



From “Deliberate Ignorance: Choosing Not to Know,” edited by Ralph Hertwig and Christoph Engel.
Strüngmann Forum Reports, vol. 29, Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262045599

5

The Deep Structure of Deliberate Ignorance

Mapping the Terrain

Barry Schwartz, Peter J. Richerson, Benjamin E. Berkman,
Jens Frankenreiter, David Haggmann, Derek M. Isaacowitz,
Thorsten Pachur, Lael J. Schooler, and Peter Wehling

Abstract

This chapter explores the “deep structure” of deliberate ignorance, defined as an individual’s or collective’s intentional choice to create a short- or long-term barrier to information for the individual or collective who made the choice. This definition is used to identify clear cases while acknowledging that the key terms of the definition (deliberate and ignorance) admit of ambiguity. It is argued that the frequency, forms, and functions of deliberate ignorance may vary across individuals as well as domains of information. Potential causal variables are suggested (e.g., the utility of the information, the nature of the information environment, the level of relevant parties who initiate and are affected by deliberate ignorance, and the legal, ethical, and social context within which deliberate ignorance occurs) and possible consequences are explored for the actors who engage in deliberate ignorance. Finally, the potential time course of deliberate ignorance is discussed within an episode of deliberate ignorance itself, across life-span development as well as cultural and biological evolutionary time.

Introduction

“It is always better to have more information than less.” “The more information you have, the better the decision you will make.” Statements like these

Group photos (top left to bottom right) Barry Schwartz, Pete Richerson, David Haggmann, Derek Isaacowitz, Peter Wehling, Benjamin Berkman, Jens Frankenreiter, Lael Schooler, David Haggmann, Thorsten Pachur, Benjamin Berkman, Barry Schwartz, Lael Schooler, Pete Richerson, Derek Isaacowitz, Jens Frankenreiter, Thorsten Pachur and Lael Schooler, Barry Schwartz, Peter Wehling, Jens Frankenreiter, Pete Richerson

seem uncontroversial, especially if information is easy to come by and potential benefits are high. It is the obviousness of the value of information that led Hertwig and Engel to describe a set of circumstances in which people deliberately avoid information that is easy to come by (see Hertwig and Engel, this volume, 2016). They call this set of phenomena “deliberate ignorance.” For discussion of a related set of phenomena, which they refer to as “information avoidance,” see Golman et al. (2017); for a case study of “intentional ignorance” in the medical domain, see Owens (2017).

Hertwig and Engel define deliberate ignorance as “a conscious choice not to seek information or knowledge, especially where acquisition costs are small and potential benefits are large.” Consider the following cases:

- A person takes an HIV test, receives an envelope with the results, but does not open the envelope.
- A person receives a quarterly statement from a retirement investment fund, opens the envelope but puts the document, unread, in a file folder.
- A scholarly journal requires that all submitted manuscripts must be free of any information that identifies the authors.
- A personnel department scrubs gender, race, and ethnicity from its job applications.

These are all examples in which information possibly, or even probably, has high signal value, yet it is avoided. What might explain this desire for less information?

Building on Hertwig and Engel’s attempts to characterize deliberate ignorance, describe various contexts in which it occurs, and delineate various functions it might serve, our goals here are as follows:

- To define and delimit deliberate ignorance, distinguish it from other (perhaps related) phenomena, and begin to characterize mechanisms and cognitive functions that implement deliberate ignorance. Our aim is to identify the clear cases of deliberate ignorance and leave more ambiguous cases for future inquiry.
- To explore the extent to which the phenomenon of deliberate ignorance, the functions it serves, and the factors that affect it are common across different actors as well as different domains of experience (e.g., personal medical information, personal financial information), and the extent to which there may be domain specificity to deliberate ignorance.
- To discuss some of the psychological and cultural mechanisms that may be involved in the phenomenon of deliberate ignorance.
- To identify potential key variables, both with respect to causal factors that influence and consequences that follow deliberate ignorance. We will emphasize consequences of deliberate ignorance that are, at present, least well studied and least well theorized.

- To examine deliberate ignorance as it may operate at different time scales, both developmental and evolutionary.

Hertwig and Engel identified six different types of functions that deliberate ignorance might serve:

1. To heighten suspense and surprise (e.g., when we don't want to know the ending of a thriller when we sit down to watch it, or the sex of an unborn child).
2. To provide a strategic advantage in certain competitive bargaining situations, an idea that originated with the counterintuitive observations of Schelling (1956; see also Dana 2006).
3. To enable people to manage and sustain their limited cognitive resources (e.g., Crawford 2015).
4. To maintain impartiality and fairness (see MacCoun, this volume).
5. To enhance performance (e.g., Kluger and DeNisi 1996).
6. To regulate emotions, such as avoidance of worry or regret (e.g., Howell and Shepperd 2013; Karlsson et al. 2009; Yaniv et al. 2004).

In this chapter, we focus on a subset of these functions. We do not discuss surprise and suspense maintenance because we think in cases like this, it is not the information itself, but the anticipation of having it revealed in the future (temporary information avoidance) that seems crucial. We also do not discuss cognitive resource management because this is already reasonably well-studied and well-theorized (see Sims 2003; for a paper that launched the study of the “economics of information,” see Stigler 1961), and we do not discuss strategic ignorance, because it has been a topic of research for more than half a century in experimental game theory and other contexts. However, there are two qualifications to this last exclusion. First, the focus of research on strategic ignorance has largely been on interpersonal settings, like experimental games or negotiations. Less studied and theorized about is what might be called *intra-personal* strategic situations, in which, for example, ignoring some information might make it easier for someone to fulfill a long-term goal or execute a plan (e.g., Carrillo and Mariotti 2000; Wooley and Risen 2018). Second, in Schelling's (1956) groundbreaking work on bargaining and negotiation, many of his examples involved strategic *binding* (i.e., deliberately limiting one's options) rather than strategic *blinding* (i.e., deliberately limiting one's information). The latter is a clear case of deliberate ignorance whereas the former is not.

