

Deliberate Ignorance

Present and Future

Christoph Engel and Ralph Hertwig

In 2016, as we penned our conceptual and explorative article on the phenomenon of deliberate ignorance (Hertwig and Engel 2016), we felt a bit like explorers setting sail for an unknown destination. Our spirits were high and we were ready for an intellectual adventure. Fascinated by the richness of the phenomenon, we soon noticed that others in the fields of economics, sociology, law, and medicine had been travelling in a similar direction, guided by terms such as “information avoidance,” “willful blindness,” and even deliberate ignorance (e.g., Robbins 1990). Still, we found the vast territory of deliberate ignorance to be mostly uncharted and hoped that our article would serve as an inspiring travelogue, describing our attempt to survey the lay of the land and inviting others to join us in exploring the phenomenon further.

To our delight, our excitement proved contagious, as demonstrated by the lively discussions that emerged at this Ernst Strüngmann Forum. In this final chapter, we reflect on specific areas that have left their mark on us both. We begin with an observation that ran throughout all discussions, and then present our thoughts, organized around the four thematic areas of the Forum.

The Power and Perils of Interdisciplinarity

Statements calling for interdisciplinary analysis of a research topic are ten a penny. Yet, to comprehend the phenomenon of deliberate ignorance and its implications requires exactly that. Deliberate ignorance is a human behavior with strong normative and institutional implications: it plays out individually and collectively, is subject to temporal dynamics and changes in norms, and can be investigated by means of experiments, surveys, interviews, modeling, or archival work. No single discipline has a full command of these tools, concepts, and dimensions. Against this background, we submitted a proposal to the Ernst Strüngmann Forum, which has a reputation for promoting truly interdisciplinary discourse, and were delighted when our proposal was accepted.

Interdisciplinary discourse is often hard. Typically, the humanities and social sciences do not share the behavioral sciences' commitment to logical positivism, behavioral experimentation, modeling, and quantification. The paradigmatic theories of individual and collective behaviors (e.g., expected utility theory and game theory) adopted by many economists and psychologists are not necessarily compatible with, say, the terminology and explanatory concepts used by social scientists, or with historians' foci on the dynamics of change across time and the intricate interdependency of individual and collective processes. Relatedly, the descriptive concept of deliberate ignorance may feel like an intellectual affront to a historian for whom the study of history is at its core enlightenment, understood as a necessary condition of human liberty (Nipperdey 1980). Finding common ground amidst all this was not easy, as many discovered at the Forum. Equally, though, such intellectual provocation can also be the starting point for something new. We hope that this volume will be viewed as an attempt to submit the polymorphous phenomenon of deliberate ignorance to an analysis without borders. Others are necessary and should follow.

What Exactly Is Deliberate Ignorance?

Definitions are simplistic constructions with a purpose. They draw boundaries because classification has instrumental value. They can stimulate thought, guide investigation, enable understanding, suggest evaluations, and even inform the design of policy interventions. Whether or not a case belongs to the territory of deliberate ignorance is therefore not an ontological question. We do not mean to propose a cut-and-dried test analogous to that for blood types, where only, say, rhesus positive counts as the phenomenon in question. The definition must fit the intended research purposes.

Admittedly, these purposes are diverse and need not be fully aligned in their definitional implications. If the aim is to catalog ever more instances of deliberate ignorance, a wider definition may be preferable to ensure that no interesting case is overlooked. If the coverage is too wide, however, the concept risks losing its bite. In addition, it becomes increasingly difficult to understand the phenomenon, let alone to model it rigorously. Narrowing the scope may be necessary to detect the fine-grained structure of the phenomenon. Yet too narrow a definition may hinder normative appraisal of instances at the margin that are normatively no less troublesome than the prototypical ones. Finally, institutions are lumpy responses to lumpy perceived problems (North 1990). An institutional designer may therefore choose a different criterion for the trade-off between precision and breadth than would be used by a modeler or surveyor of deliberate ignorance. In some contexts, a narrow and precise definition may be necessary to design an institutional intervention and convince policy makers to implement it. In other contexts, a well-intended institution with an overly

narrow definition of the cause for intervention may miss the target and prove counterproductive.

