



Contextual Integrity as a General Conceptual Tool for Evaluating Technological Change

Elizabeth O'Neill¹

Received: 8 March 2022 / Accepted: 1 August 2022 / Published online: 12 August 2022
© The Author(s) 2022

Abstract

The fast pace of technological change necessitates new evaluative and deliberative tools. This article develops a general, functional approach to evaluating technological change, inspired by Nissenbaum's theory of contextual integrity. Nissenbaum (2009) introduced the concept of contextual integrity to help analyze how technological changes can produce privacy problems. Reinterpreted, the concept of contextual integrity can aid our thinking about how technological changes affect the full range of human concerns and values—not only privacy. I propose a generalized concept of contextual integrity that is applicable to a broader variety of circumstances, and I outline a new, general procedure for technological evaluation. Among the attractive features of the proposed approach to evaluating technological change are its context-sensitivity, adaptability, and principled presumptive conservatism, enabled by the mechanism the approach supplies for reevaluating existing practices, norms, and values.

Keywords Contextual integrity · Privacy · Values · Norms · Ethics · Disruptive technologies

1 Introduction

The fast pace of technological change necessitates new evaluative and deliberative tools.¹ Among other things, the effects of new technologies can be difficult to predict and may vary by context, many decisions about technologies involve stakeholders with partially conflicting values, and technological change has the potential to cause social and moral disruption and to threaten the things people

¹ This is not to say that the fast pace of technological change is inevitable: the maintenance of such a fast pace depends on human decisions.

✉ Elizabeth O'Neill
e.r.h.oneill@tue.nl

¹ Philosophy & Ethics, Eindhoven University of Technology, Eindhoven, The Netherlands

value most.² Given the complexity of human ends, and the complexity of our world, neither kneejerk conservatism nor kneejerk pursuit of novelty is sufficient for guiding our decisions about specific technological changes. Furthermore, there is a substantial gap between the guidance supplied by traditional, high-level moral theories such as deontology, consequentialism, or virtue ethics and the information one needs to evaluate and deliberate on issues relating to technology (see, e.g., Jacobs & Hultgren, 2021). Even in domains with apparent convergence on abstract values or principles, decision-making (at an individual or collective level) about whether to support or oppose a technological change often still involves substantial empirical, evaluative, and deliberative work.³

This article outlines a step-by-step procedure for assessing whether a technological change is likely to facilitate social or moral disruption and, ultimately, whether the technological change is likely to advance or threaten our most important ends.⁴ When more fully elaborated, this tool is intended for use by a range of stakeholders, including designers creating technologies, individuals or groups deciding whether to use or adapt a technology, and governments deciding whether to permit or support the development of a technology. The procedure builds on Nissenbaum's (2004, 2009) concept of contextual integrity (CI). Nissenbaum (2004, 2009) introduced the concept to help us understand how technological changes can produce privacy problems. Reinterpreted and broadened, the concept can aid our thinking about how technological changes affect the full range of human values and concerns—not only privacy.

Nissenbaum (2004, 2009) developed her theory of CI as an account of when we can anticipate that people will view something as a privacy violation and when we should judge that something is a privacy problem.⁵ In many established social contexts, there are entrenched norms that govern the flow of personal information. On Nissenbaum's view, when those norms are contravened, there is a *prima facie contextual integrity* violation, and thus a *prima facie* reason to think a privacy violation has occurred. Descriptively, the account is intended to predict and explain people's judgment that something is a privacy violation—we can expect that people in the relevant social contexts will view violations of entrenched informational norms as privacy violations. However, the theory also has a normative component. There is

² On the difficulties associated with making predictions about technologies, see, e.g., Collingridge (1980), Hansson (2011), and Brey (2012); on varying ethical problems across contexts, see, e.g., Nissenbaum (2009) and Forge (2010); on socially disruptive technologies, see Carlsen et al. (2010), Brey et al. (2019), Hopster (2021a; 2021b), and Nickel et al. (2021); on morally disruptive technologies, see Baker (2013; 2019) and Nickel (2020); on the relationship between technological change and moral change, see Swierstra et al. (2009), Verbeek (2011), Pols (2013), Swierstra (2013), Kudina & Verbeek (2019), van de Poel (2021), Hopster et al. (2022), and van de Poel & Kudina (2022).

³ See, e.g., Morley et al. (2021), Whittlestone et al. (2019), and Mittelstadt (2019) for comments on the gap between practice and high-level, abstract principles in the context of artificial intelligence research.

⁴ By "important" ends, I mean ends that the agent possessing the ends views as important.

⁵ Nissenbaum is specifically concerned with "privacy as it applies to information about people" (individual, identifiable persons) (2004, p. 106), so what she offers is an account of information privacy, rather than a full account of privacy.

always a question about whether entrenched norms are legitimate. It could be that people view some norm violation as a privacy violation and hence as bad, yet the entrenched norm is itself problematic, with the result that the flow of information that violates the norm is not bad after all. On Nissenbaum's view, privacy is defined as the *appropriate flow of personal information*. To evaluate for ourselves whether a violation of informational norms constitutes an inappropriate flow of personal information, we must evaluate the norm itself. We can do this by assessing whether entrenched informational norms advance contextual and more general shared ends, and by evaluating such ends by appealing to more fundamental ethical principles and other considerations. If we conclude the norms are legitimate, we will conclude that a flow of information that violates those norms constitutes a (problematic) CI violation: the norm-violating flow of personal information is *inappropriate* and a privacy violation.

I propose to broaden Nissenbaum's concept of CI so that it applies not solely to informational norms and privacy but to norms and values generally.⁶ Nissenbaum's point about the importance of entrenched norms can be broadened in the following way. Human groups have built up many norms over time through a complicated set of processes, with the result that some of our entrenched norms exist because they advance shared ends. Of course, other norms originated or are maintained because they advance the ends of powerful sub-groups; in addition, some norms do not advance the ends of anyone but are instead simply "along for the ride," maintained by general-purpose mechanisms and never subjected to selection or other forces strong enough to eliminate them (cf. Boyd & Richerson, 1992). These phenomena are crucial to consider in one's evaluation of technological change, but I set them aside for the moment. Within a social system, entrenched norms that advance shared ends serve a *function*—specifically, they serve a function for those people whose ends they advance. This is the broad point. How does the broad point apply to privacy norms? Since many important human practices rely on the transmission of information about people, we have developed a special body of norms that regulate such information interactions. We have come to characterize many of these norms (in English) as having to do with "privacy." When people violate these norms, for instance, because of technological or other changes that facilitate such violations, a special kind of ethical problem—namely, privacy problems—may result. If we consider normativity more broadly, we can see that there are also many other classes of norms and values that human cultures and sub-cultures have developed, and these norms and values can be similarly disrupted by technologies. These include norms and values relating to loyalty, fairness, justice, liberty, property, beneficence, etc. (Haidt & Joseph, 2008; Curry et al., 2019; for a survey of kinds of norms, see O'Neill, 2017). What I propose is that if we broaden our understanding of the CI

⁶ I am not the first to suggest something along these lines—at least some others have previously proposed applying Nissenbaum's CI idea beyond privacy: Kim and Werbach (2016) suggest in the conclusion to their article on gamification in the workplace that CI may be useful for analyzing the ethics of gamification.

concept, we could interpret technological disruption to any type of entrenched norm within an established social context as a potential violation or reduction of CI.

Why would this be useful? Looking for threats to CI, broadly construed, can point us to types of technological change that people in the context are likely to protest, whether for reasons relating to privacy, fairness, beneficence, or any other entrenched values, and can help us identify technologies that are likely to be socially or morally disruptive. Thinking in terms of CI can help us evaluate a technological change in a systematic and step-by-step way, inform our decision about whether to embrace or oppose a technological change, and help us determine where compensatory actions may be warranted if a technological change does occur.