What Counts as Deliberate Ignorance

In generating a definition of deliberate ignorance that helps us identify clear-cut examples, we do not claim that cases falling outside of this narrow definition should be excluded from future analyses. We merely wish to start with the

cleanest cases to help us understand what is special about the concept. Each of the key terms in the definition, *deliberate* and *ignorance*, are associated with complexities and uncertainties.

To begin, let us consider the definition of deliberate ignorance provided by Hertwig and Engel (p. 5, this volume):

[T]he conscious individual or collective choice not to seek or use information (or knowledge; we use the terms interchangeably). 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 (Stigler 1961)—acquiring information would seem to be rational (Martinelli 2006).

Expanding on this definition (our additions are italicized),

an individual's or collective's *intentional choice to create a short- or long-term barrier to information for the individual or collective who made the choice,*

we will explain how our definition differs from that by Hertwig and Engel.

First, we consider all decisions that create barriers to information as potential candidates for deliberate ignorance. This definition covers most decisions not to seek or use information as described by Hertwig and Engel, but it covers other situations as well. Most importantly, we consider situations in which an actor adopts measures that make it harder to access certain information in the future as falling under the definition of deliberate ignorance. For example, an investor might opt not to receive quarterly portfolio reports online. A patient might opt not to be sent a report of blood work results unless there is a problem.

One particularly important dimension where our definition is potentially broader than that adopted by Hertwig and Engel lies in its treatment of decisions to make previously accessible information inaccessible in the future. Hertwig and Engel's definition seems to exclude such acts from the definition of deliberate ignorance. Our definition includes them. Our reason for this decision is that we think that the factors that motivate actors not to access information might also motivate them to destroy it. For instance, if not knowing certain information conveys a strategic advantage, it does not matter whether an actor does not access it in the first place or manages to effectively "forget" it (see Schooler, this volume). Such decisions could also raise questions that require special attention.

We exclude from the definition those decisions that affect exclusively the ability of others to access information. It is possible to think of many cases in which it is in an actor's interest to withhold information from others or even actively deceive them. If successful, such actions might lead to ignorance in others. From the perspective of the targets of such actions, however, this ignorance cannot be considered "deliberate." By contrast, were an actor to employ other agents in an attempt to erect barriers to information, we would consider this an act of deliberate ignorance by the actor.

Perhaps the most important element in our definition of deliberate ignorance is the adoption of an intentionality requirement. We define intentionality as comprising a voluntary element in addition to knowledge about the potential consequences of an action. In the context of our definition, intentional choices refer exclusively to choices the goals of which include the creation of barriers to information. Most importantly, this requirement is not fulfilled if the creation of barriers to information is a side effect of some other action. For example, a decision to involve an agent will almost invariably create some barriers to information for the principal. However, these barriers to information are likely not the primary goal of the choice to involve an agent. Similarly, we understand intentionality to exclude most cases in which information is not accessed for reasons related to the cost of accessing and processing this information. This also implies that those cases on which Hertwig and Engel are not particularly focused (cases in which the marginal acquisition costs are nonnegligible or the potential benefits are small) are generally excluded from our definition of deliberate ignorance. Perhaps more accurately, while not excluded from the definition, they are not a focus of interest, being neither surprising nor puzzling. Finally, we use the word “intentional” rather than “conscious” (as in the original definition) to highlight that the goal of a decision is to erect barriers to information. One may be quite conscious of the decision to file away a quarterly investment report without reading it and yet do it without the intent of avoiding information. People may file away unread reports because they are too busy to attend to them, or think they don’t know enough to act on the information or that they will study the report on the weekend when they have more time. In none of these cases is it the person’s intention to withhold the information from scrutiny. The intentionality requirement clearly narrows the scope of instances of deliberate ignorance, but we think it puts the focus on the most interesting and puzzling phenomena.

Complexities of “Deliberate”

“Deliberate” implies intentionality and usually consciousness. That said, we believe that there can be something of a continuum of deliberateness. For example, consider the possibility that what starts out as deliberate becomes habitual. After some reflection, a person decides not to read their first quarterly investment statement. The same decision is made quicker the next time, and the next, and so on, until stuffing it in a folder becomes a mindless habit. In this sort of case, the ignorance is no longer deliberate, though it was at first. Similarly, imagine a professor who decides, as a matter of policy, not to consult student grades on previous assignments in grading the current one. This is a case of deliberate ignorance. Subsequently, the professor simply enacts the policy, without thinking. Just as buckling a seat belt can go from deliberate to automatic, so can decisions to ignore information.

As another type of example, consider deliberate ignorance in certain social situations. A person decides not to ask a friend how her troubled marriage is going. This may be quite deliberate—an adherence to norms of politeness and personal privacy. This example is not discontinuous with many examples in which people do not ask questions so as not to pry or be impolite without engaging in much reflection about the matter. One typically adheres to norms of politeness without giving the matter much thought.