Like many other concepts, deliberate ignorance consists of a hard core—represented by a repository of paradigmatic examples (see Appendix 14.1 in Krueger et al., this volume)—and a somewhat fuzzy periphery. Where to draw the line of demarcation will depend on the research question and its implied decision criterion. Below, we illustrate this point by reference to a few cases discussed at the Forum.

Heuristics and Deliberate Ignorance

Brown and Walasek as well as Kornhauser (both this volume) have asked whether the nonuse of information, which is not only known to exist but has been encoded in memory, implies deliberate ignorance. Defined thus broadly, the concept would also encompass any heuristic (e.g., Gigerenzer et al. 2011) that is used deliberately. Indeed, one definition of a heuristic is “a strategy that ignores part of the information, with the goal of making decisions more quickly, frugally, and/or accurately than more complex methods” (Gigerenzer and Gaissmaier 2011:454). There is a vast literature on heuristics, to which both of us have contributed. The debate remains controversial, with one key point of contention being the rationality of heuristics: Does heuristic decision making lead to more—or to more serious—errors than “rational” procedures, as defined by logic or statistical models, or can it perhaps outperform more complex strategies when applied in the right environments (the “ecological rationality” of heuristics; Gigerenzer and Gaissmaier 2011; Hogarth and Karelaia 2007; Spiliopoulos and Hertwig 2019)?

Is heuristic decision making an instance of deliberate ignorance? In our view, the answer must be no. Defining any act of ignoring information as deliberate ignorance would miss the novelty of the concept and overlook what makes it psychologically so interesting. A person may rely on a heuristic to offset cognitive limitations, as suggested by the heuristics-and-biases view of heuristic decision making (Kahneman 2011), or because they have learned that doing so often leads to better outcomes, as suggested by the ecological rationality view of heuristic decision making (Gigerenzer et al. 2011; Hertwig et al. 2019). Yet both reasons fail to produce the sense of perplexity often generated by instances of deliberate ignorance. As we wrote (Hertwig and Engel 2016:360):

We are particularly interested in situations where the marginal acquisition costs are negligible and the potential benefits potentially large, such that—from the perspective of the economics of information [...]—acquiring information would seem to be rational.

In our view, the tension resides precisely here: in the individual or collective choice to not consult information that could be acquired at negligible costs with potentially substantial benefits. Heuristic decision making is not typically characterized

by this tension. In fact, it is the nonuse of additional information that has potentially substantial benefits in heuristic decision making, as it helps to escape, for instance, the curse of overfitting (Gigerenzer and Brighton 2009).

Yet, the usefulness of a definition depends on its purpose and the research question under consideration. Those interested in normative reasons for not generating or retrieving available information from memory or the external world may benefit from a normative discussion of heuristic decision making and its relationship to deliberate ignorance. From some normative perspectives, knowing about the content of possibly troublesome information but ignoring it (heuristic decision making) may be even more problematic than knowing about the existence of possibly troublesome information but not exposing oneself to its content (deliberate ignorance).

Forgetting, Expungement, and Deliberate Ignorance

The deliberately ignorant individual has reason to expect that decision-relevant information is available, and yet chooses not to access that information. A functionally similar effect is achieved if the individual has been in possession of the information but through some means—such as (directed) forgetting—successfully removes it from memory before facing the need to act on it (see Schooler, this volume, for a discussion of the family resemblance between forgetting and deliberate ignorance).

Institutionalized deliberate ignorance via “purposeful forgetting” also plays a role in a legal context. A classic example is expungement, common in juvenile criminal court proceedings. It entails erasing or removing from state or federal records the information that a minor has been convicted for a crime, by sealing or destroying the record. Some jurisdictions also make it illegal for private parties such as future employers to request this information (e.g., New Hampshire Criminal Code Chapter Section 651:5 X (f)). Someone who has forgotten something cannot use that information in the context of a given decision, but they may be able to retrieve the semantic or biographical fact on another occasion. Expungement is more radical: information is completely removed from the record. The effect of this institutional intervention thus transcends the concrete instance. The focus is no longer on a specific decision, but on any decision that might be affected by the information in question.

Individual forgetting and institutional expungement are processes that occur after the fact. Both reset an individual, collective, or institutional information status to a state of ignorance. In this sense, one may argue that forgetting and expungement are not identical to deliberate ignorance. They stop the information in question from entering the system. Yet, phenomenologically, these processes appear to belong to the same functional family.

