



Online consent: how much do we need to know?

Bartłomiej Chomanski¹ · Lode Lauwaert²

Received: 13 July 2023 / Accepted: 25 September 2023
© The Author(s) 2023

Abstract

This paper argues, against the prevailing view, that consent to privacy policies that regular internet users usually give is largely unproblematic from the moral point of view. To substantiate this claim, we rely on the idea of the right not to know (RNTK), as developed by bioethicists. Defenders of the RNTK in bioethical literature on informed consent claim that patients generally have the right to refuse medically relevant information. In this article we extend the application of the RNTK to online privacy. We then argue that if internet users can be thought of as exercising their RNTK before consenting to privacy policies, their consent ought to be considered free of the standard charges leveled against it by critics.

Keywords Internet ethics · Digital privacy · Right not to know · Informed consent

1 Introduction¹

When you visit a website and click a button that says, ‘I agree to these terms’—do you really agree? Many scholars who consider this question (Solove 2013; Barocas & Nissenbaum 2014; Hull 2015; Pascalev 2017; Yeung 2017; Becker 2019; Zuboff 2019; Andreotta et al. 2022; Wolmarans and Vorhoeve 2022) would tend to answer ‘no’—or, at the very least, they would deem your agreement normatively deficient. The reasoning behind that conclusion is in large part driven by the claim that when most people click ‘I agree’ when visiting online websites and platforms, they do not really know what they are agreeing to. Their lack of knowledge about the privacy policy and other terms of the online agreements thus makes their consent problematic in morally salient ways.

We argue that this prevailing view is wrong. Uninformed consent to online terms and conditions (what we will call, for short, ‘online consent’) is less ethically problematic than many scholars suppose. Indeed, we argue that uninformed online consent preceded by the legitimate exercise of the

right not to know (RNTK, to be explained below) is *prima facie* valid and does not appear normatively deficient in other ways, despite being uninformed.

The paper proceeds as follows. In Sect. 2, we make more precise the concept of online consent and summarize the case against it, as presented in the literature. In Sect. 3, we explain the arguments for the RNTK in bioethics and show that analogous reasoning leads to endorsing the RNTK in online contexts. In Sect. 4, we demonstrate that the appeal to the RNTK helps defuse the critics’ arguments against online consent. Section 5 concludes: online consent is valid (with caveats, to be explored in what follows).²

2 The alleged problems with online consent

To a first approximation, a consumer of an online product or service gives online consent when they indicate, through some action, that they agree to the terms and conditions of the use of the product or service. Typically, such conditions include the provider’s ability to collect, track, and transfer the digital data generated by the user. The user expresses agreement by either clicking a specific button, labeled with

✉ Bartłomiej Chomanski
b.chomanski@gmail.com

Lode Lauwaert
lode.lauwaert@kuleuven.be

¹ Department of Philosophy, Adam Mickiewicz University, Poznan, Poland

² Institute of Philosophy, KU Leuven, Leuven, Belgium

¹ Portions of this paper update and expand on the material previously covered (in Dutch) in *Online consent en het recht op onwetendheid*. Tijdschrift over Cultuur & Criminaliteit, (13)2, forthcoming. Permission to reuse it is gratefully acknowledged.

² Nothing in the paper is intended to suggest that we think major internet companies behave unobjectionably by offering extremely long and difficult to parse privacy agreements. This is a separate question from that of the validity of user consent.

'I agree to these terms' or some similar phrase, or simply by continuing to use the website.³

Consent, including online consent, is morally (and legally, but here our focus is on its moral dimensions) significant. Canonically, when A consents to B Φ -ing, A makes it morally permissible for B to Φ , where Φ -ing without consent would have been morally impermissible.

A's consent is thought to be valid when a number of conditions are met. Specifically, A must be competent to give consent; A's consent must be freely given; and, crucially for our purposes, A must possess an adequate level of knowledge about what they are agreeing to.

Online consent can be criticized for failing to meet any of these three conditions; scholars argue that typical internet users are not free enough to give it (because they lack meaningful alternatives or because the use of the online services is addictive), that they suffer from severe decision-making biases (undermining their competence), or that they do not know enough about the terms and conditions. Consent in such circumstances is alleged to be normatively deficient to a significant degree, so that its power to transform moral reality is vastly diminished. In this paper we focus on the latter criticism. Thus, even if we are successful, we will not have vindicated online consent by replying only to the informed-ness.

The *locus classicus* of arguments against online consent is Daniel Solove's (2013).⁴ Solove writes that "consent is not meaningful in many contexts involving privacy" because "(1) people do not read privacy policies; (2) if people read them, they do not understand them; (3) if people read and understand them, they often lack enough background knowledge to make an informed choice; and (4) if people read them, understand them, and can make an informed choice, their choice might be skewed by various decision making difficulties" (2013: 1888).

It is not entirely clear what Solove means by consent not being meaningful. A literal interpretation of the phrase would suggest that such consent is, simply, invalid, and entirely lacks its morally transformative force. Akin to a meaningless gesture, it signifies nothing.

A weaker interpretation of the term is also available. The idea could be that online consent does not achieve one of the core functions that consent is intended to achieve. In the case

³ In law, expressions of consent can also take varied forms. E.g., the GDPR defines consent as follows: "'consent' of the data subject means any freely given, specific, informed and unambiguous indication of the data subject's wishes by which he or she, by a statement or by a clear affirmative action, signifies agreement to the processing of personal data relating to him or her" (Article 4, Sect. 11).