The importance of considering CI in one's evaluation of technological change derives from several characteristic human features. First, humans pursue many shared ends, and pursuit of many individual and shared ends necessitates cooperation. Second, our continued existence and pursuit of ends are heavily reliant on cumulative culture (cf. Henrich, 2016) and the numerous aspects of ourselves and our environments that are products of past human actions. Third, the nature of our ends is such that the question of how to achieve them is often not answerable in advance by the lone theorizing individual. In sum, CI is valuable because of the type of beings that humans are: advancement of individual and shared ends relies in many ways on entrenched, human-influenced aspects of our world, and we do not have a very good understanding of all of the ways in which advancement of our ends rests on those entrenched aspects of the world. It is for this reason that we cannot afford to be heedless innovators. At the same time, it is imperative to recognize the existence of the numerous entrenched elements of human life that *hamper* the advancement of many important individual and shared ends, including those aspects of the world that have been intentionally implemented to exploit people, to preserve unequal distributions of power, and so on. When evaluating potential technological change, then, we must also avoid kneejerk conservatism, with no principled mechanisms for making changes to established traditions.

In the next section, I elaborate on Nissenbaum's notion of *context*, the idea of shared ends, and the idea that contexts can have integrity. In Section 3, I introduce a procedure for evaluating technological change, which draws on a generalized concept of CI. In Section 4, I supply a historical case to illustrate how CI broadly construed can help us analyze a situation in which members of a community considered switching from the use of one technology to another—namely, a case from seventeenth- and eighteenth-century France involving the use of sickles versus scythes for harvesting grain. In Section 5, I highlight some of the advantages of a functional, CI-based approach to the evaluation of technologies. Section 6 concludes.

2 Contexts, Shared Ends, and Integrity

Nissenbaum relies on a particular notion of context in her account of contextual integrity.⁷ She is concerned with “structured social settings with characteristics that have evolved over time (sometimes long periods of time) and are subject to a host of

⁷ The concept is closely related to other ideas from sociology and social theory, such as social spheres, fields, domains, and institutions (Nissenbaum 2009, p. 130–131).

causes and contingencies of purpose, place, culture, historical accident, and more” (2009, p. 130). Such social settings are “characterized by canonical activities, roles, relationships, power structures, norms (or rules), and internal values (goals, ends, purposes)” (Nissenbaum, 2009, p. 132). Contexts in this sense are something more than concrete situations; they are “evolved abstract spheres of activity” (Benthall et al., 2017, p. 19). As examples of contexts, Nissenbaum lists “health care, education, employment, religion, family, and the commercial marketplace” (2009, p. 130), noting that contexts will often be nested and that the same sorts of contexts will vary by culture and setting. Contexts may be more or less formalized and institutionalized (compare courtrooms and open-air markets [Nissenbaum, 2009, p. 135]), and the norms involved “may be explicitly expressed in rules or laws or implicitly embodied in convention, practice, or merely conceptions of ‘normal’ behavior” (Nissenbaum, 2018, p. 838). In any case, a key feature is that there are “objectives around which a context is oriented” (Nissenbaum, 2009, p. 134).

The idea of contextual objectives can be difficult to pin down. We can identify examples: Nissenbaum writes that “the purpose of educational contexts include transmitting knowledge, knowhow, and arguably, social values to a society’s young...” (2009, p. 134). But a context over time is composed of a shifting set of people with many different aims; not all of the individuals playing a role in the context endorse the objectives of the context (e.g., a child can be coerced to attend school; an employee may work for a paycheck without supporting the mission of the corporation that employs them). The causes of the context’s existence and persistence may have little to do with the values and concerns of the people in it; furthermore, the values and concerns of the people in the context may differ substantially from each other’s. How, then, should we think about the relationship between contextual objectives and the evaluative attitudes of those who participate in or simply encounter a context?

I will use “ends” as a term of art to refer to an agent’s full set of evaluative attitudes.⁸ I suppose that an agent’s full set of ends often is highly complex, supervenes on multiple psychological systems, and is not something that the agent would be able to fully articulate. It is worth emphasizing that there are many routes by which people can share ends. A parent may have the end of their child’s ends being advanced.⁹ Two people who have different, non-conflicting desires may discover that both their desires will be advanced by the same means, such that in those means, they have a shared end. Simply in virtue of being so inter-reliant, humans have a great many shared ends. For instance, group members that rely on the group for advancement of their own individual ends may derivatively value the survival of the group and thereby share an end. In a range of conditions, humans seem to have

⁸ As Street (2009) puts it, an agent’s evaluative attitudes include “desires, attitudes of approval or disapproval, unreflective evaluative tendencies such as the tendency to experience X as counting in favor of or demanding Y, and consciously or unconsciously held evaluative judgments, such as judgments about what is a reason for what, about what one should do or ought to do, about what is good, valuable, or worthwhile, about what is morally right or wrong, and so on” (p. 110).

⁹ If the parent’s ends are described with a utility function, one can think of the utility function as recursive, so that incorporated within the parent’s utility function is a term representing the child’s utility function. See Kleiman-Weiner et al. (2017) for a proposal along these lines.

a general motivation to advance the ends of others when doing so comes at little cost to their own ends.¹⁰ Furthermore, via an assortment of developmental mechanisms (instruction, imitation, conformity, reasoning, etc.), humans characteristically come to share many moral and cultural values with others in their social groups, and across cultures, there are many shared values and norms (which sometimes—though not always¹¹—can result in shared ends). Even strangers waiting for a bus usually have overlapping ends—e.g., they may desire that ordinary social norms be observed, that the situation does not erupt in conflict, or that no one in the situation is harmed, and they may be motivated to take actions to advance those objectives under the right conditions (e.g., giving a dirty look when someone cuts in line, or intervening in a case of apparent harassment).¹²

Of course, it is often also the case that some human ends conflict. One common circumstance can be characterized in terms of conflictual coordination games, such as the so-called battle of the sexes game (Hankins & Vanderschraaf, 2021). In this scenario, two people prefer to cooperate but face multiple possible ways in which they might do so, and each person prefers a different cooperative arrangement. For instance, two traders might each strongly prefer to make a deal rather than not, but each would most prefer the deal that is more financially favorable to them. They thus have a combination of shared and conflicting ends. Likewise, at the population level, some of the ends of a sub-group may conflict with the ends of other sub-groups.

For the individual, contextual objectives may be among the ends she possesses. If she (for example) approves of, assents to, or believes in the value of the contextual objectives, then those contextual objectives are among her ends. For someone who believes in the value of education, for example, the objectives of the education context may be among her ends. However, as noted earlier, a person may also participate in a context to which she is indifferent or that she abhors. An enslaved person forced into the context of slavery does not have the objectives of that context among their ends. Furthermore, there are some contexts (and associated contextual objectives) that advance no one's ends, and some contexts that, on balance, do more to undermine than to promote shared ends. Such contexts may nonetheless be maintained and their objectives promoted—for example, as a result of a default conservatism and failure to reevaluate contexts in changing circumstances, or as a result of one of the many hurdles to bringing about social change, such as ignorance of others' beliefs and preferences. For instance, Sunstein (2019) discusses research from

¹⁰ See, e.g., Tomasello's discussion of helping behavior in children from the age of 14 months (2019, Ch. 8) and Hare & Woods (2020) on intragroup strangers.

¹¹ For instance, two groups may each value care for in-group members, yet this common value is precisely what leads the two groups to possess conflicting ends in a case of disputed resources, where each group believes the resource will allow them to better care for their in-group members.

¹² Catarina Dutilh Novaes and Hein Duijf offer a very simple preliminary proposal for how one might model degrees of adversariality or conflict of interests between two parties, based on whether each agent involved wants some state of affairs to obtain, does not want that state of the world to obtain, or is neutral. Dutilh Novaes (2020) is concerned with measuring the degree to which parties' interests overlap (interest alignment) for the purpose of assessing whether two people engaged in argumentation have largely shared goals; in such a case, they can engage in cooperative argumentation. A model like this could also be useful for thinking about the shared goals required for CI.