Deliberate Ignorance Is More Than Information Avoidance

The term “information avoidance”¹ has been used in the health domain to describe behaviors such as parents-to-be avoiding genetic testing on an unborn child, gay and bisexual men declining to learn their HIV status, or women avoiding regular pelvic checkups (Howell and Shepperd 2012; Sweeny et al. 2010). The term is also used in Golman et al.’s highly informative review article (Golman et al. 2017). The pivotal reason behind our alternative terminological choice is that “information avoidance” suggests what Howell and Shepperd (2012) characterized as “defensive responding” (p. 259), turning the act of not seeking or using available information into a form of psychological reactance, possibly even a “public health concern” (p. 262) in need of therapy.

Yet the choice not to know, as many of the examples discussed in this book attest, is not invariably dysfunctional. Indeed, there are numerous individual and institutional contexts in which deliberate ignorance affords a strategic advantage (e.g., Auster and Dana, this volume), attenuates the impact of selection biases (McCoun, this volume), constitutes a legal right (Berkman, this volume), or seems imperative as a means of keeping transitional societies together (Ellerbrock and Hertwig, this volume). Consider, for illustration, the choice of someone who has been diagnosed with a serious illness but decides not to ask about their prognosis. In one framing, this choice can be seen as a form of denial, the irrational avoidance of information fueled by fears of a bleak future. In another framing (Miller and Berger 2019), it is a positive choice:

When faced with serious illness, being able to make decisions about the flow of information is one of the most life-affirming things you can do. It’s a way to declare: I am alive and it’s still my right to choose what’s best for me.

Patients may legitimately want to shield themselves from a menacing, and not necessarily accurate, timeline against which each day is ticked off. Calling this behavior “deliberate ignorance” does not negate the associated detrimental effects that actively avoiding information may bring (see Krueger et al. and Teichmann et al., this volume). It is this inescapable ambiguity that makes the phenomenon of deliberate ignorance so interesting and, in our view, takes it beyond the normatively charged concept of information avoidance.

The example of not wanting to know one’s medical prognosis raises another issue. As a patient, the choice not to know does not necessarily mean that the information should be concealed from everybody. Often it means that the patient needs or wants somebody else (e.g., a physician or partner) to process it. A tool not (yet) implemented in standard medical practice enables patients to communicate to their physicians their preferences to know or not know, ranging from “Tell me everything” to “I don’t wish to know any information about my prognosis but I authorize you to speak with [blank] about my case

¹ To the best of our knowledge, Frey (1982) was the first to use the term “information avoidance” but others had previously referred to the avoidance of dissonant information (e.g., Mills 1965).

and for you to answer any question that this person may have about my likely prognosis and treatment” (see Miller and Berger 2019). Is this kind of delegation a way of deliberately ignoring information, or is it a way of deliberately using information?

Obfuscating Information

Obscuring access to information is a strategy used to overcome bias, increase impartiality in selection processes, and improve the quality of the decisions reached. Many orchestras, for example, utilize “blind auditions” in the preliminary rounds of the selection process for new members. Briefly, an auditioning musician plays behind a screen, so that visual cues are removed from consideration. The intent is to force the selection committee to focus on a candidate’s musical performance, not on gender, race, or a person’s affiliation to certain teachers or musicians² (see MacCoun as well as Krueger et al., this volume). In academia, the peer review process constitutes another example. Here, a double-blind procedure is used to increase fairness in the evaluation of scientific performance (MacCoun, this volume). Conceptually, both examples shield the identity of individuals, so that decision makers must rely on content.

A related strategy deliberately adds noise to information, as is a standard practice in some scientific disciplines. Here, empirical research aspires to make causal statements about a population, along the lines of “whenever process *A* is observed, phenomenon *B* will happen.” For the most part, though, scientists are unable to observe an entire population; they can only observe a sampling of it, and this sample may not be representative of the population. There is thus a risk of overinterpreting the sample and wrongly inferring a causal relationship from a random co-occurrence in the sample. This problem is known as *overfitting*. To safeguard against this, data is deliberately perturbed (e.g., random noise is added to each data point) before it is analyzed. Scientists will not report an observed effect unless it stands the test of this deliberate obfuscation (MacCoun, this volume).