⁴ While Solove is mostly focused on the legal dimensions of online consent we interpret his arguments as bearing directly on the question of whether such ought morally to be considered valid.

under consideration, that would be protecting the user's sovereignty over their digital information, or, in Solove's words: "[providing] people with control over their personal data, and ... [the ability to] decide for themselves how to weigh the costs and benefits of the collection, use, or disclosure of their information" (2013: 1880). In the absence of adequate knowledge of the contents of privacy policies, such control is illusory. "Not meaningful" consent in this case is simply the consent that does not fulfil its role. This notion can be further filled out in a number of ways.

One such way is to maintain that uninformed consent makes reasonable decisions about whether to share digital data impossible. Andreotta et al. (2022) and Wolmarans and Vorhoeve (2022) argue, for instance, that online consent, given its shortcomings, may fail to secure the users' ability to perform reasonable cost-benefit estimations about the trade-offs involving their digital data. As the former put it:

One of the features of AI and Big data, after all, is that surprising or unexpected information, or correlations, can sometimes be revealed from existing data sets... This makes it hard for data subjects to *assess the risk* of consenting to sharing certain data because it is hard to predict how their data could be used in the future (2022: 1720; references omitted, emphasis added).

Wolmarans and Vorhoeve express the same point thus: "It is ... near impossible for a user to arrive at a reasoned assessment of the likelihood of possible implications of data processing when they consent to it" (2022: 97, emphasis added).⁵ The thought seems to be that, in the absence of good evidence about the probabilities of different outcomes of the decision to consent to data sharing, giving one's consent is unreasonable.

Marcel Becker makes a related point when he writes that "Individuals' control and knowledge about the flow of information are lost. As we all are keenly aware, requiring people to agree with terms and conditions does nothing to solve the problem" (2015: 310, emphasis added). Online consent, on this view, fails to solve "the problem" of the lack of knowledge of, and control over, one's own information.

Gordon Hull declares that online consent "fails to protect privacy" partly because "users do not and cannot know what they are consenting to" (2015: 91), while at the same time sharing the concern that such ignorance does not allow users to make reasonable cost-benefit estimates. Hull questions whether digital consent can even be rational.

Shoshanna Zuboff's (2019) arguments against online consent closely follow earlier work carried out by other legal theorists, Radin (2013) being a particular inspiration. Radin

⁵ We take it that the authors use "likelihood" in the colloquial sense, roughly synonymous with probability.

sees the crucial role of informed consent in protecting user autonomy: “the notion of informed consent [in medical ethics, which Radin assumes is analogous to ethics of online consent] ... relies squarely on autonomy theory, which is why the information required about what is happening to the patient must be detailed and understood by the patient before consent will be deemed to exist” (2013: 89). Consequently, relying on uninformed consent fails to respect the user’s autonomy.

At any rate, we can distinguish at least three distinct manners in which the lack of information undermines online consent and renders it normatively deficient. The first is quite direct and radical: if online consent is uninformed, then it is invalid (this may be attributed to Solove, as we saw above; to Radin, who calls such digital contracts a degradation and perversion of contract law; and to Hull, who denies that such consent is rational).

The second interpretation of the case against online consent attributes the following, weaker, claim to the critics: if online consent is uninformed, then it does not allow users to have control over their digital data; they are thus deprived of their autonomy and sovereignty over their information. This leaves open the question of whether such consent ought to count as valid—but it is unquestionably ‘problematic’ and in need of reform.

A third complaint against online consent identified by the authors cited above can be formulated as follows: if online consent is uninformed, then the users are unable to make reasonable decisions about trading off privacy for access to digital services (where ‘reasonable’ means something like ‘based on good evidence and sound decision-making strategies’).

Regardless of which interpretation one chooses, we think the arguments against online consent fail to establish its normative deficiency in any of the three senses. We argue for this claim by deploying the concept of the RNTK. First, we demonstrate how the RNTK can be justified in online contexts; second, we show how the application of the RNTK disarms the arguments against online consent we have just summarized.

3 The right not to know

The conception of the RNTK first arose out of discussions of informed consent in bioethics. In essence, the RNTK is supposed to be the patients’ right not to receive medically relevant information, if they decide against it. There are two dimensions of this right (for more on rights in general, see Wenar 2021): on the one hand, we might think of the RNTK as a liberty right—as the mere absence of the patient’s duty to know, generating no duties of any kind in others; on the other hand, we might think of it as a claim right—the right

held against others that they do not provide the patient with information against the patient’s will, i.e., the right against significant nonconsensual interference with the patient’s ignorance (or perhaps, with the contents of the patient’s mind).

We assume that these rights, whatever their form, are best thought of as *pro tanto* rights—rights that may be overridden in exceptional circumstances, provided there are excellent reasons to do so (one such reason could be that the information is needed to protect others from serious harm). The description of the exact circumstances when such rights can be overridden is beyond the scope of this paper. We merely signal here that it is implausible to think of them as absolute.

We are not alone, of course, in making the assumption that the RNTK is best thought of as a *pro tanto* right, and it seems thoroughly reasonable to us. However, in an insightful piece, John-Stewart Gordon (2017) raises a problem with thinking of the right not to know as a *pro tanto* right. Gordon argues that “[p]roponents of the right not to know ... cannot claim that the right not to know is a [*pro tanto*, rather than absolute] right, because then they would allow for the possibility that they might be unable to always justify their own decision in cases of conflict between principles (such as between autonomy and non-maleficence)” (2017: 51). The idea, as we understand it, is that there are circumstances when an agent’s ignorance (e.g., when prospective parents refuse to undergo tests for a genetic condition they may pass to their offspring) can harm an innocent third party (the prospective parents’ prospective offspring—assume for the sake of argument that being born with the disease would harm the child). *Pro tanto* rights can be restricted to avoid (some cases of) harm to innocent third parties. But when the parents exercise their right not to know, they cannot claim that this decision will not harm anyone, precisely because they do not know whether it will. Thus, they will be unable to justify their exercise of the right, since they cannot claim that the non-maleficence exception does not arise.