Bursztyn et al. (2018) that suggests that many men in Saudi Arabia are privately in favor of women's participation in the workforce, but that they believe most other men to be opposed to women working. As a result, despite privately being in favor, they voice opposition to women working.¹³ One can imagine something similar with an ornate coming of age context and set of traditions, some of which are physically harmful to the individuals involved, where there was formerly some function served by the context (it advanced some ends of specific individuals, sub-groups, or perhaps the group as a whole) but where changing circumstances mean that the ends of no one in the group are advanced by the tradition anymore. We can still talk about the objectives of these contexts, but since they are no longer endorsed by anyone in the group, no one in the group will take them to have normative force. If the group performs a CI analysis, they are likely to conclude that the context does not advance their ends (but instead threatens some individuals' ends), and so they lack reason to retain it.

Some of the main objections to Nissenbaum's CI theory emphasize the role that factions and powerful sub-groups play in creating and maintaining contexts, norms, and other elements of normative life; some argue in particular against the idea that contexts often involve settled objectives that reflect shared goals (see, e.g., Rule, 2019 and Benthall & Haynes, 2019). Nonetheless, there are some contexts with relatively clear objectives to which participants do explicitly assent, and there are also contexts that advance shared ends, despite participants not explicitly assenting to the internal ends of that context (e.g., because they have not thought about whether the context advances their ends). In such contexts, at least, we can think in terms of CI and we will be able to apply the evaluative approach that I propose. As I have noted, there can also be value in applying a CI analysis even in cases where it turns out the context advances no one's ends. The trickiest cases presumably are those that involve few shared ends, and where the context advances some people's ends and hinders others'. The CI approach I am proposing will certainly not lead to group agreement in every case, and I do not expect it to eliminate the need for political struggle. However, even in challenging cases with few shared ends, and even if a group does not accept the conclusion of an evaluator's CI analysis, performing the analysis can have value—it can, for example, help group members to uncover shared ends of which they were unaware, lead group members to question aspects of their normative life that they previously accepted without thought, and help locate the source of disagreement between factions (e.g., disagreement on empirical questions or interpretations of values).

So much for contexts and shared ends. What about the *integrity* of contexts? Nissenbaum does not explicitly supply a definition of integrity in an abstract sense that is separable from privacy. For the purposes of applying the CI concept beyond privacy, I propose to define contextual integrity in the following way: a context has integrity to the extent that the shared ends of the individuals participating in the context are advanced through the pattern of practices, norms, and other normative elements that are characteristic of the context.

¹³ This example has to do most directly with a practice, but it has indirect implications regarding how people conceive of the contexts of the home and work.

3 A General, Functional Procedure for Technological Evaluation

To investigate whether a technological change, such as the introduction of a new technology, the modification of a technology, or a new application, reduces CI, and to inform a decision about what to do in response, we can use the procedure described in Table 1. Here, the procedure is presented in broad strokes. The idea is that the procedure will be eventually fleshed out with more specifics as we build up a better picture of which types of technologies tend to affect which contexts, and which values and other normative elements tend to appear in which contexts. Once elaborated, the procedure is intended for use by a range of actors, including groups engaged in collective deliberation, group members determining whether to oppose or support a technological change, and individuals who are observing a group (of which they are not a member) and aiming to evaluate a technological change on the basis of a rich understanding of the socio-technical situation of that group. I will refer to the person implementing the procedure as “the evaluator,” but it is important to note that an evaluator may sometimes be a decision-maker, and sometimes an observer.

For convenience of presentation, the procedure is primarily formulated as something to be performed prior to the technological change; however, some version of the procedure can be performed before, during, and after a technological change occurs, and it will frequently be advisable to perform the procedure iteratively, including as one monitors the effects of a technological change that has already been implemented. Additional comments on each step follow after Table 1.

3.1 Comments on the Steps in the Procedure

3.1.1 Step 1

With regard to Step 1 and the question of what constitutes a technological change, we are interested in the creation of new technologies, as well as in changes to the production, maintenance, disposal, use, or application of existing technologies. Any of these changes can be morally or socially significant (see, e.g., Palm & Hansson (2006) on assessment over a technology’s life cycle).

3.1.2 Step 2

In Step 2, one identifies relevant contexts and entrenched normative elements that are likely to be affected by the technological change. This can be a substantial undertaking, particularly when one is dealing with a new or emerging technology. It is of course difficult to anticipate the possible consequences of a technological change, whether to the normative elements in human life or to other aspects of the world (Hansson, 2011; Palm & Hansson, 2006). Within this step, as well as Steps 3 and 4, given that they all involve anticipation of consequences, one may beneficially

incorporate numerous existing tools, such as technomoral scenario-building (Swierstra et al., 2009), risk analysis, technology assessment, scenario planning (Hansson, 2011), forecasting, futures studies (Brey, 2012, 2017), and generation and consideration of diverse narratives (see the efforts of Dihal et al. (2021) or the Future of Life Institute's "worldbuilding" competition [2022]). One may also want to invest in careful investigation of which values are "embedded" in a given technology, or, which outcomes a technological design will tend to promote or will make more likely within particular contexts (Klenk, 2021). For this, one may wish to draw on "disclosive ethics" (Brey, 2000). Even when a change has already occurred, it is not simple to identify the full range of relevant consequences it has had so far. For the purpose of identifying contexts to consider and for the purpose of predicting and recognizing how elements of normative life are affected by the technological change, an analysis will be more complete if informed by a diverse group, including those affected in different ways, experts and laypeople, the vulnerable, and the historically marginalized.¹⁴ This is especially true in cases in which a technological change is likely to affect many contexts. Furthermore, obtaining an adequate picture of the likely and actual effects of some technologies may require substantial laboratory research, involving both experimentation and simulation, as well as (incremental) experimentation within society (van de Poel, 2011, 2016). A crucial part of this process of anticipation and monitoring involves attempting to identify (possible, likely, or actual) unintended effects and other secondary effects of the technological change (van Eijndhoven, 1997).

What does it mean to say that elements of normative life—e.g., practices, roles, norms, concepts, values, and value-laden artifacts—will be affected by the technological change? This could mean that the technological change itself will have *direct* effects on normative elements—for instance, a new dam makes a traditional, sacred practice of fishing downstream impossible, or medical implants violate the sanctity of the human body—or that the technological change may facilitate rather *indirect* effects on normative elements: it might facilitate threats to practices, norms, or things people value by making alternative practices more attractive, incentivizing norm violations, making a norm less salient, introducing a new value that is in tension with an old value, introducing a new context that lacks norms and thereby endangers shared values, etc.

The phenomena mentioned thus far are all attenuating influences on elements of normative life, but in other cases technological change may reinforce normative elements, for example, by facilitating compliance with norms, supplying new ways to execute an old practice, or making it easier to promote something one values. In such a case, the technological change may increase CI—but one cannot be sure until one has checked to see whether the technological change threatens other normative elements as well and (in Step 4) checked whether the entrenched normative elements that are enhanced actually promote shared ends.

¹⁴ See the discussion of standpoint theory in Grasswick (2018) and see, e.g., Anderson (2015) on how social power can produce moral epistemic disadvantages. See also the multitude of work on the topic of how to incorporate stakeholders into technological design and policymaking.

Table 1 A general, functional procedure for evaluating technological change, using CI, broadly construed

Steps	CI implications
1. Identify a (possible) technological change	
2. Compare (predictions about) the new, changed situation with the old situation. (a) Identify social contexts that are likely to be affected by the change. (b) Within contexts, examine how entrenched normative elements—practices, roles, norms, concepts, values, value-laden artifacts, etc.—are likely to be affected. (c) More generally, examine how cross-contextual normative elements are likely to be affected	Identification of <i>prima facie</i> CI reduction—when entrenched elements of normative life are threatened
3. Examine how the <i>ends of individuals</i> are likely to be affected (on the basis of [2] and other information, including how background factors that are conducive to individuals' ends are likely to be affected)	
4. Examine how <i>shared ends</i> (both contextual and general) are likely to be affected (on the basis of [2–3] and other information, including an examination of how background factors that are conducive to shared ends are likely to be affected)	Identification of <i>genuine</i> CI reduction Identification of other threats to shared ends
5. Having identified a genuine CI reduction, if the evaluator assents to the shared ends at stake, then she will conclude that the <i>CI reduction constitutes a problem</i> . If she objects to the shared ends, then she will conclude that although there is a CI reduction, the CI reduction is not actually problematic because the group's shared ends are misguided or wrong	Evaluation of CI reduction, or other threats to shared ends, via evaluation of shared ends: conclusions about whether the effects of the technological change are <i>problematic</i>
At this stage, the evaluator may also invoke ends or premises unrelated to CI to reach an evaluative conclusion	