The social norm of “respect for privacy” may be quite powerful. We avert our gaze from a friend’s open browser window. We avoid looking at our romantic partner’s emails and texts. We may do so to maintain social ties or to avoid shame and embarrassment. As Elias (1978) argues in his classic work, *The Civilizing Process: The History of Manners*, the history of Western society from the Middle Ages to the nineteenth century represented a gradual transformation in people’s ideas concerning manners and bodily propriety. Central to this transformation were decisive changes in feelings of shame, repugnance, and embarrassment that attended a wide range of human bodily functions such as eating, spitting, blowing one’s nose, urinating, or defecating. These changes in manners and associated feelings may change what we deliberately do not want to know. We do not want to know, for instance, what other people do in their bathrooms or bedrooms, or things that trigger feelings of shame on someone else’s behalf. Moreover, as social norms change, the tendency to pursue deliberate ignorance may change. At the same time that privacy regarding bodily functions is enhanced, the explosion of social media may already have transformed people’s notions of privacy regarding other information about their daily lives.

Related to politeness, relations of trust are partly defined by one’s unwillingness to check up on someone to make sure something has been done. This may be deliberate, in the sense that merely asking the question would be a violation of trust. Alternatively, it may be quite automatic, wrapped up in the very notion of trust. A slogan often heard in the domain of foreign relations—“trust, but verify”—is quite literally a contradiction.¹

As with most social processes, behavior dictated by concerns about politeness or trust can change as social norms change. Consider the adage that “it takes a village to raise a child.” This implies that one’s fellow villagers are authorized to step into the parental role if action is needed. Whereas this intrusiveness may be socially acceptable in some places, and may have been socially acceptable at some times in other places, presently, many people in many places would regard adopting a parental role as a member of the “village” as a deep violation. Cuddihy (1974, 1978) has written extensively about how the

¹ Trust, of course, is a very broad phenomenon. Trust does not necessarily equate to an unwillingness to check up on someone. Controlling other people is often costly from an economic point of view; that is, it requires time, effort, and potentially sophistication that may not be available to the principal. Thus, trust may be cost minimizing rather than norm preserving.

boundaries between public and private, or between self and other, have differed historically among religious groups in the United States.

From the perspective of our definition of deliberate ignorance, socially embedded practices like politeness and trust pose a problem in that they may not (always) be intentional. Nonetheless, we think they should be included in the definition because the range of phenomena they encompass may be vast.²

Complexities of “Ignorance”

The central distinction to be made here is between ignorance and error. Ignorance is not knowing. Error is false belief: “knowing” what isn’t true (see Lewandowski, this volume). Being wrong is not the same as being ignorant. It is unremarkable that people make errors. More remarkable is that people are ignorant, by choice. In this connection consider the pervasive phenomenon of confirmation bias (e.g., Wason and Johnson-Laird 1972). The “bias” is that people overvalue evidence that can confirm a hypothesis and undervalue (or do not value at all) evidence that can falsify a hypothesis. By not valuing falsifying evidence, are people being deliberately ignorant? We think the right answer to this question is that it depends. *Unmotivated* confirmation bias reflects error; people erroneously evaluate the evidentiary usefulness of various pieces of information. They overvalue some information and undervalue other information. *Motivated* confirmation bias, by contrast, may be ignorance, not error, in that under these circumstances, it is at least possible that people appreciate evidentiary usefulness of information but choose to ignore it nonetheless (see Dawson et al. 2002). We hasten to add in this context that although the word “bias” in “confirmation bias” implies that information seeking is nonoptimal, there may be some circumstances in which confirmation bias serves people well (see, e.g., Hahn and Harris 2014; Oaksford and Chater 2001). Imagine, for example, a researcher embarking on a new line of investigation of a phenomenon that is difficult to produce in the laboratory. A selective focus on successes may keep the researcher engaged in efforts to produce the phenomenon more reliably and robustly. Attention to disconfirmation might nip the research enterprise in the bud.

In summary, we think that deliberate ignorance is best thought of as a kind of “natural category” (Rosch 1973; Wittgenstein 1953). There are clear, prototypical examples: Should the HIV test result envelope be opened or the quarterly investment report read? There are also other examples whose membership in the category is graded, in large part, by how deliberate they are and how much they reflect ignorance rather than error. Natural categories possess no “necessary and sufficient features” in the way that artificial, scientific

² It should be noted that choosing not to ask a question may be due to a fear of the social consequences of asking. Effectively, the actor is choosing ignorance, but not because of a lack of wanting to know.

categories do (e.g., a geometric shape either is or is not a triangle). These categories possess instances that resemble each other just as members of a family resemble each other. Very good examples of a category are prototypical. Other examples become increasingly less good as they are less like the prototype. Wittgenstein famously identified “tools” and “games” as examples of natural categories. A prototypical tool might have been a hammer or a screwdriver when he made his observations. These days, it might be a computer, a cell phone, or an app. We think it is a wise research strategy to focus studies of deliberate ignorance on prototypical examples and then extend research outward to less-clear examples as the concept of deliberate ignorance becomes better understood.

Domain Specificity of Deliberate Ignorance

As we have defined it, deliberate ignorance can be found across widely different domains of information, yet it is still possible for its incidence, character, causes, and consequences to be domain specific. People may be more likely to pass up information in domains where expertise has been acknowledged than in domains that are more matters of personal preference. For instance, people might defer to financial advisors or doctors but pore over available information about restaurant options in a city they will be visiting. In addition, the degree to which individuals can understand information, or think they can understand information, may vary across domains. A geneticist, for example, is likely to have substantial background knowledge relevant to health but may not have any knowledge related to finances. A given piece of information (e.g., the results of a genetic test or the composition of an index fund) may therefore have different value to the geneticist than to an economist. Consequently, interest in receiving information may differ across the domains: the geneticist may be more willing to trust an expert and remain ignorant about finances, while taking an active interest in the detailed results of the genetic test. The general point here is that instances of deliberate ignorance may be a function of the difference between how much one already knows and how much one needs to know for the information to be effective. Since individuals’ depth of knowledge will vary across domains, as will the informational complexity of domains, domain specificity seems to be highly likely. Interestingly, the geneticist may still exhibit greater *deliberate* ignorance in the health domain than the financial domain: deliberate ignorance requires awareness that the information exists in the first place, thus the geneticist has more opportunities to be “deliberately” ignorant in the health domain.