Whether obfuscation qualifies as deliberate ignorance depends on the research question being asked. If one focuses solely on obfuscation, it is clear that the critical information is, by definition, available. This speaks against broadening the concept of deliberate ignorance. On the other hand, obfuscation increases both the cost of information retrieval and the risk that critical information will be missed. If obfuscation is, at least in principle, viewed as deliberate ignorance, its boundaries must be defined. Does that which counts as deliberate ignorance depend on the horizon of the intended recipient? A

² How good a hiring committee’s judgments about musical performance are without visual information is another question. Although auditory information is commonly assumed to be the most important information in the evaluation of music, experimental studies suggest that people “depend primarily on visual information when making judgments about music performance” (Tsay 2013:14580).

mere nuisance for the savvy user may prove an insurmountable obstacle for the novice. Should self-obfuscation be considered differently from inducing a third party to hide information, or benefiting from an outsider making it more difficult to access information?

The Construction of Reality

Berger and Luckmann (1991) hold that reality does not simply exist, it is socially constructed. Take, for instance, the construction of reality in the U.S. impeachment trial that is ongoing as we write this chapter. Did President Trump press legitimately for an investigation into a political rival's son in an effort to fight a corrupt elite? That is one socially constructed reality. Another is that he solicited the interference of a foreign government to help him win the 2020 election. Deliberately accepting one of the two constructed realities as the full and objective truth may have multiple effects. For instance, it allows people to remain comfortably ignorant of any other information and/or revelations that may emerge in the future that are more consistent with the other reality.

Should constructing and/or adopting one narrative in this way be subsumed under the heading of deliberate ignorance? Again, it depends on the research question. From an individualistic perspective, it may be important to distinguish between deliberate ignorance and the production and dissemination of a false narrative. The distinction between omission and commission may also cast a different normative light on such constructive efforts. Yet communication theorists are likely to argue that all forms of information processing are constructive. From this perspective, individuals do not mechanically integrate multiple pieces of information; instead, they make sense of communicative acts. Accordingly, drawing a strict boundary between looking the other way and telling an alternative narrative would thus be fallacious.

Summary

The concept of deliberate ignorance consists of a hard core of meaning and a fuzzy periphery. How the boundary is drawn depends on the research question as well as on the decision criterion. Many fascinating phenomena, some of which we have raised here, are located at the fuzzy periphery.

It is important to note that the term "deliberate ignorance" suggests a simple dichotomy: the decision maker either knows or chooses to remain ignorant. In many contexts, however, knowledge and ignorance are matters of degree. For example, a newspaper reader who reads the first paragraph of an article that describes graphically the impact of industrial-scale beef farming on animal welfare may choose not to read the whole article, not wanting to know any more about the provenance of their affordable supermarket beef. They thus know something but not everything, replacing complete ignorance with ambiguity.

How to Model Deliberate Ignorance?

Behavior defined as deliberate ignorance can be analyzed and modeled in more than one way.

Individualistic versus Holistic Models

The most fundamental conceptual divide lies between the individualistic and holistic perspectives. From an individualistic perspective, the focus is on an agent's decision to ignore information. Here, ignorance is a deliberate choice: The agent had the freedom to generate, retrieve, or use a defined piece of information. The question, however, becomes: Why did they choose not to do so?

A holistic perspective, by contrast, focuses primarily not on agency, but on why and how information remains unused that otherwise, in a counterfactual world, could have been available. The object of investigation is the social process by which a common understanding of social reality is forged. Why could a community not see an alternative interpretation of a set of facts? From this perspective, the construction of one understanding—and the non-construction of an alternative understanding—is a political act.

Individuals versus Higher-Order Agents

If an individualistic perspective is taken, the first step is to define the agent of interest. Is it the individual deciding in isolation, say, not to get tested for a genetic risk? Or is the agent a group of people, such as the individual's family, who would potentially also be affected by the test results? If the latter, does deliberate ignorance require that each member of the group not know?

Is the agent an institution, such as a firm? If so, whose knowledge is attributed to this legal entity: that of the board members only or that of any employee? Is it a sufficient condition for deliberate ignorance that the agent has not taken the necessary steps to ensure that decision-relevant knowledge become available?