One of the ways in which this step requires substantial work is that we may need to draw on moral psychology, sociology, and anthropology and perform considerable theorizing, in order to develop an adequate articulation of the pertinent norms that may be affected by a given technological change. Here again, Nissenbaum's account suggests a valuable strategy. An important part of Nissenbaum's contribution on the question of privacy comes in her identification of *parameters* that can be used to systematically articulate privacy norms. The key parameters that Nissenbaum identifies are actors (information subject, sender, and receiver), information type, and transmission principles. The information type parameter has to do with what the information in question is about. In Nissenbaum's theory, this parameter highlights the point that different norms regulate the flow of different types of information. For instance, in the medical context, different norms regulate information about patient medical conditions, patient attire, insurance provider, and account balance (Nissenbaum, 2009, p. 143). Transmission principles are the conditions under which information may or must be transmitted. Some examples: one may be allowed to transmit information only via sharing, not selling; a transmission may occur only if the recipient pledges to keep the information confidential; or one's medical information must be transferred to the

patient if requested. A more fully elaborated account of the generalized CI procedure that I am proposing will develop sets of parameters for use with different domains of norms. For instance, we can develop sets of parameters for articulating norms in the following categories: fairness norms and the distribution of costs/benefits; norms governing punishment; care norms; property norms; honor norms; norms of respect; and norms of liberty, among many others.

On the topic of entrenched elements of normative life, I want to emphasize that entrenchment is a matter of degree and that some things that may seem highly entrenched at first glance may in fact not be. We may take something for granted because we inherited it (e.g., the consent model of privacy, or norms that permit disposal of yard waste in landfills), yet it is relatively new in human history and we still have not realized all of the ways in which it affects our full range of ends, or some people may have become habituated to something (e.g., group instant messaging, or posting consumer reviews on the internet) that is still far from well integrated into our larger normative system. It will often be worth questioning how deeply entrenched the normative elements are that are threatened by the technological change. When the entrenched element is relatively newer, as a rule of thumb we may expect that disruptions to it will be less socially or morally disruptive (cf. Hopster, 2021b), and we should pay extra attention in Steps 3 and 4 to whether the element in fact advances our ends—because it has not gone through the longer historical process of adjustment and screening that older elements have gone through.

3.1.3 Step 3

In Step 3, one examines how the *ends of individuals* are likely to be affected. This step draws on information obtained in Step 2 as well as other information, including information about how background factors that are conducive to individuals' ends are likely to be affected. Why are the ends of individuals worth considering? They will matter for the CI evaluation, and for collective deliberation, inasmuch as they matter for *shared* ends—e.g., if our society aims to protect autonomy and well-being, then it is relevant to consider the ends of the individuals whose autonomy and well-being we wish to protect. It is worth devoting a distinct step to the investigation of individual ends to help us disentangle the various reasons that group members have to care about the ends of individuals, and to aid our subsequent investigation of which ends are shared with others.

A consideration of *background factors* is an important component of Steps 3 and 4. To understand how a technological change affects individual and shared ends, we must look not only at how the change affects the entrenched elements of normative life that have been put in place over time for the purpose of advancing those ends, but also at background factors which no human may have ever put in place purposefully, but which nonetheless play a role in facilitating achievement of our ends. For instance, suppose that someone ordinarily uses a stovetop kettle to prepare her tea each morning. The kettle takes 5 min to boil water, and during that time, she has got into the habit of watering her plants. To be more environmentally friendly, she decides to replace her old kettle with an electric kettle. Now her water

boils in 1 min. For the purpose of advancing her various ends (in this case, keeping her plants alive), it will be advantageous if she notices straightaway that the technological change she made to advance one end has removed a background feature of her environment (a 5-min forced pause in activities, compelled by the constraints imposed by the physics of her kettle) that was instrumental for bringing about one of her other ends. She will probably want to make an alternative plan—for instance, cultivating a new habit triggered by some other prompt during some other time interval—which ensures that she waters her plants each day. In both Steps 3 and 4, it is important to investigate the effects of a technological change on background factors that matter for ends—anticipating or predicting such effects, monitoring and/or experimenting as the change is implemented, and taking compensatory measures in cases where one's ends are threatened due to changes in background factors.

3.1.4 Step 4

In Step 4, one ascertains whether a *prima facie* CI reduction is a *genuine CI reduction*, and one investigates whether the technological change also threatens to undermine shared ends in other ways. In a case where the change threatens entrenched normative elements, one must now check whether those entrenched normative elements in fact advance shared ends—if yes, one concludes that the *prima facie* CI reduction is a *genuine CI reduction*.

During this step, one looks at both shared contextual and shared general ends—there is practical value in examining standard sets of values associated with particular contexts and activities, as well as considering general shared ends such as virtue or welfare. For instance, in the CI approach to privacy, Nissenbaum lists some of the values that have been identified as worth considering when we are looking at systems featuring flows of information: the fairness of power shifts, democracy, unfair discrimination, informational harm, equal treatment, civil liberties, and individual autonomy (2018, p. 842). For the purpose of applying the CI concept beyond privacy, we can draw on existing ethical and social science research to compile comparable lists of shared values to consider when analyzing systems featuring resource distribution, punishment, relationships, care settings, etc. In addition, separating the examination of whether contextual norms and practices advance shared contextual ends and whether contextual ends advance shared general ends allows one to draw conclusions such as the following: there is CI within some subsystem, but the system as a whole lacks CI.

One may conclude that a *prima facie* CI reduction is *not* a *genuine CI reduction*, if one finds that entrenched normative elements are disrupted, but that those elements in fact do not advance shared ends. In this case, there was little CI to begin with. This type of outcome reflects the fact that as circumstances change and as people's conceptions of their ends change, entrenched norms that may have advanced the shared ends of the group in the past may no longer serve that function. If traditional norms are no longer (or never were) conducive to shared contextual or general ends, we may conclude that we have reason to adjust those norms. As discussed in Section 2, sometimes entire contexts, even, are not conducive to general ends; we may conclude that we have reason to revise or eliminate those contexts as well.

In this step, one can also pick out cases where the technological changes will not disrupt normative elements themselves, but will change background factors such that entrenched normative elements will *no longer* suffice to advance shared contextual ends—this is another form of CI reduction. Suppose a synthetic chemical industry has emerged for the first time, and someone proposes to introduce a new chemical that poses no threat to a society's entrenched norms or social contexts. This new chemical, however, does threaten the survival of a remote animal species, and the society in question values species diversity. Thus far, the protection of this species has not required any action from that society—no norms or practices have been required. But now that someone is proposing to introduce the new chemical, a background factor has changed in a way that threatens that society's shared end of species diversity. To protect this shared end, the group may conclude that it needs to implement new norms, such as a norm that the introduction of new chemicals cannot proceed without substantial testing.

3.1.5 Step 5

In Step 5, the evaluator takes a position on the shared ends in question and reaches an evaluative conclusion on the possible technological change. An evaluator who is a member of the group will reflect on the shared ends of her group, and she will reach a conclusion about whether she still assents to those ends. The evaluator who is merely an observer has thus far been engaged primarily in investigating people's normative lives and ends descriptively, but now will need to evaluate them. The evaluator will assent or object to shared ends on the basis of the usual mechanisms of moral cognition and established practices of evaluation: consideration of the information obtained thus far, intuition, moral perception, concern for further normative criteria (e.g., coherence), application of abstract moral theories, etc.

During this step, the evaluator may also look beyond CI and shared ends and appeal to additional considerations, such as ethical principles, theories, or values that diverge from those accepted by either the society under study or her interlocutor. If she cannot convince her interlocutor to accept the considerations she is appealing to, her argument will presumably have little impact on that particular interlocutor,¹⁵ but she will have reached an evaluative conclusion for herself. Such an evaluative conclusion will be richer and better informed as a result of completing Steps 1 to 4 of the CI analysis.