Information in some domains may, moreover, be inherently more informative independent of relevant knowledge. A genetic test for Huntington disease, for example, is a near perfect diagnostic of the underlying

condition, leaving no room for doubt once results are obtained. Other information, particularly in cases that involve others' judgments or evaluations, can be subject to substantial noise. Learning that a colleague is (stubbornly) unpersuaded by a new theory does not resolve whether the theory is indeed useful.

Moreover, information in some domains may inherently be more informative in guiding decisions. Someone who learns of a treatable or curable medical condition, for instance, can take concrete actions to improve future outcomes. Opportunities to respond to information that one is unattractive, on the other hand, may be limited. While holding accurate information may still confer some advantages (e.g., better calibrating expectations on the dating market), these benefits are more nebulous. People also may have (potentially motivated or biased) beliefs about how actionable or informative information is. For instance, people's beliefs about how painful it is to talk to those with opposing political views may be exaggerated (Dorison et al. 2019).

The desire to seek or avoid information may further differ across domains as a result of prevailing social norms. If friends are likely to discuss current political events, choosing not to read the latest news may impose a social cost. Not only may it preclude participation in discussion, but you may be judged adversely for failing to adhere to a "duty" to be informed. In other cases, the decision *not* to remain ignorant may violate social norms, a point we made above regarding privacy and trust. When we see an open browser window with our significant other's emails displayed, succumbing to the temptation to read the emails (and hence acquire potentially new and useful information) is likely to be judged unfavorably.

These examples raise the question of whether the motivations and strategies for deliberate ignorance differ fundamentally across domains or whether there are indeed commonalities reflective of "information preferences" that hold across domains. Previous work in psychology, economics, and other disciplines has found substantial avoidance across consequential domains: information about health (Oster et al. 2013) or finances (Sicherman et al. 2016), among others.

There are several possible domain-specific influences on deliberate ignorance. Information in some domains may have instrumental value, in others it may have hedonic value, and in still others it may have both. Deliberate ignorance may be a function of what kind of value the ignored information may have. Knowledge of one's portfolio can have both instrumental and hedonic value; knowing you were well thought of by one of your college professors some years ago may have only hedonic value; knowing which route from Berkeley to Palo Alto has less traffic may have only instrumental value.

In addition, there may be domain specificity in what one is *expected* to know. When someone says that some issue is "above my pay grade," domain specificity regarding "who is in charge or who is the expert" may partly be at play.

Domain specificity may also vary across time and culture. Consider attitudes about whether people should trust or rely on experts.³ “Trust the expert” as guidance may change as a culture’s attitude about “expertise” changes. “Rely on the expert” (Wegwarth and Gigerenzer 2013) may change as a culture’s attitude about who bears responsibility for outcomes changes (for the distinction between “trusting” as a social practice of deliberate ignorance and “relying on” as a form of imposed and seemingly unavoidable ignorance, see Townley 2011). For instance, healthcare in the United States has become much more consumer driven over the last fifty years. Prescription drugs are advertised to patients, who, of course, cannot go out and just purchase them. Doctors are admonished not to be paternalistic and to make certain that patients realize that the decision rests with them. This change in cultural attitude toward the role of expertise in decision making may have made patients less “deliberately ignorant” in 2019 than they were in 1969.

Finally, domain specificity may play a role in whether an individual’s attitude toward information is “I don’t need to know,” “I can’t know (it’s too complicated),” or “I don’t want to know.” Which of these responses a particular piece of information provokes may affect the frequency of deliberate ignorance as well as its consequences. What is called the “illusion of explanatory depth” might be relevant to choosing not to know (e.g., Keil 2006; Sloman and Fernbach 2017). The illusion of explanatory depth refers to the fact that most people say yes when asked if they know how a toilet (a zipper or a bicycle) works. Yet when asked to explain how it works, most people’s knowledge is very shallow. Discovering this fact about their ignorance encourages people to change their view about how well they understand. If some deliberate ignorance stems from an “I don’t need to know” attitude, and some originates from an “I already know enough about that” attitude, then explanatory depth-type manipulations may reduce deliberate ignorance.

The importance of domain specificity is, of course, unknown at this time. Also unknown are the dimensions along which we can most usefully characterize domains (e.g., by function, by complexity, or something else). Asking “Who is the expert here?” might help parse domains usefully. Other possible useful distinctions among domains are whether the information is actionable or not, and whether it is hedonically charged or not.

Despite well-documented information-avoidance behavior, some people routinely get tested for sexually transmitted diseases, frequently (and maybe excessively) check the value of their portfolios, or expose themselves to political views contrary to their own. This suggests that information preferences may be an important source of individual differences, similar to time and risk preferences. In addition, whether the potential benefit of information is small

³ The distinction between “trust” and “rely on” is meant to capture the possibility that we may rely on others in some domains whether or not we trust them and their expertise.

or large may be answered quite differently by different social actors. Consider the example of the genetic test for Huntington disease: For some individuals “at risk,” there might be a huge benefit in definitively knowing whether or not they carry the respective genetic variant. For others, this information might be a source of potential harm in that it will destroy uncertainty as an indispensable resource for leading a self-determined life. There are, in other words, no unequivocal and easily generalizable criteria for judging whether (and for whom) the benefit of knowing something is small or large.