Let us consider product liability: A manufacturer or seller of a defective product can be held liable for injuries arising from its use. Many legal orders define liability in such a way that it does not suffice for the manufacturer simply not to know about the defect. Rather, they must ensure that product development and production is organized in such a way that they would be alerted to any sign of a defect. Should it count as deliberate ignorance if, for instance, the safety analysis for a new flight control system is organized in such a way that critical flaws are not detected? Is a professional deliberately ignorant if they do not run tests that are standard in the profession? Is a group deliberately ignorant if it excludes from its membership an individual who would very likely have known critical information? Are corporate actors deliberately ignorant if

they fail to organize the flow of information within the corporation in such a way that relevant information is brought to the attention of the board?

Utility versus Strategic Interaction

Again, assuming an individualistic perspective, should the focus be placed on the motives of an individual who decides in isolation, or on understanding the strategic advantage of not knowing? The strategic perspective is technically more involved, as the conditions for equilibria need to be defined. Higher-order effects need to be considered. If, for instance, deliberate ignorance affords a strategic advantage, that advantage presupposes that the counterpart knows or believes that they are alone in having the relevant knowledge. The “game of chicken” is a paradigmatic example. In this model of conflict in game theory, two players head toward each other; if both stay on track, they will collide. The logic of the game is that the player who yields first loses the game. Yet if neither gives in, both will perish. Consider two vehicles: one self-driving, the other driven by a human. The vehicles approach an unmarked intersection and need to negotiate for priority. If the human driver believes that the self-driving car does not know that a collision could be fatal and will therefore press for priority, the onus is on the human driver to stop and (in game theoretic terms) lose.

Deontological Motives

Why might an individual in a nonstrategic situation prefer not to know? Individualistic modelers need assumptions about people’s motives to generate predictions. These motives may be utilitarian. The individual expects to be better off, in whatever sense, if they do not acquire a piece of information. They may, for instance, be concerned that they will be unable to not use the information and feel obliged to make choices they do not want to make. Alternatively, the motives may be deontological. An individual who treasures enlightenment values may feel morally obliged to access and use the information. At the same time, they may also hold privacy or secrecy in high regard and balk at the idea of invading another individual’s legally protected private sphere, even if they could exploit the knowledge gained to their own benefit.

What Are the Normative Implications of Deliberate Ignorance?

Deliberate ignorance eludes categorical normative conclusions and recommendations. Its manifestations are neither always normatively suspect nor always in accord with principles of ethics and rationality. Recently, it has been asked whether deliberate ignorance calls for interventions to protect the interests of those who desire to ignore information and those who do not desire so (Sharot and Sunstein 2020). We suggest that the high degree of context specificity

requires case-by-case analysis, thus making normative investigations of deliberate ignorance intriguing and regulatory interventions challenging.

Deliberate Ignorance: The Problem or the Answer?

Recent years have seen alarming developments in the form of deepening ideological divides and rising political polarization. In many countries, politicians, activists, and indeed voters appear to be deeply divided on issues such as inequality and immigration, with the divisions falling increasingly along party lines (e.g., DellaPosta et al. 2015; Iyengar and Westwood 2015; Sides and Hopkins 2015). One obvious fear is that this dynamic of polarization is intimately connected with deliberate ignorance. If asked to name a single problematic aspect of human reasoning that overrides all others, many psychologists will probably cite confirmation bias (see, e.g., Evans 1989): the tendency to seek or interpret evidence “in ways that are partial to existing beliefs, expectations, or a hypothesis in hand” (Nickerson 1998:175). It is easy to see how this tendency insulates people from views that contradict their preexisting beliefs, creating a fertile ground for political polarization.³ As one possible cause and motive of deliberate ignorance, the confirmation bias can foster undesirable behaviors.