3.2 Important Points of Divergence Between the Proposed Procedure and Nissenbaum's CI Decision Heuristic

The general procedure for evaluating technologies that I have proposed in this article differs substantially from the CI decision heuristic that Nissenbaum proposes (2009, pp. 182–183). In summary, Nissenbaum's (2009) heuristic has the

¹⁵ At least not directly, via the reasons that the evaluator offers. It may nonetheless affect the interlocutor via other routes, such as via a process of testimonial deference.

evaluator (1) identify a practice that has been altered by a technology, (2) identify the prevailing context for the practice and contexts that are nested within that broader context, (3–5) identify entrenched (privacy) norms that the technology may have influenced, using the parameters that Nissenbaum has identified as building blocks for privacy norms, (6) perform a *prima facie* assessment of the altered practice: “A breach of informational norms yields a *prima facie* judgment that contextual integrity has been violated” (p. 182), (7) perform evaluation I: “Consider moral and political factors affected by the practice in question” (p. 182), (8) perform evaluation II: “Ask how the system or practices directly impinge on values, goals, and ends of the context. In addition, consider the meaning or significance of moral and political factors in light of contextual values, ends, purposes, and goals” (p. 182), and (9) draw a conclusion: “On the basis of these findings, contextual integrity recommends in favor of or against systems or practices under study.” (pp. 182–183).¹⁶

In my proposed procedure, we distinguish *prima facie* CI reductions (i.e., threats to entrenched elements of normative life), genuine CI reductions (i.e., threats to elements of normative life *that advance shared ends*), and subsets of genuine CI reductions that the evaluator judges to be *problematic*. This does not map perfectly onto Nissenbaum’s use of the CI concept. In some places, Nissenbaum’s summary of her view suggests that we can diagnose CI violations just by looking for norm violations, without assessing whether the norms advance shared ends: “Contextual integrity is defined in terms of informational norms: It is preserved when informational norms are respected and violated when informational norms are breached” (2009, p. 140). In Nissenbaum’s, 2018 summary of her view, CI (and appropriateness, too) appears to be evaluable simply with reference to whether entrenched norms are violated, without requiring further assessment of the norms themselves. However, it is possible that her way of speaking in these instances rests on an assumption that the norms involved promote shared ends, and the ends involved are legitimate. In any case, I propose distinguishing *prima facie*, genuine, and problematic CI reductions in the way Table 1 depicts because I believe it is useful to explicitly break up the process of evaluation into multiple steps.

Another important difference is that Nissenbaum’s heuristic recommends identifying a singular *prevailing context*. By contrast, the general procedure proposed in this article supposes that there will typically be multiple important contexts to consider. This makes the evaluative task more complex, but I believe it is necessary because so many technologies involve or have implications for multiple contexts. For instance, think of the question of whether a social media site most resembles a context of friends interacting, a workplace, a networking venue, a news media entity, etc. Social media has potential implications for each these contexts, and each involves different and potentially conflicting sets

¹⁶ In her description of (9), Nissenbaum adds, “(In rare circumstances, there might be cases that are sustained in spite of these findings, accepting resulting threats to the continuing existence of the context itself as a viable social unit).”

of norms, roles, and other normative elements. Thus, when evaluating social media sites, we have reason to consider their relationship to entrenched normative elements in multiple contexts.

4 The Case of the Scythe, Gleaning, and Pasture Norms

I now want to offer a historical case as an illustration of a situation with CI (broadly construed) that was threatened by a socially disruptive technology.

In some communities in Europe during the Middle Ages and early modern period, entrenched sets of norms and practices developed around the harvest context, including norms regulating gleaning.¹⁷ Gleaning is a practice in which, after a landowner's harvest is complete, others are permitted to take various harvest remainders for their own use. These remainders may include grain but also materials used for thatch and bedding. The specifics of the practice, including who may do the gleaning, when, and using what tools, have varied over time and place. For instance, sometimes gleaning was treated as the right of the poor and those incapable of undertaking more substantial agricultural work; sometimes gleaning was treated as a right of harvest-workers' families (Ault, 1961; Vardi, 1993). I want to suggest that in some regions a possible technological change—a switch from using the sickle to using the long-handled scythe—promised to disrupt entrenched harvest-related norms in a way that can be usefully interpreted as a threat to CI. In particular, in some regions of seventeenth- and eighteenth-century France, the use of the scythe for harvesting grain was met with enough opposition by interested parties that it was prohibited by custom and sometimes explicitly by law (Root, 1990, p. 353; Root, 1992, pp. 109–112; Bloch, 1966, pp. 46–47; see also Postan et al., 1965, pp. 155–156).

Although the long-handled scythe had been used in some regions of Europe since Roman rule (Roberts, 1979), it was initially used only for harvesting hay and not grain (Postan, et al., 1965, pp. 155–156). At the end of the thirteenth century, a new forging technique made it possible to tilt the scythe blade differently, prompting some to use scythes to harvest grains: oats and barley in the fourteenth century (Comet, 1997, p. 24), and subsequently wheat and rye (Roberts, 1979, p. 16). In general, for landowners' purposes, various factors might count in favor of or against use of the grain scythe in different economic and ecological conditions: In its favor, a harvest obtained using a scythe required fewer workers than the sickle, and the harvest could be brought in faster (reducing risks associated with bad weather) (Roberts, 1979); against it, the scythe required more steel than the sickle and hence was more expensive (Comet, 1997, p. 24), and the scythe required more strength to use than the sickle—harvesters using scythes were generally men and required a higher wage (Roberts, 1979).

¹⁷ I came across this historical episode via a brief discussion of it in Foley (2019).

Another important difference between the scythe and sickle, which mattered less (directly) to landowners and more to gleaners and the community at large, was that a scythe cut lower on the plant than the sickle: the scythe left less material in the field that gleaners could take for thatch, bedding, and fuel.¹⁸ Furthermore, in many areas, after gleaners had taken what they could, community members with livestock could use the fields as pasture; the scythe left less stubble in the field that could be used for pasturing purpose.

In regions of France with a tradition of open (unenclosed), narrow fields and community grazing (Bloch, 1966, pp. 35–48), in particular, whether a harvest was done with sickles or scythes clearly mattered for shared community ends. Bloch presents these regions as having a particular type of *agrarian regime*—an agricultural system characterized by a form of crop rotation as well as “an intricate complex of techniques and social relations” (1966, p. 35). Not all agricultural activity falls under an agrarian regime in his sense; there must be some level of (possibly informal) regulation involved. Of the agrarian regime of this region and time, Bloch writes, “It is difficult to imagine a more coherent system, and even in the nineteenth century its ‘harmony’ could still arouse the grudging admiration of the most sophisticated critics” (1966, p. 44). Over a period of decades and sometimes centuries, the communities had developed relatively stable, interconnected systems of agricultural norms, traditional entitlements, practices of community deliberation, and so on. Presumably, many of these elements of the social system were products of power struggles, and many advanced the interests of sub-groups without thereby benefiting anyone else. But some of the norms served a function for the community as a whole: community-regulated practices of coordinated crop rotation helped preserve soil fertility and likely helped with pest management; rules governing community pasturing helped ensure a sufficient source of food for the livestock of community members while limiting overgrazing; and gleaning functioned as a de facto insurance mechanism for those who had suffered temporary misfortune (Root, 1992, p. 111), aiding not only those individuals themselves but also those who benefitted from their future labor or other contributions to the community.