We may learn about individual differences and domain specificity by measuring people’s willingness to remain ignorant in different domains. Is the decision to learn about one’s health predictive of the decision to learn about one’s attractiveness? If so, this suggests there is an underlying common factor. Relatedly, we may wonder whether findings on deliberate ignorance are driven by a small subset of the population that wishes to remain ignorant about information very broadly, or whether a large fraction (or even a majority) of people deliberately choose to remain ignorant in at least some settings.

One measure of information preferences is a scale developed by Ho et al. (2018): respondents are asked to imagine themselves in a series of hypothetical scenarios in which they can choose to obtain (or not obtain) information. The scenarios cover three domains that span many high-stakes decisions, and for which there exists empirical evidence of avoidance: personal health (e.g., the choice to obtain information about life expectancy); personal finance (e.g., the choice to learn about alternative investments that could have been pursued); and personal characteristics (e.g., one’s attractiveness). Ho et al. (2018) find that items from each of the domains load onto domain-specific latent factors, and these latent variables load onto a general factor. Moreover, the general factor is predictive of consequential information acquisition in the three domains. Suggestive of general information preferences, the scale is also able to predict the decision to acquire information outside of these particular domains; namely, the decision to learn about the gender wage gap in one’s industry, the consequences of climate change to one’s local community, and unfavorable information about one’s political party. In addition, there are small to moderate correlations between information preference and dispositional factors, with some variation across domains. These findings suggest that domain specificity may be superimposed on some person-level preferences related to domain-general deliberate ignorance.

Psychological Mechanisms That Implement Deliberate Ignorance

Since deliberate ignorance can manifest itself in various ways and circumstances, there are also likely to be somewhat different types of mechanisms involved. On a general level, one important distinction is between (a) mechanisms for situations in which the agent knows about the existence of the

relevant information but has not yet acquired it, so that the information is not yet internally represented (e.g., the HIV test result), and (b) mechanisms for situations in which the information has already previously been acquired and is internally represented (e.g., the example of the flute episode discussed by Schooler, this volume). Below, we discuss these two types of mechanisms and their cognitive requirements in turn.

In situations in which the information is not internally represented, ignoring it means not including it in the search rule one uses during information search, although it may very well be considered relevant in principle for the task at hand. Implementing deliberate ignorance in such a situation involves, in particular, executive functions of planning and selective attention. In addition, it is important that when the environment is searched and the to-be-ignored information is encountered by chance, it has to be recognized as such, which requires constant monitoring and matching with the current task goal. Importantly, relative to purely exploratory information search, targeted information search (i.e., search that focuses on some specific kinds of information and explicitly excludes others) has been shown to entail cognitive costs. For instance, Fechner et al. (2018) developed a model that implemented various decision strategies that differed in their search rule within a given cognitive architecture. One decision strategy gathered all relevant attributes it could find in the environment and its search rule did not mandate a particular order in which attributes had to be inspected. The other strategy processed attributes in a particular sequence and stopped information search as soon as a given attribute allowed it to discriminate between the options; all other attributes were ignored. Because the strategies were implemented in a common cognitive architecture, Fechner et al. could determine the cognitive costs (in terms of predicted response time) that the strategies produced for processes of information retrieval, action coordination, perception, and motor responses. It turned out that although the second strategy often ignored some of the attributes available in the environment, it produced higher cognitive costs than the strategy which considered all available information. These costs were produced, in particular, by the strategy's search rule, which mandated a focus on specific attributes and the exclusion of others. These results were subsequently confirmed in an empirical study.

Situations in which the to-be-ignored information is already represented in the cognitive system involve further complexities. Due to the architecture of the cognitive system, if a piece of information is generally relevant for a task at hand, its activation in memory will be enhanced. Ignoring this information, therefore, requires an active downregulation of the cognitive system, for instance, by processes of suppression and inhibitory control. These processes involve considerable mental effort (with pronounced individual differences). In addition, processes of inhibition decline in older age (Hasher and Zacks 1988; but see Rey-Mermet and Gade 2018).

In summary, deliberate ignorance will often make specific demands on the cognitive system. Based on the requirement profile, one can make predictions regarding individual differences in engagement in deliberate ignorance, how it might change across the life span, and which situational variables (e.g., time pressure, dual task load) will modulate engagement in deliberate ignorance.

Causal and Consequential Variables

In the consideration of possible variables that might influence deliberate ignorance, we divide our discussion into two major parts. First, we identify variables that may affect the likelihood of deliberate ignorance—potential causal variables. Then we identify possible effects of deliberate ignorance. This discussion is intended to be the first, not the last, word on how we might subject deliberate ignorance to further empirical investigation.

Potential Causal Variables

We identified four categories of causal factors: utility of the information; information environment; level of relevant parties; and legal, ethical, and social considerations. Virtually all of these variables have implications for normative assessment of deliberate ignorance.⁴

Utility of the Information

As indicated in our discussion of possible domain specificity, the type of information (e.g., health, financial, personal characteristics, political) that is potentially being ignored might be salient. We might make different normative or policy judgments about cases depending on this variable. Across domains, however, we suspect that judgments about the significance and appropriateness of deliberate ignorance will focus on the extent to which the ignored information is perceived to have net utility. This judgment will likely have both causal and normative impacts.

There are a number of factors relevant to perceived net utility (see Table 5.1). Across these factors, the core task is to assess the possible *benefits and harms of having the information*. While it might be difficult to list comprehensively and measure precisely all the possible effects a piece of information could have, the goal is to try to determine at least whether or not knowledge of the information would create a net benefit or harm. Relevant informational

⁴ It should be noted that in the context of the Enlightenment-based attitude (where more information is always better than less), instances of deliberate ignorance tend to carry almost automatic negative normative judgments. That is, one must defend or justify decisions to remain ignorant, whereas decisions to acquire knowledge need no justification.