At the same time, disinformation, propaganda, and fakery—particularly, but not only, in the digital ecosystem—are matters of growing concern worldwide (see Lewandowsky, this volume). According to a large-scale analysis of Twitter data, the “amount of false news online is clearly increasing” (Vosoughi et al. 2018:1150). The persuasive power of a falsity resides, among other factors, in the insidious fact that false information tends to be novel, and novelty elicits what is, under normal circumstances, an adaptive response: it grabs people’s attention. Analyzing all 126,000 major news stories distributed on Twitter from 2006 to 2017 and verified to be true or false, Vosoughi et al. found that the truth simply cannot compete with hoax and rumor. Falsehood consistently dominates the truth on Twitter: it reaches more people, penetrates deeper into social networks, and spreads much faster. When “falsehood flies, and the truth comes limping after it,” as Jonathan Swift so elegantly wrote three centuries ago, the competence to discern true from false news becomes essential. By the same token, the competence to exercise deliberate ignorance is becoming a critical civic skill. Once a person, a news source, a website, or an organization has been identified as regularly communicating falsity, users need to resist it. They need to withstand the temptation to fall for novelty, surprise, and the deceptive promise of relevance. Here, deliberate ignorance is anything but the

³ Interestingly, a recent study found initial evidence of possible backfire effects of exposing people to opposing views on social media. Attempts to introduce users to a broad range of opposing political views on social media sites such as Twitter might not only be ineffective but also counterproductive: they may actually increase political polarization (Bail et al. 2018).

sign of an intellectually incurious and lethargic cognitive system that yearns for comfort and consistency; it requires executive control and a system that strives for veracity rather than consistency. This is one type of deliberate ignorance that we had in mind when referring to its function as a “cognitive sustainability and information-management device” (Hertwig and Engel 2016). Deliberate ignorance, therefore, can be both the problem and the solution.

Who Decides and the Problem of Externalities

Another key normative aspect of deliberate ignorance is that the choice not to know often affects the well-being of others (and, to use the terminology of economics, leads to externalities), delegates responsibility, and ultimately raises the issue of who has the (political) power to decide. Consider, for illustration, Berkman’s (this volume) discussion of the right not to know one’s genetic makeup and the debate that has broken out in the medical community over this right. In response to the wider availability and improved utility of large-scale genomic sequencing (e.g., relating an increasing number of genetic variants to clinical phenotypes), the American College of Medical Genetics and Genomics (ACMG) in 2013 issued a recommendation for the handling of “incidental findings” (Green et al. 2013). These are pieces of “information (typically clinically significant and medically actionable) that is generated during a test or procedure but which does not relate to the original purpose for which the test or procedure was conducted” (Berkman, this volume, p. 200). The recommendation was that labs should actively search “for a ‘minimum list’ of variants that predispose patients to risk for disorders that ‘would likely have medical benefit for the patients and families of patients undergoing clinical sequencing’ ” (Berkman, this volume, p. 202). But how could such an active search be aligned with a right not to know? As Berkman describes, the ACMG Working Group controversially argued against soliciting patient preferences on receiving (or not receiving) incidental findings. In other words, patients would no longer be given the choice not to learn about clinically important and actionable findings. Berkman (this volume, p. 213) lays out the grounds for this recommendation as follows:

It is a vexing problem to possess genetic information that one deems to be clinically important, but to be precluded from disclosing it because a patient has exercised their RNTK. These medical professionals are apt to experience what we can colloquially call the “I-can’t-sleep-at-night” problem. More technically, they are experiencing a phenomenon known as moral distress.

Without wanting to downplay the physicians’ distress, we see an irony in the Working Group’s hardnosed paternalistic recommendation; namely, the privilege not to know was to be transferred from one stakeholder to another. Under the status quo, it was the patient’s right not to be informed about incidental findings; the physician had to carry the potentially distressing burden of

knowing a patient's genetic predispositions and risks. Had the new recommendation been put into effect, physicians would have been granted the right not to know patients' preferences. The ACMG eventually retreated from the Working Group's recommendation in response to criticism within the research ethics community. As this complex negotiation over who should be accorded the privilege not to know illustrates, deliberate ignorance is frequently and intricately intertwined with power (see also Ellerbrock and Hertwig, this volume) as well as the delegation of responsibility. Knowledge is power—but so can be the right not to know.

To conclude, an individual's choice to ignore information rarely affects just their own well-being. Externalities are a powerful justification for third-party intervention, yet caution is warranted here. In societies that have a universal health care system, for example, even decisions about individual health will eventually impact everybody else, as everyone shoulders the costs for the health service. However, if such externalities, which only come into being in the first place due to institutional intervention, are seen as sufficient grounds to condemn others' information preferences and behaviors, the right to intervene will become pervasive. The more severely third parties are affected, the more consideration may be given to interventions that make it harder for an individual (not) to generate, retrieve, or use an influential piece of information.