In sum, within this agrarian regime, important and interconnected community practices relied on particular uses of the landowner’s land. As a result, if landowners made a change to the harvesting technologies used on their land, it could disrupt practices and norms in a way that threatened community ends. I claim that this is a case in which a system with CI was threatened by a change in technology use. Here, the norms and values involved relate not to privacy and information

¹⁸ Another interesting difference was that if extra effort was not taken to ensure otherwise, use of the scythe left more *grain* on the field than the sickle—the scythe made a mess (Comet, 1997, p. 24). Had this been left to gleaners, gleaners would have been better off when harvests occurred via scythe than via sickle. But the substantial loss of grain produced by the scythe meant that if landowners were going to switch to using the scythe, they had a significant incentive to attempt to obtain that lost grain themselves (Hussey, 1997, pp. 66–67). This could be done by hiring workers to rake the field following the scything, or by giving workers a share of what they collected if they raked or hand-gathered what the scythe had left on the field.

flows but rather to control over land, labor obligations and rights, obligations to neighbors, charity, and rights to food and other resources. In this case, community members avoided the threat to CI by prohibiting the new use of the technology.¹⁹ Of course, prohibition is not the only possible response. An alternative would have been for the community to accept the new use of the scythe, and over time to reach a new form of CI by modifying norms, practices, and other elements of the system as the consequences of the technological shift became apparent. Perhaps the wealthy would need to donate more to church coffers, while the church would do more to help the laborers and the poor who had been deprived of the gleaning resources on which they relied. Additional compensatory action would be needed to ensure that livestock acquired sufficient food, and so on.

I have supplied only a short sketch of this historical episode. There are many more internal and external players and causal factors involved than can be discussed here.²⁰ Eventually, the agricultural regime of this region of France did change, for multiple reasons. What I have sought to make plausible is that these communities' agricultural regime, and their encounter with a possible technological change, offers an example of a context that had a degree of integrity that was threatened by a socially disruptive technology.

To be clear, I do not intend to romanticize these agricultural systems or to claim that the communities involved were admirable, all things considered. It is entirely consistent with the general CI approach to characterize a context as having a high degree of integrity with regard to some shared values, but at the same time to argue that in many other ways the community failed to live up to values they professed to have (e.g., perhaps the maintenance of inequalities in some of these communities was incompatible with their Christian values), or to argue that some norms or practices in the community (e.g., the subjugation of women) were morally wrong, regardless of the ends the community members possessed.

5 Some Attractions of a Functional, Contextual Integrity-Inspired Approach to the Evaluation of Technologies

5.1 Context-Sensitivity and Adaptability

A defining feature of Nissenbaum's CI approach is that it is context-sensitive. The procedure I propose supplies an even greater degree of adaptability than

¹⁹ Presumably, at least some landowners were tempted to use the scythe or it would not have been prohibited. There is an interesting question about what motivated those landowners and why their preference diverged from that of the community—was it an instance of short-sightedness, lack of awareness of how their decision might affect the larger system, or perhaps a tragedy-of-the-common-type case (e.g., where each individual landowner prefers that he but few others use the scythe)? Similar cases are discussed in Elinor Ostrom's *Governing the Commons* (1990).

²⁰ In the eighteenth century, other important players included the monarchy, various reformers who aimed to increase agricultural production, and elements of the King's bureaucracy whose concern with tax collection and minimizing the number of itinerant poor motivated them to protect commons traditions (Root, 1992).

CI by also enabling us to consider the relationship between technologies and a wide range of values. In her original proposal, Nissenbaum focused on just a single value or type of concern, namely, privacy. By generalizing CI, we can think both in terms of a normative domain (a value or set of concerns), as well as a context. Corresponding to Nissenbaum's privacy-CI, we can have justice-CI, care-CI, liberty-CI, and so on, depending on which broad normative domains concern us, or which are most likely to be disrupted by a particular kind of technology. In cases in which technologies modify information flows, Nissenbaum argued that privacy becomes relevant. With a general concept of CI, we can extend this idea—for instance, perhaps in cases where technologies affect distribution of resources, we can anticipate that fairness becomes relevant; in cases involving expansion/restriction of capacities, we must consider liberty-related norms; or in cases involving threats to welfare, harm-related norms become relevant. For any given technology, we can mix and match from what we know about contexts and normative domains, to generate a set of normative elements—values, norms, and so on—that the technology may disrupt.

Thinking in terms of CI offers a method for analyzing technologies with reference to the ends of those involved and permits moving the level of analysis up and down in the hierarchy of nested contexts and systems that may interest us. It recognizes that individuals and groups have a multiplicity of ends, and that a change to a system can advance some ends and hinder others, such that one may want to make further changes if one does accept a technological change, and such that one may need to reconcile conflicts between ends by appealing to further ends. Furthermore, the proposed procedure is consistent with the idea that moral systems are dynamic—it accommodates the fact that a group may reevaluate its norms over time and modify them if its shared ends have changed.

Evaluative and deliberative questions about new technologies are asked by people in different positions and at different scales—e.g., at the family, neighborhood, school, city, or country level. The general procedure I have proposed can be applied at these various levels and by people in different positions. If the members or leaders of a social group have a conception of the group's ends, can obtain a sense of which norms and practices currently advance or support those ends, and can anticipate how the technology may affect those norms and practices (or if they can monitor its effects), then they can implement a version of the proposed procedure for technological evaluation. The notion of CI allows us to talk about technologies that are disruptive at different scales—a technology could be socially disruptive to a tribe or a nation, to a religious community, or to primary education. We can make sense of this using CI theory because we can talk about CI in each of these contexts: there are entrenched norms and practices, promoting shared ends, and participants in the context may dispute whether potential changes that threaten CI will be for the good, overall, or not.

5.2 A Principled Presumptive Conservatism, with a Mechanism for Reevaluation

CI theory explicitly features an initial presumption in favor of the status quo. Nissenbaum (2004) states:

A presumption in favor of the status quo for informational norms means we initially resist breaches, suspicious that they occasion injustice or even tyranny. We take the stance that the entrenched normative framework represents a settled rationale for a certain context that we ought to protect unless powerful reasons support change. The settled rationale of any given context may have long historical roots and serve important cultural, social, and personal ends. (p. 127)

I view CI's presumptive conservatism as another attractive feature of the theory. It reflects some basic facts, including the following: advancement of many individual and shared ends relies on existing entrenched, human-influenced aspects of our world; we exist in a world of complex systems—biological, social, ecological—that we do not understand; human ends are complex and it is difficult to intentionally construct systems that advance those ends; and there are limits to humans' capacity to adapt to changing environments and limits to the speed with which we can do so. Benthall et al. observe, "To the extent that CI has a conservative impulse, it is to warn against the agitation caused by disruptive technologies that change the environment too quickly for social evolution to adapt" (2017, p. 47). The general procedure for evaluating technologies that I have proposed retains this feature.

CI theory thus supports something like the lesson of Chesterton's fence: if we do not have a good understanding of why an entrenched element of our normative life exists, we should be cautious as we go about modifying or eliminating it, because it may have been implemented or may have arisen and been maintained because it serves a function with respect to shared ends. Chesterton writes:

In the matter of reforming things, as distinct from deforming them, there is one plain and simple principle... There exists in such a case a certain institution or law; let us say, for the sake of simplicity, a fence or gate erected across a road. The more modern type of reformer goes gaily up to it and says, 'I don't see the use of this; let us clear it away.' To which the more intelligent type of reformer will do well to answer: 'If you don't see the use of it, I certainly won't let you clear it away. Go away and think. Then, when you can come back and tell me that you do see the use of it, I may allow you to destroy it' ... the truth is that nobody has any business to destroy a social institution until he has really seen it as an historical institution. If he knows how it arose, and what purposes it was supposed to serve, he may really be able to say that they were bad purposes, or that they have since become bad purposes, or that they are purposes which are no longer served. (1929, pp. 26–27)

When analyzing CI, though, we must also think beyond individual or group intentions and reasons. If we are considering changing an old or entrenched element of the world, it is also relevant to consider whether that element, though not put in place intentionally, is a product of a cultural evolutionary process and still serves a valuable function within our system, or is an important background factor on which other valuable parts of our system rely. CI allows us to talk about the fact that using technology to make a change to an established system with relatively stable dynamics, which has been adapted over time through a multicomponent, multilevel process

of intentions, planning, learning, feedback, cultural and natural selection, etc., can produce disruptions that upset the usual dynamics—and consequently upset important ends. Given a disruption to such a system, one can expect there will be a period of adjustment—possibly a long period—before a new stable set of dynamics or equilibrium of some type (conducive to advancing shared ends) emerges. Language for this kind of phenomenon is especially important for the purpose of understanding the impacts of socially or morally disruptive technologies, because one way in which they are likely to be disruptive is via disruptions to dynamic systems that are relatively well adapted to past circumstances.