While we recognize the strength of this argument, we would make two points in response: first, as it will emerge in further discussion, it does not seem to us that decisions whether to agree to privacy policies would in general carry risks of substantial harm to innocent third parties. Consequently, at least when narrowed down to our concerns—RNTK in online contexts—Gordon’s argument need not apply.

In making our second reply, we will help ourselves to Gordon’s thought. In a closing passage of the article we have mentioned, Gordon says:

I strongly believe that one should respect a person’s wish not to be informed if he or she indicates this preference. At the same time, I also strongly believe that in many practical cases the so-called right not to

know is overridden by other considerations, such as the principle of non-maleficence. The latter principle does not necessarily always rule out the supposed right not to know; ethical judgements certainly depend on the particular case, and cannot be assumed beforehand (2017: 52).

By and large, we find ourselves in agreement with Gordon here (perhaps we would prefer to substitute “some” for “many” in the second cited sentence). What we say below strikes us as basically an elaboration of this line of thought in the specific context of online consent: people’s wishes⁶ not to be informed ought to be *prima facie* respected, though there are cases—that cannot be decided in advance—in which other considerations outweigh the strength of the wish. But since there is no logical conflict between Gordon’s argument about the right not to know we cited earlier, and Gordon’s concluding remarks we cited just now, we also think there is no logical conflict between Gordon’s argument about the right not to know and our subsequent discussion.

3.1 Consistency

Bioethicists over the years have offered a number of justifications for the RNTK. In what follows, we will briefly summarize some of these arguments.

One argument for the RNTK rests on what one might call “justificatory consistency.”⁷ The idea is that since we are justified in recognizing the patients’ right not to consent to a beneficial surgery, the same considerations should, *a fortiori*, also justify their right not to receive potentially beneficial relevant information. After all, refusing surgery in most cases will be more detrimental to the patient’s health than refusing health-related information. Graeme Laurie (2014) makes this point as follows:

If there are good reasons to respect persons through recognition of the right to self-determination as reflected in the absolute right to refuse medical treatment as exists in many countries—irrespective of the consequences—then it follows *a fortiori* that other refusals with less potentially drastic outcomes—such

as refusing to receive personal information—ought similarly to be respected with no questions asked. (54)

While we would not put it as strongly as Laurie (remember, we think of the RNTK as a *pro tanto*, not an absolute, right), we find the reasoning compelling. Moreover, Laurie’s point is applicable not just to the liberty-RNTK but also to the claim-RNTK—it is not just that we have no duty to undergo surgery; the right to refuse it is a right we hold against others’ interference with our bodies. The RNTK, so conceived, would then be, to borrow a phrase from Douglas and Forsberg (2021) again, a right we hold against others’ “significant nonconsensual interference” with (the contents of) our minds. This reasoning applies straightforwardly to online contexts as well. Refusing to read the terms and conditions of use of a digital service surely belongs among the “other refusals,” less impactful than refusing surgery, that Laurie speaks about.

3.2 RNTK, legitimate interests and digital consent

Other bioethicists justify the RNTK on the grounds that it protects our important interests. Exercising the RNTK specifically protects us from suffering various kinds of psychological harm and anguish. As Roberto Andorno (2004) points out, “[f]or many people, the discovery that [for example] they have a genetic condition that places them at a high risk of suffering certain untreatable diseases could so depress them that the quality, joy, and purpose of their lives would literally evaporate” (435). Exercising the RNTK protects people from having to suffer such harms. Ignorance is, at least sometimes, bliss.

Consider, in addition, the story of a “Dr. Fearful,” as told by Jonathan Herring and Charles Foster (2012, p. 22):

Dr. Fearful is a philosopher of the highest caliber. He is a happy, optimistic person. His doctor has done routine blood tests and comments that these will indicate whether or not he has a predisposition to dementia. Dr. Fearful pleads with his doctor:

‘Don’t tell me the results. I know what I am like. I would rather live my life in ignorance of the future than run the risk that I will have to live with knowledge that I may develop a disease. If the tests prove a predisposition I will sink into a malaise. I know also that once other people know I have this disease it will affect how they treat me. They will not involve me in future plans and will look for early signs of the disease. That would be intolerable. This is an important life choice. Let me choose ignorance.’

Dr. Fearful’s request for ignorance seems to us to be a perfectly legitimate life choice, as deserving of respect as any others.

⁶ Contrary to Gordon, we believe that going against such a wish often counts as a rights violation, including in mundane cases and concerning information that is not especially important (like, as in Gordon’s example, knowing the score of a soccer game). However, just because one can come up with examples of insignificant transgressions, this does not show that, when the transgression occurs, rights are not violated (consider, for illustration, the following case: I ask you, my roommate, not to move the pen I legitimately own and that I placed on the table; nevertheless, you do; this seems to us to be both insignificant, and also a property-rights violation).

⁷ The phrase is borrowed from the work of Tom Douglas and Lisa Forsberg (2021).

We agree. The life choice to avoid mental anguish associated with a burden of such knowledge should be protected—because avoiding mental anguish is a legitimate, important life interest we have. In the absence of compelling reasons to override it (e.g., preventing significant harm to third parties), such protection is *pro tanto* justified. The protection comes in the form of a claim right against interference with the contents of our minds.

We hold that there likewise could be legitimate interests protected by the RNTK when it comes to online privacy. For instance, although this is probably only true for a small minority of internet users, one can imagine cases where knowing how one's information may be used by Big Data-reliant companies could lead to adverse psychological effects.