Should Certain Preferences Be Ignored?

The concern about externalities is utilitarian. The policy makers' concern is that individuals might increase their personal well-being at the expense of inflicting disutility on others. This would be inefficient, as total welfare is smaller than it could be. Critically, from this normative perspective, the goal is to fulfil as many individual wishes as is feasible, given the resources the economy can muster. In the textbook version of the argument, preferences are defined narrowly as the willingness-to-pay for goods or services. The global optimum is reached if the wishes of those with the highest willingness-to-pay are fulfilled. Conceptually, the apparatus of welfare theory can also be applied if utility is not equated with profit. A subbranch of economics works on such extensions, and calls for social preferences (e.g., care for the well-being of others) to be integrated into the otherwise narrow willingness-to-pay preference functions. Bierbrauer (this volume) convincingly argues that applying welfare theory to such more broadly defined preferences can lead to repugnant outcomes. For instance, if a person cares about the material well-being of another, but not vice versa, both will be better off if the government takes money from the caring person and gives it to the other one. Thus, the most socially minded get the worst deal. This result can be avoided only if policy makers deliberately ignore the possibility that some members of society might hold social preferences and respond to motives other than self-interest when making decisions. Deliberate ignorance, therefore, has a place at the heart of normative utilitarian theory.

The concern about preferences that normative theory should ignore reaches even further. Happiness research has produced a body of evidence suggesting that a person's state of affective well-being is highly adaptive. Individuals adjust surprisingly quickly when circumstances deteriorate, often bouncing back to the same happiness level as before (Frederick and Loewenstein 1999). Hedonic adaptation is a good thing. It allows people to lead a meaningful life even under dire circumstances. Yet if policy makers were to strive simply to maintain the level of happiness, they would have carte blanche. As happiness normally reverts quickly to its original level, policy makers could ignore the harms that their interventions inflict on citizens and focus instead on enriching themselves or furthering the political goals of their clientele. To produce normatively acceptable decisions, welfare theory must deliberately ignore some preferences (Bierbrauer, this volume).

Should Deliberate Ignorance Be Assessed Solely in Terms of its Consequences?

Society gives academics freedom and finances the scientific enterprise because it embraces the ideals of enlightenment. But is uncovering the secrets of life invariably good under all circumstances? There have always been conflicting normative claims. Creationists argue that it is a sin to investigate the Darwinian origins of life, contending that only the Bible holds such answers. The current debate over using CRISPR technology to edit the human genome centers on human dignity as a limitation for scientific investigation (see, e.g., Brokowski and Adli 2019). These concerns are deontological.

Deontological and consequentialist theories are disconnected as a matter of principle. Deontological theories argue from first normative principles, such as Kant's categorical imperative: "Act only according to that maxim whereby you can at the same time will that it should become a universal law" (Kant 1785/1993, 4:421). By contrast, consequentialist theories hold that only the consequences of one's acts are the basis for judging their moral rightness. Utilitarian theories are a subgroup of consequentialist theories in that they define welfare to be the ultimate goal. Scarce resources are to be used in the most productive way.

One important bridge has, however, been built in the debate between deontology and utilitarianism. On deontological grounds, it can be argued that rules should be followed at all times ("rules are rules"): it would be immoral to break a rule that has been legitimately established. Rule utilitarians argue that the same norm can also be established on utilitarian grounds if a multi-period framework is adopted: over time, everybody is best off if legitimate rules are followed. This creates order and saves the inefficient transaction cost of sanctioning rule violations (Hooker 2016). By analogy, one might ask whether there is such a thing as information utilitarianism or, more broadly, information consequentialism. To the extent that society is better off, at least in the long

run, if information is generated, retrieved, and used, utilitarian or consequentialist theorists could require that information always be generated, retrieved, and used, irrespective of the immediate benefit. Of course, there could also be rule consequentialist grounds to *not* generate, retrieve, and use information. Blind auditions, for example, may fall under this category.