Crucially, however, the procedure I have proposed (like Nissenbaum's theory) is not entirely conservative—it recognizes that existing norms, contexts, and values can be flawed or deeply wrong, and it provides a procedure for the reevaluation and rethinking of norms and values. In other words, it takes existing systems as a starting point, and it encourages us to examine what role the elements of the system are currently playing in our moral and social systems—they may be there for a reason and may still be serving some function or producing some effect that we value, or they may not.

6 Objections and Limitations

An important line of objection that one might raise against this proposal has to do with the fact that human ends change over time. In particular, several scholars have emphasized that technological change has special potential to influence values and other ends (Swierstra, et al., 2009; Swierstra, 2013; van de Poel, 2021; Verbeek, 2011). One might ask whether one should perform one's evaluation of a technological change not with one's own ends as a reference point, but with some other ends—perhaps those one is likely to acquire in the future or perhaps the ends that some distant future generations will possess. For instance, Swierstra et al. (2009) and Boenink et al. (2010) argue against "moral presentism," or the tendency to use only current morals (and not possible future morals) to evaluate technologies. My response to this is that one can only be motivated to advance ends that one does not yet have if doing so is recommended by ends that one currently has. If an evaluator believes there is a possibility that their conception of their ends is erroneous or that it may change in the future for good reasons, it will make sense by their own lights to consider the ways in which their (conception of their) ends may change in the future, and to potentially update their (conception of their) ends in advance of the future shift they anticipate undergoing.²¹ Performing a CI analysis can itself cause the evaluator to undergo changes in their own ends—e.g., in the course of considering the values of others, one may come to believe that those values are right. The procedure does not assume that the ends an evaluator begins with are the same as the ends they possesses at Step 5. Much more remains to be said on this topic, but it must be left for future work.

²¹ On this topic, see the literature on shifting preferences and changing selves, e.g., Pettigrew (2019) and Lambert & Schwenkler (2020). See also the related discussion in the literature on moral enhancement.

Another worry one might have about this approach is that it might seem to have a flavor of ineffectual relativism about it. If the evaluator possesses misguided ends, one might think their evaluation will not help anyone move closer to what is right. The reason that the evaluation procedure ends in an appeal to the evaluator's own ends has to do with the fact that the approach is meant to be practical, in the sense that it leads to a conclusion that can motivate the evaluator and anyone else who shares the ends they invoke in Step 5. A libertarian may well reach a different conclusion in Step 5 than a limitarian (Robeyns, 2022). The approach is valuable despite this, because performing the CI analysis through Step 4 is edifying for everyone, and because each person can engage with the arguments and values appealed to during Step 5 to ascertain whether to accept the evaluator's evaluative conclusion or not.

An important limitation of the approach as it has been presented thus far is that it does not address at any length the topic of multiple conflicting shared ends. In a case where a technological change will advance some shared ends and hinder other shared ends, how should one proceed? A corresponding question confronts the individual decision-maker who finds that a technological change will advance some of her ends and hinder others, and to some extent this type of issue also poses a problem for Nissenbam's account, because considerations pertinent to privacy-related shared ends will sometimes conflict. This is a bigger challenge for the general CI approach than for Nissenbaum's privacy theory, though, because the general CI approach is meant to be applied to the full range of human values.²² By necessity, a fuller account of the general CI approach must address this important and difficult question. Resources for addressing the problem may be obtained from the existing literature on hard choices (Chang, 2017) and trade-offs in ethics, but development of the general CI approach in this respect is a task that must be left for the future.

7 Conclusion

For the purpose of helping us evaluate and respond to new technologies, it would be beneficial to have an overarching approach that can be adapted in different circumstances to incorporate consideration of a variety of values, norms, and other normative elements. I have proposed a contextual integrity-inspired general procedure that can serve such a function. Much philosophical and social scientific work remains to be done in order to render the proposed procedure usable in practice across a variety of contexts. In particular, a significant next step will be to develop accounts of parameters for norms from various normative domains other than privacy (e.g., fairness, care, liberty). Nonetheless, in this article, I hope to have sketched in broad strokes how a contextual integrity-inspired general procedure can guide evaluation of technological change.

²² Thanks to an anonymous reviewer for raising this point.

Acknowledgements For comments on previous versions of this paper, I would like to thank the members of the Foundations & Synthesis line of the Ethics of Socially Disruptive Technologies program, particularly Jeroen Hopster, Michael Klenk, Sven Nyholm, and Cecilie Eriksen; the Philosophy & Ethics research group at Eindhoven University of Technology, particularly Wybo Houkes, Matthew Dennis, Emily Sullivan, and Lily Frank; and participants in a 2019 workshop on the Ethics of Behaviour Prediction and Behavioural Influence at Oxford University. Thanks also to several anonymous reviewers.

Funding This research has been supported by the Netherlands Organisation for Scientific Research under grant number 016.Veni.195.513; the Ethics of Socially Disruptive Technologies research program, which is funded through the Gravitation program of the Dutch Ministry of Education, Culture, and Science and the Netherlands Organisation for Scientific Research under Grant number 024.004.031; a fellowship at the Cornell Tech Digital Life Initiative funded by NSF Grant SES-1650589 (PI: Helen Nissenbaum); a visit to the Simons Institute at the University of California, Berkeley; and a visit to the Leverhulme Centre for the Future of Intelligence at the University of Cambridge.

Declarations

Ethical Approval Not applicable.

Consent to Participate Not applicable.

Consent for Publication Not applicable.

Competing Interests The author declares no competing interests.

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

- Anderson, E. (2015). Moral bias and corrective practices: A pragmatist perspective. *Proceedings and Addresses of the APA*, 89, 21–47.
- Ault, W. O. (1961). By-laws of gleaning and the problems of harvest. *The Economic History Review*, 14(2), 210–217.
- Baker, R. (2013). *Before bioethics: A history of American medical ethics from the colonial period to the bioethics revolution*. Oxford University Press.
- Baker, R. (2019). *The structure of moral revolutions*. MIT Press.
- Benthall, S., & Haynes, B. (2019, August 19). *Contexts are political: Field theory and privacy*. Symposium on Applications of Contextual Integrity, Berkeley, CA.
- Benthall, S., Gürses, S., & Nissenbaum, H. (2017). Contextual integrity through the lens of computer science. *Foundations and Trends in Privacy and Security*, 2, 1–69.
- Bloch, M. (1966). *French rural history: An essay on its basic characteristics*. University of California Press.
- Boenink, M., Swierstra, T., & Stemerding, D. (2010). Anticipating the interaction between technology and morality: A scenario study of experimenting with humans in bionanotechnology. *Studies in Ethics, Law, and Technology*, 4(2), 1–38.
- Boyd, R., & Richerson, P. J. (1992). Punishment allows the evolution of cooperation (or anything else) in sizable groups. *Ethology and Sociobiology*, 13(3), 171–195.