Consider the case of Sophie, who is aware that it is possible to download all the data that Google has about her. She heard from friends who had done it that Google has several times more data, presenting a far more accurate picture of their online presence, than they expected it to have. Sophie knows that finding out exactly what Google knows about her will cause her extreme anxiety. With few realistic avenues for altering or removing this information, Sophie chooses not to access Google's data about her, rather than going through the stressful experience of finding everything out.

This is a legitimate choice. Sophie's inner peace is protected by her exercise of the RNTK the information that Google collected on her. In this case too, the RNTK is a right against others' interference with our mental contents (it sure seems like a rights-violation to inform Sophie, against her will, about what she does not want to know).

Sophie's case might strike some as contrived. Surely, protecting against psychological harms of knowing too much is applicable only to a minority of real internet users. However, avoiding psychological harm is not the only interest that the RNTK could be grounded in protecting in online contexts. To see this, it is useful to keep in mind, as critics of online consent like to point out (Solove 2013; Hull 2015; Zuboff 2019), that we would need to devote an enormous amount of time to read in their entirety the privacy policies and other terms and conditions we have agreed to. The cost of doing that was estimated (already over a decade ago) to come to more than \$780 billion in potentially lost productive capacity (McDonald and Cranor 2008). In other words, to read all the information about how our data are used, we would collectively be giving up hundreds of billions of dollars.

However, thanks to the exercise of the RNTK this information, we can spend our time on other, more productive activities instead. The RNTK protects us against having to forgo all this economic value. Given that it would impede our economic interests (to the tune of nearly \$1tn), it would be a rights-violation to somehow force us to spend our time familiarizing ourselves with the contents of these policies.

So, it is not just that we have no duty to read privacy policies we agree to. The protection-of-interests rationale also implies other people's duty not to inform us of their contents against our (expressed or signaled) will.

At this point, one could critique our reasoning as follows: Suppose that after you click 'I agree' without reading, the website presents you with a small pop-up box containing the summary of their terms and conditions, explaining in a few sentences the gist of what the provider may do with your data. It would seem very odd to maintain that displaying such a box is a rights-violation. And yet, on the account we are pursuing here, it appears we are committed to making this counterintuitive claim.

Of course, the objection only applies to our defense of the claim-RNTK. The liberty-RNTK imposes no duties on others. So, the website displaying the summary of privacy policies even after the 'I agree' button was pressed would not violate the liberty-RNTK. But we need not concede ground on the claim-RNTK either. The claim-RNTK, recall, is a right that others do not significantly interfere with our minds. Plausibly, a small, easy to ignore pop-up box on the bottom of the screen containing some information we may wish to avoid does not rise to the level of significant interference.⁸

3.3 RNTK, autonomy and digital consent

Rights may also be justified by appealing to the idea that they are required to respect the capacity for purposive rational action (Gewirth 1978). In bioethics, such defenses of the RNTK are generally couched in terms of autonomy. On the face of it, this defense is relatively straightforward. Consider for example another of Andorno's (2004) arguments for the RNTK. On his view, the right "is precisely based on the idea that people should be free to make their own choices with respect to information. If we understand autonomy in this wider sense [i.e., as self-determination], then the decision not to know should be, at least in principle, as fully respected as the decision to know" (436).

Understanding the RNTK as rooted in autonomy (and, especially, as having the very same basis as the right to know) also underlies its recognition in international documents and guidelines, such as the UNESCO Universal Declaration on the Human Genome and Human Rights, and Council of Europe Oviedo Convention on Human Rights and Biomedicine. On this conception of the RNTK, if autonomy grounds the right to receive information, then, by the same token, it also grounds the right to refuse it.

⁸ What *would* count as such a violation? Perhaps forcing the customer to correctly answer some reading-comprehension questions about the policy before allowing her to express agreement.

Whether as patients or as users of internet services, we are autonomous beings, “free to make [our] own choices with respect to information”.⁹ Consequently, our RNTK privacy policies and other terms and conditions should be as respected, on autonomy grounds, as our RNTK personal medical information. It is, once again, clear to us that the RNTK is not just a liberty right—it is not just that we have no duty to read these legal documents; it is also a right against others’ interference with our mental contents. Forcibly telling us information we do not want to hear interferes with how we want to determine our mental lives to go. Absent excellent reasons, we should not be told what we do not want to know.

To be sure, some worries remain. One can legitimately wonder, for instance, whether we may really base our defense of the RNTK on autonomy if obtaining certain types of information is a precondition of autonomous choice. As John Harris and Keywood (2001) argue:

Absence of crucial information is inimical ... to the ability to control one’s own destiny, and hence inimical to autonomy... in a way that other autonomy-limiting choices are not. For where the individual is ignorant of information that bears upon rational life choices she is not in a position to be self-governing. ... [O]f course it is not necessarily irrational not to want to know one’s probable life expectancy and many would be prepared to forego autonomy rather than face the knowledge of a looming premature death. However they cannot defend the wish to remain ignorant of a fact like that in the name of autonomy (421).

Harris and Keywood’s basic idea is that without “crucial information” relevant to our life choices, we cannot govern ourselves autonomously. Since refusal to hear “crucial information” makes the exercise of autonomy impossible, it cannot be defended on autonomy grounds any more than selling oneself into slavery can (the comparison is Harris and Keywood’s).

While there is room for disagreement with the normative underpinnings of Harris & Keywood’s points (see, e.g., Bortolotti 2013), we need not engage with those. Instead, there is empirical evidence that typical internet users possess a substantial amount of information regarding how their data are used. First, as the PEW Research Center (2019) reports, “72% of Americans report feeling that all, almost all or most of what they do online or while using their cell-phone is being tracked by advertisers, technology firms or other companies. Another 19% think some of what they do is being tracked” (2019: np). Thus, 91% of Americans have at least some awareness of being tracked. Moreover,

as Caleb Fuller (2019) has found after administering a more complex series of surveys (focused specifically on Google’s data practices),

Many consumers indeed are relatively well-informed. When queried about their knowledge of Google’s information-collection model ..., respondents overwhelmingly are aware that the company gathers personal information about them as they use Google. At least regarding the existence of the practice, the extent of information asymmetry is low, with 89% ... of respondents indicating awareness of Google’s collection of personal data. ...

