asks Google to delink some legally published information about herself, the most effective implementation is to remove the links from Alice's local search engine, namely google.fr. It is mostly useless to remove them also from google.pt because virtually nobody in France will ever care to check information about Alice using the Portuguese version of Google, while the very few who may care will not be deterred by a pan-European delinking anyway. Someone who is determined to find a piece of information about Alice will simply use a search engine not based in Europe or indeed other means. Recall that the information itself is not erased. This is also why the argument that delinking personal information will be an obstacle to journalists' investigations is unpersuasive: investigative journalism is not really about googling a name. Back to the original point, it follows that, if the CJEU's goal was to "slow down" the accessibility of personal information online that an individual in Europe wants to see delinked, this much has been achieved. At least temporarily. There remains the issue stressed in the previous section: Google displays the removal notification at the bottom of the search results page for name searches in Europe. As the ruling gets discussed and people start seeing removal notices at the end of searches about individuals, savvy users may develop new search strategies.

Some have bitten the bullet and argued that all this is correct, but this is precisely why the delinking should be global, that is, applied to all versions of any search engine world-wide. I agree that a global approach to delinking would be more coherent with the ruling. One would be closing the door of the house and all its windows as well. In the case of Google, this would mean delinking the information in question also from its Brazilian version (google.com.br), for example. However, I disagree with the approach itself because global delinking is probably both unfeasible – since countries are unlikely to accept a European ruling without questions – and dangerous, since undemocratic and illiberal countries should have no right to impose a global delinking ban on personal information. Arguments in favour of a

global delinking assume that the debate on the right to be forgotten has been settled and that the CJEU is right while everyone else who may disagree with it is wrong, including the United States, or, in the previous example, Brazil. This is clearly a flawed assumption. The controversy is still unresolved and the current situation is nothing more than a compromise, and a rather messy one at that. Even the CJEU and the Advocate General disagreed among themselves. The ruling by the CJEU is more like a wrong step in the right direction: it is based on a mistaken physicalisation of a logical space, and making such space all-encompassing is not a solution but a worsening of the problem. Since a global delinking would not reject but "absolutise" the Westphalian approach – it would try to make the physical space so inclusive as to ensure that any logical space must fall within it – it would still have to follow the old rule "my place my rules, but your place your rules", but then how could one explain to Brazilians that some legally published personal information online should no longer be indexed in a Brazilian search engine because the CJEU has ruled so? Would the opposite also apply? Could Brazilians appeal? And how could one determine what is of public interest in this or that country? Maybe an investor from Brazil does need to know whether a person has (for example) had some properties repossessed in the past. Finally, consider the following scenario. The day after some global delinking starts being enforced, nothing will stop undemocratic and illiberal places from hosting a search engine that provides links to all information anyway. It would be ironic if we were to find information using a search engine based in North Korea because it were more complete than the local ones. Geographical space is no longer the solution, so the approach recommended by the Advisory Council is a good compromise, which adapts an outdated answer to a new question. It does not work very well but it is the classic "better than nothing" solution. Opting for a global delinking would be, instead, the classic "perfect is the enemy of good". It would be just another way of demanding that the world adapts to European decisions. Some people may see it as a form of digital neo-colonialism.

When it came to finalising the text of the report, I was happy, pragmatically, to concede the point about a pan-European delinking because this simply adds nothing to a national one, in terms of effective protection of individuals. At the moment, a national or European approach may still be preferable. Google already deals with comparable cases of defamation complaints on a national basis only. So it might have been sufficient to delink "Mario Costeja González" from google.es only. As I stressed above, delinking it from other European search engines seems ineffective but not harmful anyway because Spanish users are no more likely to use them than they are to use google.com. It would be a very different story if one were to argue that some legally published personal information online should be removed (i.e. the information itself, not just the link) altogether, or blocked at the source, e.g. by not allowing any search engine to index it in the first place. I am not against similar options, quite the opposite, but I suspect that, in order to consider them, we would have to have a serious debate about how really *harmful* the information in question needs to be to justify such a drastic solution. Yet this is something with which few

people in favour of the right to be forgotten seems to be willing to engage and prefer to swing back to "it's just a link" kind of position.

If we are serious about enabling individuals to have more and better control of their personal information online, we need to think in terms of designing the right constitutive rules of the logical space of online information. This is not easy, but more collaborative approaches may be envisioned, in some wiki-like way, through tagging or ranking. And policies about firewalls, robot.txt files, or IP addresses show that we have already begun the process. Geolocation could be used to remove links from google.com only for searches conducted within the EU, for example. This would mean that the same search about personal information on google.com would return all links if carried on in Washington but not when carried on in Brussels. Of course, there are always technical ways of by-passing such limitations. VPN providers, for example, let you change your IP location. The fact that Chinese authorities know this too well casts an odd light on the whole debate, one concerning power and control. This is why an international agreement on the right to be forgotten should be pursued. The question is who is or should be in charge of such deep transformations in our "onlife" experiences, the topic of the next section.