Table 5.1 Variables that may affect the utility of information.

Variables	Attributes
Benefit/risk of knowing	<ul style="list-style-type: none"> • Magnitude • Duration • Marginal value compared to baseline
Uncertainty	<ul style="list-style-type: none"> • Quality of evidence • Applicability to an individual • Decisiveness of information
Timing of relevance	<ul style="list-style-type: none"> • Temporal gap between risk and benefit manifestation • Life stage
Actionability	<ul style="list-style-type: none"> • Magnitude • Direct medical or preventative action • Indirect action • Small marginal impact
Accessibility	<ul style="list-style-type: none"> • Understandability • Held by experts

characteristics might include the magnitude and duration of the risk or benefit. It is also relevant to consider the marginal value of the information, particularly as it compares to the baseline of already acquired information. Information that makes someone newly aware of an issue will be more valuable than information that merely adds detail to an already established area of knowledge. The magnitude and direction of the information's utility will be extremely salient and will often serve as an initial threshold question when analyzing an instance of deliberate ignorance. When information has higher potential net benefit it will be easier to make a normative claim that deliberate ignorance is inappropriate; information with lower or negative value will not generally be associated with disapproval.

A related consideration is the *uncertainty associated with a piece of information*. We use "uncertainty" in a number of different ways:

1. Uncertainty related to the quality or amount of evidence available (which raises questions about the certainty with which conclusions can be drawn): Even if a patient definitely has a particular genetic variant, there might be weaker (e.g., single case report) or stronger (e.g., population-level data) evidence about the link between that genotype and a particular pathogenic phenotype.
2. Uncertainty about the applicability of information to a particular person: Well-characterized genetic variants are often still only partially penetrant; that is, only a subset of people with that variant actually manifest the disease. In any particular patient, there is uncertainty about whether or not having that variant will prove to be relevant to health.
3. Uncertainty about the decisiveness of a particular piece of information: Particularly when information is associated with future consequences,

it might be unclear how relevant the expected outcome might be for a person's future given the unpredictability of life. For example, knowing at age thirty that you are predisposed to develop cancer in your sixties would be irrelevant if you die of a different cause beforehand.

4. Uncertainty about the importance of the gap between how much one knows already and how much one needs to know: If a person already feels well informed, that person may justifiably ignore new information when it is offered.

As with utility, uncertainty will be relevant to the moral judgments we make about deliberate ignorance. Highly uncertain information will be less useful and will be more open to morally justified deliberate ignorance.

This relates to a third variable: *timing of relevance*. Certain kinds of information will be immediately relevant, while others might have delayed relevance. This will create situations where the harms and benefits of the information might be distant in time from each other. For instance, learning that you will likely develop a disease carries an immediate psychological cost, but the medical intervention might not be possible until some distant future time. Conversely, a company might not rigorously investigate a possible safety issue to protect short-term share prices, leading to a more distant risk of liability. When thinking about timing of relevance, it is also important to consider the life stage of the person involved (e.g., childhood, reproductive years, retirement age) because harms and benefits can shift over time.

A fourth variable, *actionability*, is particularly relevant when potential benefits are assessed. It is a key factor when making moral judgments about deliberate ignorance. We will usually only condemn someone for choosing to remain ignorant when they have foregone the opportunity to take an important action that would have been prompted by knowledge of the information. Actionability can be classified in four different ways:

1. Specific actionability, where there is an intervention one can take that is directly related to the piece of information. An example of this would be a medical intervention taken as a result of the news that one has a particular diagnosis.
2. Indirect actionability, where one does something in response to information, but not directly related to the possibility of altering the condition revealed by that information. This would include a decision to spend a significant portion of your retirement fund to travel the world after learning that you will soon become debilitated by an illness.
3. Actionability related to socially relevant information (where an individual's action can only have a miniscule impact). A clear example of this is climate change; one can take a direct action in response to information about climate change, but that individual action has little meaning on its own. In cases like this, strong norms may arise to stigmatize such ignorance. When dilemmas of cooperation exist, it may be

necessary for a large majority to act in the common good even if each individual's impact is negligible.

4. Actionability of personal interest (where the information is only relevant to satisfy a curiosity). There are potential things one can do with this information (e.g., tell a friend or relative about one's ancestry test results) but these actions have only limited consequences.

A fifth variable relates to the *accessibility of the information*. If information is not readily accessible, people may feel like it is not worth having. For instance, people might not read the important financial disclosure information that their mutual fund sends out each year because of a perception that they will not be able to understand it. Relatedly, if the information is being held by an expert (e.g., doctor, financial planner), there might be a tendency to think that it is not one's responsibility to acquire it. This again raises the distinction suggested above between trusting an expert and relying on an expert.

Information Environment

As our brief discussion of cognitive mechanisms involved in deliberate ignorance showed, both acquiring and ignoring information can be cognitively effortful and costly. Externalizing the information, with permanent storage and easy access, can lower some of these costs. Thus, the information medium (e.g., written, oral, digital), as well as the costs (metaphorical and financial) of storage and retrieval, may affect the likelihood of deliberate ignorance. High permanence, easy access, and low cost may convey that a deliberate ignorance decision is easily reversed and thus increase the likelihood of deliberate ignorance. This implies, of course, that instances of deliberate ignorance will only increase as the digital sourcing and storage of information increases, suggesting the apparent paradox that the more information is made available to people for easy access, the less inclined they will be to avail themselves of it.