- Brey, P. (2000). Method in computer ethics: Towards a multi-level interdisciplinary approach. *Ethics and Information Technology*, 2(2), 125–129.
- Brey, P. (2017). Ethics of emerging technology. In S. O. Hansson (Ed.), *The ethics of technology: Methods and approaches* (pp. 175–191). Rowman & Littlefield.
- Brey, P., et al. (2019). *Ethics of socially disruptive technologies*. Project proposal for Netherlands Organization of Scientific Research.
- Brey, P. A. (2012). Anticipatory ethics for emerging technologies. *NanoEthics*, 6(1), 1–13.
- Bursztyn, L., González, A.L., & Yanagizawa-Drott, D. (2018). *Misperceived social norms: Female labor force participation in Saudi Arabia* (National Bureau of Economic Research Working Paper No. w24736). <https://doi.org/10.3386/w24736>
- Carlsen, H., Dreborg, K. H., Godman, M., Hansson, S. O., Johansson, L., & Wikman-Svahn, P. (2010). Assessing socially disruptive technological change. *Technology in Society*, 32(3), 209–218.
- Chang, R. (2017). Hard choices. *Journal of the American Philosophical Association*, 3(1), 1–21.
- Chesterton, G.K. (1929). *The thing: Why I am a Catholic*. Aeterna Press.
- Collingridge, D. (1980). *The social control of technology*. Martin's Press.
- Comet, G. (1997). Technology and agricultural expansion in the Middle Ages: The example of France north of the Loire. *Medieval Farming and Technology*, 11–39.
- Curry, O. S., Mullins, D. A., & Whitehouse, H. (2019). Is it good to cooperate? Testing the theory of morality-as-cooperation in 60 societies. *Current Anthropology*, 60(1), 47–69.
- Dihal, K., Hollanek, T., Rizk, N., Weheba, N., & Cave, S. (2021). Imagining a future with intelligent machines: A Middle Eastern and North African perspective. *The Leverhulme Centre for the Future of Intelligence, University of Cambridge*. <https://www.ainarratives.com/resources/mena-report>. Accessed 10 Aug 2022.
- Dutilh Novaes, C. (2020). Who's afraid of adversariality? Conflict and cooperation in argumentation. *Topoi*, 40, 873–886.
- Foley, M. (2019). *Farming for the long haul: Resilience and the lost art of agricultural inventiveness*. Chelsea Green Publishing.
- Forge, J. (2010). A note on the definition of “dual use.” *Science and Engineering Ethics*, 16(1), 111–118.
- Future of Life Institute. (2022). *A project by the future of life institute*. <https://worldbuild.ai/about/>. Accessed 10 Aug 2022.
- Grasswick, H. (2018). Feminist social epistemology. In E.N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/fall2018/entries/feminist-social-epistemology/>. Accessed 10 Aug 2022.
- Haidt, J., & Joseph, C. (2008). The moral mind: How five sets of innate intuitions guide the development of many culture-specific virtues, and perhaps even modules. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind, volume 3: Foundations and the future*. Oxford University Press.
- Hankins, K., & Vanderschraaf, P. (2021). Game theory and ethics. In E.N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/win2021/entries/game-ethics/>
- Hansson, S. O. (2011). Coping with the unpredictable effects of future technologies. *Philosophy & Technology*, 24(2), 137–149.
- Hare, B., & Woods, V. (2020). *Survival of the friendliest: Understanding our origins and rediscovering our common humanity*. Random House.
- Henrich, J. (2016). *The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter*. Princeton University Press.
- Hopster, J. (2021a). What are socially disruptive technologies? *Technology in Society*, 67(101750), 1–8.
- Hopster, J. (2021b). The ethics of disruptive technologies: Towards a general framework. In J.F. de Paz Santana & D.H. de la Iglesia (Eds.), *Advances in intelligent systems and computing*. Springer.
- Hopster, J., Arora, C., Blunden, C., Eriksen, C., Frank, L., Hermann, J., Klenk, M., O'Neill, E. and Steinert, S. (2022). Pistols, pills, pork and ploughs: The structure of technomoral revolutions. *Inquiry*, 1–33.
- Hussey, S. (1997). ‘The last survivor of an ancient race’: The changing face of Essex gleaning. *The Agricultural History Review*, 45(1), 61–72.
- Jacobs, N., & Huldgtren, A. (2021). Why value sensitive design needs ethical commitments. *Ethics and Information Technology*, 23(1), 23–26.
- Kim, T. W., & Werbach, K. (2016). More than just a game: Ethical issues in gamification. *Ethics and Information Technology*, 18(2), 157–173.

- Kleiman-Weiner, M., Saxe, R., & Tenenbaum, J. B. (2017). Learning a commonsense moral theory. *Cognition*, 167, 107–123.
- Klenk, M. (2021). How do technological artefacts embody moral values? *Philosophy & Technology*, 34(3), 525–544.
- Kudina, O., & Verbeek, P. P. (2019). Ethics from within: Google Glass, the Collingridge dilemma, and the mediated value of privacy. *Science, Technology, & Human Values*, 44(2), 291–314.
- Lambert, E., & Schwenker, J. (Eds.). (2020). *Becoming someone new: Essays on transformative experience, choice, and change*. Oxford University Press.
- Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1(11), 501–507.
- Morley, J., Floridi, L., Kinsey, L., & Elhalal, A. (2021). From what to how: An initial review of publicly available AI ethics tools, methods and research to translate principles into practices. In L. Floridi (Ed.), *Ethics, governance, and policies in artificial intelligence* (pp. 153–183). Springer.
- Nickel, P. J. (2020). Disruptive innovation and moral uncertainty. *NanoEthics*, 14(3), 259–269.
- Nickel, P. J., Kudina, O., & van de Poel, I. (2021). Moral uncertainty in technomoral change: Bridging the explanatory gap. *Perspectives on Science*, 30(2), 260–283.
- Nissenbaum, H. (2009). *Privacy in context*. Stanford University Press.
- Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, 79(1), 119–158.
- Nissenbaum, H. (2018). Respecting context to protect privacy: Why meaning matters. *Science and Engineering Ethics*, 24(3), 831–852.
- O'Neill, E. (2017). Kinds of norms. *Philosophy. Compass*, 12(5), 1–15.
- Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.
- Palm, E., & Hansson, S. O. (2006). The case for ethical technology assessment (eTA). *Technological Forecasting and Social Change*, 73(5), 543–558.
- Pettigrew, R. (2019). *Choosing for changing selves*. Oxford University Press.
- Pols, A. J. (2013). How artefacts influence our actions. *Ethical Theory and Moral Practice*, 16(3), 575–587.
- Postan, M. M., Rich, E. E., & Miller, E. (Eds.). (1965). *The Cambridge economic history of Europe, Vol. Economic organization and policies in the Middle Ages*. Cambridge University Press.
- Roberts, M. (1979, March). Sickles and scythes: Women's work and men's work at harvest time. *History Workshop Journal*, 7(1), 3–28.
- Robeyns, I. (2022). Why Limitarianism? *Journal of Political Philosophy*, 30(2), 249–270.
- Root, H. (1990). The "moral economy" of the pre-revolutionary French peasant. *Science & Society*, 54(3), 351–361.
- Root, H. L. (1992). *Peasants and king in Burgundy: Agrarian foundations of French absolutism*. University of California Press.
- Rule, J. B. (2019). Contextual integrity and its discontents: A critique of Helen Nissenbaum's normative arguments. *Policy & Internet*, 11(3), 260–279.
- Sunstein, C. R. (2019). *How change happens*. MIT Press.
- Swierstra, T., Stemerding, D., & Boenink, M. (2009). Exploring techno-moral change: The case of the ObesityPill. In P. Solli & M. Düwell (Eds.), *Evaluating new technologies: Methodological problems for the ethical assessment of technology developments* (pp. 119–138). Springer.
- Swierstra, T. (2013). Nanotechnology and technomoral change. *Ethics & Politics*, 15(1), 200–219.
- Tomasello, M. (2019). *Becoming human: A theory of ontogeny*. Belknap Press.
- van de Poel, I. (2021). Design for value change. *Ethics and Information Technology*, 23(1), 27–31.
- van de Poel, I., & Kudina, O. (2022). Understanding technology-induced value change: A pragmatist proposal. *Philosophy & Technology*, 35(2), 1–24.
- van de Poel, I. (2011). Nuclear energy as a social experiment. *Ethics, Policy & Environment*, 14(3), 285–290.
- van de Poel, I. (2016). An ethical framework for evaluating experimental technology. *Science and Engineering Ethics*, 22, 667–686.
- van Eijndhoven, J. C. (1997). Technology assessment: Product or process? *Technological Forecasting and Social Change*, 54(2–3), 269–286.
- Vardi, L. (1993). Construing the harvest: Gleaners, farmers, and officials in early modern France. *The American Historical Review*, 98(5), 1424–1447.

- Verbeek, P.P. (2011). *Moralizing technology: Understanding and designing the morality of things*. University of Chicago Press.
- Whittlestone, J., Nyrup, R., Alexandrova, A., & Cave, S. (2019, January). The role and limits of principles in AI ethics: Towards a focus on tensions. *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 195–200.

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