VI. Power

One of the crucial questions of our age is who may legitimately exercise what power over which kind of information. The question is complicated because one may interpret "who", "what power" and "which information" in many ways and end up talking at cross-purposes. In the previous sections, we saw what "what power" and "which information" may mean. Regarding the "who", in the context of the debate on the right to be forgotten there are *seven* entities that may be legitimately involved:

1. the person to whom the information refers, e.g., Mario Costeja González;
2. the publisher of the information, e.g., the Spanish newspaper *La Vanguardia*;
3. a search engine, e.g., Google Spain;
4. search engine users, that is, the public;
5. a national Data Protection Agency, e.g., the Agencia Española de Protección de Datos;
6. a national court of justice, e.g., the Audiencia Nacional; and finally,
7. the CJEU.

I argued at the end of the previous section that there should also be an eighth entity, which is currently missing: an international body that could ensure that a delinking legislation may be agreed upon by all liberal and democratic countries, as it happens in many other contexts.

In terms of power, # 5, # 6 and # 7 are legal entities that do not generate the personal information or its links – the two relevant types of information under discussion – but can determine how the other entities manage both.

The public (# 4) is the only entity in the list that has no direct power in either generating or controlling the personal information in question or the links provided by a search engine. However, the public exercises a decisive "passive power", insofar as the debate must refer to the "public interest", or rather a lack thereof, in order to determine whether the links to some personal information in question should remain available. In order words, the public is the reason why we are having a debate at all.

The mess concerns # 1, # 2, and # 3. Sometimes a person has both creative and controlling power with regard to the personal information that he or she wishes to see delinked. One may have generated the information and made it public. For example, so far the highest number of URLs removed by Google from its search results (8009) concerned Facebook.[15] In this case, it seems reasonable to expect the person to seek the removal of the personal information itself first, before asking to see it delinked.

The search engine has no creative power with respect to the personal information it indexes, but it has both creative and controlling power over its links and how they are presented. Such power could be exercised more imaginatively, but the ruling dictates that the links have to be removed, not, for example, lower-ranked, ranked historically, or annotated. This is a pity because much more could be done, in terms of sedimentation of information, if ISEOs were asked to find alternative solutions to the mere removal of links to legally available information, such as an old Anti-Social Behaviour Order (ASBO) now irrelevant.[16]

The publisher, such as *The Guardian* or the *BBC*, used to be the most powerful of all seven entities involved, because it has both creative and controlling power over the personal information in question. A publisher can regulate or block access to personal information quite easily. Its intervention makes any gerrymandered delinking no longer a problem. And it can operate discretely, without creating any Streisand effect. However, we saw that the ruling has had the self-defeating effect of forcing ISEO to act as gatekeepers, giving them even more power to decide what is or is not left on the map. At the moment, publishers are disempowered. They are merely notified about links removed, and even this is under discussion. They have no clear right of appeal, although they can informally negotiate a re-linking with a search engine, as it has happen with both *The Guardian* and the *BBC*. All this is unsatisfactory and in need of rectification. I remain convinced that, while we work on a better solution for the logical space of the Internet, a sensible procedure would be

---

[15] http://www.google.com/transparencyreport/removals/europeprivacy/?hl=en-US.

[16] In the United Kingdom, an ASBO is a civil order made against an individual proved to have engaged in anti-social behaviour, normally some minor incident that does not warrant criminal prosecution.

for an individual to request first the publisher to remove or block the *availability* of the personal information in question. Failing that, one could ask the search engine to delink it. If that does not work, an appeal to the national DPA or the relevant court in one's country would be the next step. And if that does not work, one could appeal to the CJEU. At each stage, the fact that the previous request was rejected should be documented and make a difference. In short, publishers should be the first to be consulted, not ISEOs, and their evaluation should matter.

VII. Limits

The ruling suggested some limits to which personal information may be delinked. Delinking is not admissible when the personal information, which would no longer be easily accessible, matters to the public good. This seems obvious at first sight, but a moment of critical reflection shows that there are plenty of borderline cases in which it is quite difficult to decide when information is of public interest. This is why every request needs to be carefully analysed individually, case by case, further empowering ISEOs. A disappointing solution is to reduce information that is of public interest to information that concerns public figures. This step merely doubles the difficulties. On the one had, it runs the risk of making some personal information – that about individuals that may be of public interest even if the individuals are no public figures – more easily delinkable because no longer subject to the exception. On the other hand, it shifts now the problem to defining when an individual is a public figure. In the information society everyone can be a public figure for 15 minutes, to paraphrase Andy Warhol, but precisely because an individual goes through many stages in life, it seems very difficult to identify the public nature and interest of some personal information without referring to the context within, and the purpose for which that information is being sought. Yet these two variables easily change through time and circumstances and cannot be easily forecast. Personal information that was of public interest may no longer be so for some time, and then turn out to be so again in the future. A coherent conclusion is that delinking should be equally dynamic, yet this, if feasible, would generate a rather chaotic situation.