The data suggest that most respondents possess a relatively high degree of awareness [about more specific collection practices] ... 87% ... know that Google keeps a record of searches and 80% know that Google registers a browser’s physical location. ...

Lastly... survey responses suggest that consumers are significantly less well-informed about [what Google does with their data], but not completely uninformed about them. ... While 81% ... correctly identify that Google collects information ‘to target ads based on your search history and location’, many of them consistently overestimate the number of uses to which Google puts their data (361-2).

The findings belie the claim that (at least Google) users are unaware of crucial information about their data (see also Mills 2022 for additional, somewhat more impressionistic, arguments along those lines—still, some debate remains to be had as to what information about data practices should count as crucial). So, on the face of it, Harris & Keywood’s concerns do not apply here. Exercising the RNTK privacy policies is consistent with possessing crucial information about the most important aspects of data collection and processing that online platforms do. Indeed, per Fuller, users appear to think that more of their data are collected than actually is. This type of ignorance would presumably lead to less extensive data sharing—the opposite of what is traditionally thought to be the users’ most concerning behavior (we assume a priori that the findings could be extended to platforms other than Google).

Lastly, users may be aware that, regardless of particular contents of privacy policies, there exist regulatory frameworks intended to prevent the most egregious of privacy abuses; they may also know that other people’s use of online products and services has not typically resulted in a concerning number of dangerous privacy violations. Such knowledge would not necessarily be captured in various surveys of what customers know about the contents of privacy policies, but it may well play the role of the crucial information on which to base the decision to consent to them.

⁹ Notice that Andorno does not specify that it is *medical* information.

4 What does the RNTK do?

As we saw, the claim made by scholars quoted in the beginning is that online consent in the absence of (a certain amount of) knowledge of the privacy policies is normatively deficient in at least three ways: uninformed consent is invalid, implies a loss of control over data, and leads to unreasonable behavior. In response, we advance the following three claims: the less-than-fully informed consent preceded by an exercise of the RNTK (1) should be considered *pro tanto* valid, (2) need not be detrimental to autonomy, and (3) need not be unreasonable.

4.1 Valid consent and the RNTK

In this subsection we argue against the following principle: if online consent is uninformed, then it is invalid. Specifically, we argue that consent preceded by a legitimate exercise of the RNTK ought to be at least presumed valid.

That consent preceded by the exercise of the RNTK is valid seems to be an intuitive implication of recognizing the RNTK. Consider the quick example used by Beauchamp and Childress (1979):

If a deeply committed Jehovah's Witness were to inform a doctor that he wishes to have everything possible done for him, but does not want to know if transfusions or similar procedures would be employed, it is hard to imagine a moral argument to the conclusion that he must be told (79) [quoted in Ost 1984: 301-2].

Intuitively, the patient would not have a legitimate ground for claiming that his doctors' subsequent treatment constituted assault in virtue of not being consented to. Rather, the physicians would seem to be acting permissibly, and—crucially—without violating their patient's bodily rights when proceeding with treatment while respecting his pleas for ignorance. This suggests that the patient's uninformed consent preceded by the exercise of the RNTK should count as valid and worthy of respect—i.e., as successful in transforming the moral reality. To see this more clearly, imagine a case in which the patient does not know that transfusions are likely to be employed during surgery and does not mention not wanting to know this—while making it clear that being a Jehovah's Witness is very important to him. In such a case, his consent to the surgery, while uninformed, would be normatively suspect. But the only difference between this case and the previous one is that here the patient's consent was not preceded by exercising the RNTK.

More theoretically: one important aspect of rights is that they give their possessor the power to change others'

rights and entitlements. By consenting to surgery, the patient gives the doctor the right to interfere with his body. Similarly, with the RNTK. By exercising it, we endow others with certain liberties they did not previously have. For instance, by exercising the RNTK the privacy policies, we give the internet companies the (liberty) right to collect our data without ensuring that we are informed about the specifics. This changes what they may permissibly do.

Now consider the following case:

DIGITAL CONSENT Upon landing on the European Parliament's website (<https://www.europarl.europa.eu/portal/en>), Ignoramus encounters the following pop-up privacy notice: “We use analytics cookies to offer you a better browsing experience. You have the choice to refuse or accept them. [underneath this text, there are two buttons, one saying ‘I refuse analytics cookies’ the other saying ‘I accept analytics cookies’. Then, underneath those, the website further says:] For any information on the other cookies and server logs we use, we invite you to read our data protection policy, our cookies policy and our cookies inventory” and provides links to the policies. Without reading the cookies policy, Ignoramus clicks the ‘I accept’ button.

Two things seem to be happening here. First, Ignoramus decides not to read the policies. Second, Ignoramus signifies consent to cookies being used. We admit: this is not informed consent. But it is preceded by a legitimate exercise of the RNTK. Ignoramus clearly has an unimpeded ability to access the policies—they are available at a click of the mouse. His choice to click ‘I agree’ without reading signals the exercise of the RNTK. This is analogous to the patient in Beauchamp & Childress' case, except here the exercise of the right is signaled through a non-verbal action, whereas in their case, it is verbalized. This is not a relevant difference (if, instead of talking things through with the patient, the doctors in Beauchamp & Childress' case left a closed envelope with the description of the procedures and what they involve, and the patient, aware of what's inside, decided not to open it—our intuitions remain). Consequently, we hold that Ignoramus's consent ought to be respected as much as the patient's in the original scenario.