Level of Relevant Parties

The range of actors or parties who can decide to remain ignorant or who are impacted by a decision to remain ignorant is great (see Figure 5.1). Individuals can decide to remain ignorant in a way that only (or predominantly) impacts themselves. For example, an individual can choose not to have a suspicious mole looked at by a doctor, or choose not to examine the performance of his retirement account. Individuals can also join with close third parties (e.g., relatives) to remain ignorant in a way that directly impacts themselves as well as the third party. A couple, for instance, might decide not to investigate the cause of their infertility problems, or putative siblings could decide not to seek genetic testing to see if they are actually related. An organization can decide

		Actor Who Decides to Be Ignorant			
		Only individual	Individual plus 3rd party (e.g., a family)	Organization	Society
Person Directly Kept Ignorant	Only individual	1. Skin mole 2. Investment			
	Individual plus 3rd party (e.g., a family)		1. Cause of infertility 2. Genetic relatedness		
	Organization			1. Facebook and Russian trolls 2. "Blind" auditions	
	Society				1. Stasi files 2. Gun injuries

Figure 5.1 Levels of action of deliberate ignorance. In this matrix, only the cells on the major diagonal are instances of deliberate ignorance, strictly defined. However, deliberate ignorance will often have effects on others (externalities) so that the other matrix cells are relevant, especially with regard to normative considerations. Items in the cells along the diagonal refer to examples in the text.

to remain ignorant in a way that impacts itself as well as individuals in that organization or unrelated third parties. For example, Facebook could opt not to rigorously investigate Russian election interference on its platform, or a symphony orchestra could adopt a policy to have blind auditions. Similarly, a society (i.e., its political representatives and dominant social groups) might decide to remain ignorant in a way that impacts itself as well as the individuals and organizations in that society as well as unrelated third parties. For example, the United States currently prohibits funding of public health research related to gun injuries, or some transitional societies choose not to render accessible existing intelligence on citizens collected by past regimes (see Ellerbrock and Hertwig, this volume).

It is important to note that our conception of *deliberate* means that the choice to not know something can only be made by the entity that will remain ignorant, as depicted in Figure 5.1. A choice to remain ignorant can certainly have a profound impact on other parties, but those other parties will not have engaged in deliberate ignorance.

Characteristics of the Actor

In some contexts, the characteristics of the actor (e.g., tolerance for ambiguity, need for closure, openness, neuroticism, and other personality variables) may also play a causal role in deliberate ignorance. So, also, do people’s beliefs

about the domain in question as well as people's hedonic needs and goals. Also relevant is whether an actor's decision to remain ignorant is an isolated instance or a repeated action. A single decision might simply be evidence of a temporary choice to remain ignorant, but serial decisions may provide evidence of stronger feelings. This variable could have interesting moral relevance, as we might more strongly condemn someone who repeatedly refuses to acquire important information (or we might more strongly praise someone who repeatedly refuses to acquire information with a net harm). In addition, repeated instances of deliberate ignorance may dramatically increase the magnitude of the consequences of that ignorance.

Legal, Ethical, and Cultural Considerations

As a final category, it is important to consider the legal, ethical, and cultural contexts of information. Certain kinds of information are socially valued or might have morally relevant social benefits such that one might be more expected to acquire the information. Social and cultural valuation of information and knowledge is, however, not always "rational" or "functional." It is not necessarily the most important and useful knowledge that is most appreciated. In contrast, certain types of knowledge (e.g., of esoteric arts or haute cuisine) are often culturally valued because they express a position of social superiority or exclusivity. Similarly, certain types of information are socially disvalued or have morally relevant social costs such that one might be discouraged to seek or accept it. These factors can impact the cost and benefit both of acquiring and of having information (for discussion on divergent social and cultural evaluations of knowledge and ignorance, see Gross and McGoey 2015; High et al. 2012).

Legal and ethical considerations can be a lever if we wish, as a matter of policy, to discourage (or, in rarer cases, encourage) deliberate ignorance. Relatedly, moral culpability for the consequences of actions flowing from having a piece of information can influence one's decision to know or not to know.

Consequences of Deliberate Ignorance

There is a wide variety of possible consequences that are worthy of study, on both scientific and policy/welfare grounds. We enumerate some potential consequences here.

In some cases, the instrumental consequences of ignorance will be positive overall, as in the examples of trusting or of blinding in peer reviews or job applications. In others, the instrumental consequences will be mostly negative, as in failing to take or get the results of medical tests that might reveal conditions that are treatable. In these types of cases, the potentially positive hedonic consequences of ignorance might be more important to the actor than the negative instrumental consequences. This suggests that there is utility, or benefit,

in holding some beliefs that new information might threaten (e.g., “I might not have cancer” or “I might be able to retire in material comfort”) quite apart from what one does as a result of those beliefs. Golman et al. (2017, 2019) has modeled such belief utilities. There may also be utility in self-consistency, leading to deliberate ignorance of information posing a threat to the image that one is consistent.⁵ Under a different facet of consistency, ignorance may help one maintain consistency with the values and attributes of one’s desired community. Community solidarity is often important to people, and discovering that you are less attuned to your group, or that your group is less attuned to you, can have a utility impact. Indeed, such information can damage social relationships.

We make this distinction between instrumental and hedonic consequences as if they are easily distinguished, but sometimes they are not. In many, if not most, cases, the consequences of deliberate ignorance may be both instrumental and hedonic. Consider the example of testing for Huntington disease: Not wanting to get tested does not simply have a “hedonic” value for people at risk (“I might not have the Huntington gene”). It is also “instrumental” in that it enables them to live a life with an (almost) open future. Likewise, knowing that one carries the Huntington gene has limited clinical instrumental value since, to date, there is no prevention, treatment, or cure for the condition.