The situation hardly improves when the debate refers to the relevance of the personal information in question. Recall that the ruling indicates that a delinking may be granted if it is "irrelevant or no longer relevant, or excessive in relation to those purposes and in the light of the time that has elapsed". As philosophers and logicians know too well, determining the relevance of some information is a very hard problem to solve. In this case too, it is impossible to fix the relevance once and for all, by devising a simple and universal rule, for relevance is dynamic and it is certainly a mistake to reduce it to a mere chronological matter – as the ruling seems to invite one to do – as if old information were irrelevant when compared to new one. Old news is not synonymous with irrelevant news. What someone had for breakfast this morning may be totally irrelevant compared to what that person did fifty years ago.

If we combine these difficulties – public interest in the personal information, public role of the referent of the personal information, and relevance of the personal information – paradoxically, one may argue that Mr Costeja has now become quite a public figure and that the personal information about his past is both relevant and of public interest and so should be re-linked.

VIII. Processing

The last problem I wish to analyse concerns whether an ISEO such as Google is correctly defined as "data controller". The definition of "data controller" in the EU Data Protection Directive (Directive 95/46/EC), is based on a definition of data processing (Article 2) that is so inclusive and generic that it cannot fail to support the CJEU's ruling:

> (b) 'processing of personal data' ('processing') shall mean any operation or set of operations which is performed upon personal data, whether or not by automatic means, such as collection, recording, organization, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, blocking, erasure or destruction;

In the UK, the Information Commissioner's Office had to issue a 20-page guidance in order to "explain the difference between a data controller and a data processor" (Information Commissioner's Officer 2014). The guidance rightly concluded that, according to the Directive, basically anyone doing anything with data is processing the data, and hence can qualify as a data controller. The Directive (adopted in 1995), which predates Google (founded in 1998) and the world of Social Media, does not appear to distinguish between *formatting* data, e.g. in terms of a barcode, *retrieving* data, e.g. by a scanner reading the barcode, *transmitting* data, e.g. by an electronic cashier sending the retrieved barcode data to a database, and *processing* data, e.g. by a computer that collects many retrieved barcode data and builds the profile of a customer. Data processing should be something much more specific that merely doing anything to or with data and/or information (another important distinction (Floridi 2010) underestimated by the Directive) but, at the moment, such specificity is lost in the Directive. The disappointment is that the CJEU could have followed the advice of the Advocate General, Niilo Jääskinen, who correctly recommended that "In his opinion, the internet search engine provider cannot in law or in fact fulfill the obligations of the controller provided in the Directive in relation to personal data on source web pages hosted on third party servers". (CJEU 25 June 2013) The CJEU could and should have interpreted the Directive much more stringently, concluding that a link to some legally available information does not process the information in question. It seems clear that the definitions of "data controller" and "data processor" themselves should be improved in the future.

IX. Conclusion

In this article I focused on some salient problems surrounding the debate on the right to be forgotten. They are not the only ones, but they have emerged as being among the most pressing since the issuing of the ruling. Whitehead famously remarked, in his *Process and Reality*, that "The safest general characterization of the European philosophical tradition is that it consists of a series of footnotes to Plato" (Whitehead 1929). In its turn, Plato's philosophy may be safely characterised as a philosophy of memory. From such a broad perspective it becomes obvious that the right to be forgotten is re-defining a very significant aspect of the European philosophical tradition. It is becoming a defining issue of our time because it is about how mature information societies will cope and deal with their memories.

Our culture is shifting its emphasis from the duty to remember to the right to be forgotten. I hope the shift may be evidence that some of the most tragic wounds in our European history are now healed. But I am concerned that it may be a sign that we are tired of confronting our past mistakes and responsibilities. Either way, I suspect that our technologies and their "undo" facilities might have quietly formatted our expectations about how much it is actually reversible in real life and should be rightfully archived.

Some sensitive and private information from the past needs to remain in the past, to ensure that memory does not undermine the future. This "remembering without recalling" is how I like to define *closure*. It is a difficult process that humanity has fine-tuned for a long time. It can already be found in the Bible, for example, where Isaiah (43:18) recommends: "Remember not the former things, nor consider the things of old". He is talking about the exodus, something that he certainly does not want to see forgotten, but that should not be constantly represented as the only perspective for Israel. Such a fine way of dealing with closure has become difficult on the web, a flatland lacking historical depth. Information is dynamically structured "on the fly" by our searches, and there are only endless first-pages to which we constantly add more information. We must ensure that the right kind of personal information may be remembered (no removal of past information) without being constantly recalled (no unnecessary resurfacing of past information). Getting this balance right is difficult but imperative, because we are working on one of the main roots of our culture.