Typical internet users act in ways that are plausibly like Ignoramus's ways in DIGITAL CONSENT. They exercise their RNTK, and thus their subsequent uninformed consent is, at least presumptively, valid.

4.2 Autonomy and the RNTK

In this subsection we argue against the following principle: if digital consent is uninformed, then it is autonomy-undermining. As before, we look at this question through the lens of the RNTK.

As we argued in Sect. 3, the RNTK itself can be defended on the basis of autonomy. To exercise the RNTK is to make a sovereign decision over what information we want to populate our heads with, and what information we want to exclude.

In relevant ways, this is analogous to our sovereignty over physical property like, say, our car: we get to decide who/what is, and who/what is not allowed inside. Furthermore, suppose that, by some coincidence, allowing you to get in our car increases our autonomy (perhaps you know the city better than we do, so having you in the car increases the number of diverse options we will have while driving). This makes our decision not to let you in our car no less autonomous. Similarly, when people opt not to read the terms and conditions, in spite of thus foreclosing some options, that decision is—at least presumptively¹⁰—autonomous as well.

Critics may object: do we not—by making the autonomous choice not to access information—thereby deprive ourselves of control over our data, thus undermining our autonomy on balance? Perhaps we do. But we typically do not criticize choices as autonomy-limiting merely in virtue of their resulting in diminished future control over some aspect of our lives. A typical employment contract or tenancy agreement involve, respectively, the employee's loss of control over much of their time, and the landlord's loss of much control over their property. On the face of it, all of these seem perfectly fine, despite involving autonomous choices to lose or limit one's control over an important resource one owns. It is not obvious why decisions to lose or limit one's control over the digital data one generates when, say, visiting The Washington Post's website, should be treated differently.

4.3 Reasonable estimates and the RNTK

In this subsection, we argue against the principle that if online consent is uninformed, then it prevents reasonable decisions.

On the face of it, the criticism appears unimpeachable. When you do not know what you are agreeing to, how can you be reasonable about it? However, we still think that the strength of the critique is overstated.

The problem raised by the critics relying on the italicized principle above is that uninformed consent makes it difficult to reasonably estimate the cost–benefit tradeoffs when it comes to decisions about digital privacy. Yet, there are many areas of life where we decide to forgo accessing new

information in exchange for doing something else with our scarce mental resources, so that our subsequent cost–benefit estimates are made more difficult. Nonetheless, this need not preclude our decision process from being reasonable when considered in its totality.

Let us explain what we mean by that phrase. Suppose we model agents with incomplete knowledge, who nevertheless have to make a decision, as follows¹¹:

The agent faces a two-stage decision problem: (i) What to pay attention to: The agent selects an information [search] strategy to refine her belief about the state. ... (ii) What action ... to take: This is a standard choice under uncertainty with the beliefs generated in the first stage via Bayesian updating. The objective is to maximize the expectation of [utility] less the cost of information ..., which is a function of the information strategy (Mackowiak et al. 2021: 8).

In other words, when we speak of our decision processes considered in their totality, we mean the two-stage model just quoted: first, information gathering; second, taking the action on the basis of the information gathered. We assume that the reasonability of the whole process depends on both stages.

Crucially, information search is not free. There are costs associated with acquiring new knowledge. Sometimes, these costs may be onerous (perhaps very time-consuming or requiring specialist skills). Put another way: ignorance has costs, but so does ignorance-reduction.

A marginal gain in information is a marginal reduction of ignorance. The choice to stop searching for information is simply the choice to retain the present level of ignorance. Sometimes, searching for more information prior to making a choice is simply not worth it, while making choices at some level of ignorance is. Consider these examples: as you run out of your house to drive to work, you neglect to check a detailed, hour-by-hour, weather forecast; as you join an amateur rugby team, you do not peruse detailed statistics on possible injuries; you buy a book at an airport bookstore without poring over the reviews to see if you will like it; you take vitamins without reading the leaflet on side effects too closely. To the extent that these are autonomous choices, not dictated by irresistible outside forces, they do not seem apt

¹⁰ Why “presumptively”? We grant that there could be cases where the subsequent autonomy limitation is so severe that it could invalidate consent. It would take a further argument, however, to show that these cases are relevantly similar to typical instances of online consent.

¹¹ Here we follow the model developed by economists working on rational inattention, characterized thus: “In a rational inattention model, an agent can choose in a flexible way what kind and how much information to absorb. The agent then acts based on the chosen information. This is a model most readily applicable to situations in which a lot of information is available, the key constraint is an agent's limited ability to process information, and the agent has had time to think or experiment to determine an optimal information acquisition strategy” (Mackowiak et al. 2021: 2).

for criticism as unreasonable. Rather, they show that, in each case, the cost of obtaining additional bits of information (the marginal cost of reducing ignorance) was judged (not necessarily unreasonably) to be higher than the expected marginal benefit the new information would provide.¹² (NB: we are not yet claiming that these choices are analogous in some special way to choices concerning online consent; the examples are intended to show the plausibility of assessing decisions in their totality.)

Note, however, that it would be impossible to evaluate these choices as reasonable or unreasonable without also considering the costs of information gathering (the costs of ignorance). Trivially, if the agent's "objective is to maximize the expectation of [utility] less the cost of information" (ibid.) we cannot tell whether the objective is met without knowing the cost side of the equation. It follows, then, that decisions can count as reasonable even if the agent does not have a lot of information when taking them.