Institutional Implications: How to Avert or Foster Deliberate Ignorance

From a normative perspective, preventing undesirable deliberate ignorance is no less relevant than is enabling desirable deliberate ignorance. The central goal of institutional design, however, is to prevent socially undesirable behavior. Most (formal and informal) institutions have been set up to combat pervasive undesirable behaviors rather than to facilitate desirable ones. From command-and-control regulation to disincentives, from moral suasion to nudges, there is a whole panoply of tried-and-tested tools for discouraging individuals from engaging in undesirable behaviors. Combatting unwanted deliberate ignorance is not in any principled way different from standard normative concerns, such as fighting pollution or speeding.

By contrast, the enabling function of institutional intervention is well understood only for the core of a market economy. Specifically, within the market economy, property rights define and standardize the object of trade, and contracts make trade possible. Could these standard techniques also enable socially desirable deliberate ignorance? This is not obvious. If a *Homo ignorans* desires to remain ignorant of something, it is crucial that nobody relays relevant information to him. If more than one person becomes aware of the information of which a person desires to remain ignorant, this preference will only be protected as long as all who do or might know refrain from informing him. *Homo ignorans*'s preference thus creates a one-to-many relationship. This makes it difficult for *Homo ignorans* to turn his interest not to know into an object of trade. He would have to strike a deal with all potential informants.

Formal versus Informal Institutions

Legal institutions are formal, in the sense that they are explicitly designed and, if needed, explicitly enforced. For some instances of deliberate ignorance, this formality may be desirable. If policy makers want to prevent a firm from remaining deliberately ignorant about the harmful effects of a production technology on the environment, they may want to force the firm to clear the procedure with an environmental agency before it starts producing. If the agent of deliberate ignorance is an individual consumer, however, legal intervention may at most be an institutional backstop. Unless the legal rule is mirrored by a sufficiently powerful social norm, protection is likely to be imperfect.

The Role of Education

An educational approach is particularly appealing when the goal is to enable behavior. The normatively desirable reaction will rarely consist of never or always generating, retrieving, or using information. Rather, individuals should ideally be empowered to discriminate between contexts and issues in which having more information is better and those in which it is better to refrain from accessing information. In Hertwig and Engel (2016), we gave the example that (social) media and Internet platforms have become experts in designing mental stimulants that usurp users' attention. In an informationally obesogenic environment, citizens are threatened with loss of agency over how much of their attention they allocate, and to what. One of the most important goals for future school and adult education may be to equip students with the competence to discern good from worthless information, and to detect and reject the relentless attempts to hijack their limited attentional resources.

The Role of Digital Technology

Digital technology may also offer an intriguing solution precisely because it is embedded in (computer) code. As Lessig (2009) noted, code is law. Social norms and legal rules are never perfectly enforced. There is always an implementation gap resulting from neglect, resistance, or a lack of enforcement. By contrast, computer code is self-enforcing. If a piece of information is not to be accessible (and provided that it has not yet been duplicated), a single line of code can make it disappear. Likewise, if there is concern that people might avoid a piece of information they ought to see, a few lines of code cannot only make sure that they receive the information but also document exactly when they received it.

When deliberate ignorance is implemented by code, normative conflicts that can otherwise be kept hidden become manifest. For example, antidiscrimination law prohibits discrimination on grounds of race or gender, but as long as decision makers do not openly justify their choices based on either category, it is difficult to prove that they have engaged in discrimination. An electronic decision tool can be programmed to purge a data set from informative correlations with gender or race. Although highly effective in preventing discrimination, this intervention may also reduce prediction accuracy; the more it does, the more the normatively undesirable behavior is actually correlated with gender or race. Intervention by code thus forces an open discussion of this trade-off. Society must decide how high a price it is willing to pay not to discriminate.

Deliberate Ignorance: A Wisdom Call

According to Kant (1784), "Enlightenment is man's emergence from his self-imposed immaturity." A self-determined life, grounded in knowledge

and understanding, is certainly desirable. Quite often, making the best use of the available knowledge enables people to live a meaningful life. Yet, as the contributions to this volume demonstrate, more knowledge and information are not always desirable, and deliberate ignorance cannot simply be equated with self-imposed immaturity. Individuals and societies may have good reason not to generate, acquire, access, disseminate, or use knowledge and information, even if doing so would be feasible and affordable. There are contexts and conditions under which it is better to remain deliberately ignorant. Striking a balance between the liberating and enlightening effects of knowledge and the beneficial effects of self-imposed ignorance requires individual, collective, and institutional wisdom.