Now, it is perfectly acceptable to treat online consent as the decision problem with the two-stage structure described above, where the first stage consists of gathering information about the terms and conditions of the online agreement, and the second stage is the action (agree or not agree) taken on the basis of this information. Moreover, we may also think of subsequent decisions within the online ecosystem (e.g., about sharing some particular bit of data, say liking a picture or retweeting a pasta recipe), made after the agreement is consented to, as consisting of the two stages mentioned above: gathering information about what sharing the data might do for us (what might happen to that data, how that might affect us etc.), and then making the decision whether to share.

As we pointed out, it is impossible to judge whether the decision is reasonable only on the basis of considering one of the stages mentioned above. This applies to the critique of online consent we discuss in this subsection: one cannot simply declare the agent's choice unreasonable when all one focuses on is the second stage of the decision problem (i.e., how much information the agent has when clicking the 'I agree' button). But this strikes us as precisely what the critics of online consent do. They seem to focus only on what the agents do not know.

A lingering doubt remains—surely, it would be better for people to have more decision-relevant information than less, regardless of whether they typically decide to forgo

this information. On the one hand, we see no issue with this way of thinking about online consent. Perhaps it would be, in some sense, better if people possessed more information that is, in some sense, relevant to their decisions. However, to the extent that this is what the criticism of online consent amounts to, it is not especially effective. Such criticism, after all, is applicable to pretty much every choice people make. We never have full information and there is always room to learn more.

Indeed, appealing to the RNTK can help us see this issue in a different, much more benign light. In essence, exercising the RNTK is about having the right to decide whether obtaining new potentially relevant information is worth the effort—by our own lights. Appealing to the RNTK is merely a way of recognizing that information search is costly, and stopping short of full information may in many cases be a rational, reasonable decision, depending on individual circumstances—and that it should be respected as such.¹³

Some people in some contexts will decide they need lots of information, and so will be prepared to bear substantial costs to obtain it; others in the same context, or the same people in different contexts, will make different determinations. The RNTK protects our ability to make this determination for ourselves. The interest-protection rationale (we have an interest in avoiding excessive costs), the autonomy rationale (we are sovereign over our mental contents), and the consistency rationale (if we are fine with more weighty choices being done in relative ignorance, we should be fine with the choice to retain the present level of ignorance in less impactful contexts) all justify this point.

Suppose the critics insist that there is some (objective) level of information concerning the costs and benefits of a decision that a person needs to possess no matter what, for her choice to count as reasonable. Let us concede that such levels exist. It seems to us, however, that the onus is on the critic to supply us with an argument why, in general, the kinds of ignorance people typically exhibit when agreeing to privacy policies fall below that objective threshold—especially in light of the evidence we get from Fuller, the PEW surveys, and what ordinary people can be expected to know from observing others around them who already use the online services.

In general, our response to such arguments is this. First, the level of knowledge exhibited by participants in Fuller's experiments is sufficient for reasonability, on any plausibly relevant conception thereof. Second, given the intersubjective variety in the magnitude of costs of reducing ignorance,

¹² A choice to stop searching may even be reasonable in cases where information has high marginal utility, simply because the choice is urgent. There is a fire in my department library, and I don't know whether the priceless manuscripts we house have been taken to safety. I have to decide whether to try to save them or run for my life, but I have to do it now. It would be great to give the chair a call and ask about the manuscripts. But there's no time.

¹³ Inevitably, people will sometimes make mistakes in assessing the costs of information search. But this pedestrian fact should not detract us from the broader point we're defending: the RNTK protects our ability to select from a number of search strategies ourselves, rather than rely on others to circumscribe it for us.

a very broad range of choices regarding how much information one acquires is consistent with being a reasonable decision-maker. The RNTK protects our ability to move within this range, even if it may also allow us sometimes to stray outside it.

Perhaps the critics' idea instead amounts to this: the decision-maker cannot make reasonable decisions because she is unable to assign probabilities to different outcomes of sharing her data ('what are the chances of my data being sold to unscrupulous data brokers thereby destroying my life? Who knows?'). Hence the impossibility of reasonable decisions. This objection, as we saw, is pressed by Andreotta et al. (2022) and by Wolmarans and Vorhoeve (2022). But it doesn't follow, from the fact that the likelihoods of different outcomes cannot be reasonably estimated before a decision is made, that such decisions are unreasonable. After all, there exist well-researched strategies for decision-making without the knowledge of the relevant probabilities (maximin, maximax, the Hurwicz rule, to name a few—see, e.g., Weber (1987); Buchak (2022)). Are we to count all decisions based on these strategies as unreasonable? (This may surprise some Rawlsians). Or are we to think that people do not abide by them in real-world decision-making? We find no arguments for either conclusion in the critiques we cited.¹⁴

Moreover, there are reasons to question the assumption that having less information inevitably leads to lower decision accuracy. As Gerd Gigerenzer et al. (2011), decisions with less information available can be equally or more accurate than decisions when the agent integrates all information. While it is beyond the scope of this paper to argue that this is what happens when we make decisions about digital privacy, the result shows that we cannot simply infer that less informed people will inevitably make worse decisions.

This establishes the presumption in favor of choices preceded by the RNTK being not just legitimate, but also reasonable. This, in turn, establishes the presumption that a critique of such choices as being too uninformed is at best incomplete and at worst paternalistic, as it assumes the critic's superior knowledge of the decision-maker's own circumstances and welfare, which is unlikely. In making choices prior to which they do not acquire a lot of information, decision-makers should be presumed reasonable unless there are good reasons to the contrary. Lack of information alone does not suffice for unreasonableness.

¹⁴ Alternatively, if consenting in situations where probabilities are unknown is normatively suspect, then a range of extremely important decisions—"whether to marry, who to marry, whether to have children, what career path to follow, how much time to devote to friends and family" (Roberts 2022: 1), and many business decisions—appear to inherit the same suspicion. Yet, we don't ordinarily think that our consent in such matters is problematic. Mills (2022) makes a related point using the analogy with marriage.

5 Conclusion

Overall, the appeal to the RNTK disarms arguments against uninformed online consent being problematic. These arguments fail to establish that online consent is invalid, autonomy-undermining, or unreasonable. A choice preceded by the exercise of the RNTK need not be any of those things.

It thus seems to us that uninformed online consent does not deserve its bad reputation among scholars. It is likely that typical internet users exercise their RNTK prior to agreeing to the terms and conditions of various websites (they certainly act as if they do). Their consent, then, ought to be presumed meaningful in all of the senses under consideration.

Acknowledgements The authors wish to thank Andrew Rebera and Kristien Hens for their detailed comments on a written draft of this paper. Portions of this paper were presented in seminars at Western University, London, ON and Adam Mickiewicz University, Poznan. The authors thank the audiences for their feedback. We also thank the reviewer for this journal for useful suggestions.

Author contributions BC: 70%; LL: 30%.

Funding No funding was received for this paper.

Data availability Not applicable.

Declarations

Conflict of interest The authors declare no competing interests.

Ethical approval and consent to participate Not applicable.

Consent to publish Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Andorno R (2004) The right not to know: an autonomy based approach. *J Med Ethics* 30(5):435–439
- Andreotta AJ, Kirkham N, Rizzi M (2022) AI, big data, and the future of consent. *AI Soc* 37(4):1715–1728
- Barocas S, Nissenbaum H (2014) Big data's end run around procedural privacy protections. *Commun ACM* 57(11):31–33
- Beauchamp T, Childress J (1979) *Principles of biomedical ethics*. Oxford University Press, New York

- Becker M (2019) Privacy in the digital age: comparing and contrasting individual versus social approaches towards privacy. *Ethics Inf Technol* 21(4):307–317
- Bortolotti L (2013) The relative importance of undesirable truths. *Medicine Health Care Philos* 16(4):683–690. <https://doi.org/10.1007/s11019-012-9449-x>
- Buchak L (2022) Normative theories of rational choice: rivals to expected utility, the stanford encyclopedia of philosophy (Summer 2022 Edition). E. Zalta (Ed.). Retrieved from <https://plato.stanford.edu/archives/sum2022/entries/rationality-normative-nonutility/>. Accessed 2 May 2023
- Douglas T, Forsberg L (2021) Three rationales for a legal right to mental integrity. In: Lighthart S, van Toor D, Kooijmans T, Douglas T, Meynen G (eds) *NeuroLaw: advances in neuroscience, justice & security*. Springer, Cham, pp 179–201
- Fuller CS (2019) Is the market for digital privacy a failure? *Public Choice* 180(3):353–381
- Gewirth A (1978) *Reason and morality*. Chicago University Press, Chicago
- Gigerenzer GE, Hertwig RE, Pachur TE (2011) *Heuristics: the foundations of adaptive behavior*. Oxford University Press, Oxford
- Gordon J-S (2017) On the nature of the right not to know. In: GadebuschBondio M, Spöring F, Gordon J-S (eds) *Medical ethics prediction and prognosis interdisciplinary perspectives*. Routledge, London, New York, pp 42–54
- Harris J, Keywood K (2001) Ignorance, information and autonomy. *Theor Med Bioeth* 22(5):415–436
- Herring J, Foster C (2012) Please don't tell me': the right not to know. *Camb Q Healthc Ethics* 21(1):20–29
- Hull G (2015) Successful failure: what Foucault can teach us about privacy self-management in a world of Facebook and big data. *Ethics Inf Technol* 17(2):89–101. <https://doi.org/10.1007/s10676-015-9363-z>
- Laurie G (2014) Recognizing the right not to know: conceptual, professional, and legal implications. *J Law Med Ethics* 42(1):53–63
- Mackowiak B, Matejka F, Wiederholt M (2021) Rational inattention: a review. ECB Working Paper, No. 2570. Frankfurt a. M.: European Central Bank (ECB) <https://doi.org/10.2866/417246>
- McDonald AM, Cranor LF (2008) The cost of reading privacy policies. *I/S J Law Policy Inf Soc* 4:543–565
- Mills K (2022) Consent and the right to privacy. *J Appl Philos* 39(4):721–735
- Ost D (1984) The 'Right' not to know. *J Med Philos: Forum Bioeth Philos Med* 9(3):301–312
- Pascalev M (2017) Privacy exchanges: restoring consent in privacy self-management. *Ethics Inf Technol* 19(1):39–48. <https://doi.org/10.1007/s10676-016-9410-4>
- PEW Research Center (2019) *Americans and privacy: concerned, confused and feeling lack of control over their personal information*
- Radin MJ (2013) *Boilerplate: the fine print, vanishing rights, and the rule of law*. Princeton University Press, Princeton, NJ
- Roberts R (2022) *Wild Problems*. Portfolio/Penguin, New York
- Solove DJ (2013) Privacy self-management and the consent paradox. *Harv Law Rev* 126(7):1880–1903
- Weber M (1987) Decision making with incomplete information. *Eur J Oper Res* 28(1):44–57
- Wenar L (2021) Rights. In: Zalta EN (ed) *The stanford encyclopedia of philosophy*. Spring, Cham
- Wolmarans L, Voorhoeve A (2022) What makes personal data processing by social networking services permissible? *Can J Philos* 52(1):93–108
- Yeung K (2017) 'Hypernudge': big data as a mode of regulation by design. *Inf Commun Soc* 20(1):118–136
- Zuboff S (2019) *The age of surveillance capitalism: the fight for a human future at the new frontier of power*, 1st edn. PublicAffairs, New York

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.