In addition, we must pay attention to the possibility of externalities (i.e., of effects on third parties) and their normative implications. They may also have behavioral implications, modifying the likelihood that one will actually display deliberate ignorance.

Finally, it is possible that the same act of deliberate ignorance can have different consequences (even differently valenced consequences) at different points in time. We provide two illustrations.

First, if one deliberately ignores information about possible options in a consumer decision, one could miss out on an option that is better than the one actually chosen (negative). On the other hand, one might enjoy and value the chosen option more, owing to reduced regret and/or lowered expectations (positive) (see Schwartz 2004). Thus, one makes a worse choice but feels better about it. Schwartz (2004) distinguishes between two different sorts of goals that may inform decision making, by inducing two different information search-stopping rules. When maximizing, one seeks the best, which requires exhaustive search of the options. When satisficing, one seeks good enough, which usually does not require exhaustive search of the options. In deciding to satisfice, a person is being deliberately ignorant about information regarding options not considered. Schwartz (2004) suggests that maximizing leads

⁵ Note that this presumed process of belief utility maintenance does not include the well-studied phenomenon of biased updating of beliefs since in that case, the information is not ignored; instead, only some of it is actually assimilated into the person’s self-view.

to better objective decisions but worse subjective ones, a result confirmed in a study of hundreds of college seniors looking for jobs. Maximizers got better jobs, but felt worse about the jobs they got. Thus, deliberate ignorance of options is a cost at the time of choice, but a benefit when experiencing the results of the choice (Iyengar et al. 2006).

Second, if a person does not check his/her portfolio frequently, opportunities will be missed to improve it. Frequently checking the portfolio, however, leads to more inspections in which the portfolio's value may decrease, which owing to loss aversion (e.g., Kahneman and Tversky 1979, 1984) will make a person feel worse about the portfolio. Looking at the portfolio infrequently enables day-to-day market fluctuations to smooth out, so that the historical tendency of equities to rise in value dominates what a person sees on inspection. Thus, by ignoring portfolio performance, a person might earn less money but be happier about it (Benartzi and Thaler 1995), though under some circumstances, ignoring the portfolio can equate to *more* money being earned (Sicherman et al. 2016).

It is possible that there are actually many instances of this kind in which ignorance produces multiple heterogeneous effects.

Deliberate Ignorance and Temporal Scale

There are at least three types of temporal scales that can be considered when evaluating the deep structure of deliberate ignorance. The first is the temporal space of the decision to remain deliberately ignorant itself. Within that space, at least three phases can be distinguished, borrowing from the Rubicon model (for a summary, see Heckhausen 2007):

- the deliberation phase, in which the actor considers whether or not to remain deliberately ignorant
- the implementation phase, in which that decision is enacted
- the outcome phase, in which the consequences of the decision unfold

Distinguishing among these three phases is important not only for differentiating among causes, symptoms, and consequences of deliberate ignorance; it also carries implications for the presumed functions of deliberate ignorance itself. Many of the functions involve affect regulation, such as trying to avoid being upset, worried, or regretful. As such, the decision to remain deliberately ignorant for affect regulation purposes involves a prediction—an affective forecast—about how the actor will feel having or not having some information. Ample research suggests that individuals may not be accurate in their predictions, or affective forecasting, so predictions about how particular states of knowledge may influence affective states may also not be terribly accurate (for reviews, see Schwartz and Sommers 2013; Wilson and Gilbert 2003). Affective forecasting errors include failure to anticipate hedonic adaptation. In

the case of information, this would be a failure to appreciate that the hedonic impact of bad news will be reduced over time. Errors also include what is called “focalism” (Wilson and Gilbert 2003), the tendency to inflate the importance of bad news by focusing on the aspects of life that will be affected by it and ignoring the aspects of life (possibly many) that will be unaffected. Affective forecasting errors are especially relevant to the case of deliberate ignorance because deciding not to acquire some information out of concern that it will lead to negative affect may itself cause negative affect, even in the absence of the information itself (in the case of regret, see Schwartz 2004). The important implication of the possibility of misprediction is that the affective causes of a decision to be deliberately ignorant may not map perfectly onto the affective outcomes of being deliberately ignorant. This may be especially relevant in cases where the decision to remain deliberately ignorant must be repeated over time. In the example of maximizing versus satisficing discussed above (Schwartz 2004), a person might quite reasonably expect that a better objective decision will lead to a better subjective state and thus insist on very high standards for the decision. Such a person, adopting such a strategy, will often be disappointed—not once, but repeatedly.

Distinguishing among deliberation, implementation, and consequences as well as an acknowledgment that people often mispredict affective consequences raises what might be called the phenomenology of deliberate ignorance: What does it feel like to find out that relevant information is available? What does it feel like to choose not to examine that information? Having made that decision, what does it feel like to go forward without the information? Is a decision to ignore, once made, then forgotten? Or is it revisited again and again? For example, having decided not to review the quarterly statement of your retirement portfolio, do you cast the matter aside and forget about it, or do you continue to pay attention to the refusal to know? Do you impose recriminations on yourself later, when you discover that your portfolio has suffered a downturn? The phenomenology of deliberate ignorance is of interest in its own right, but it may also be of interest in determining whether ignorance actually achieves the affect regulation it is meant to provide.

A second, broader temporal scale involves the lifespan of the actor. There may be interesting developmental shifts involving deliberate ignorance. For example, children may be especially curious; evidence from Gigerenzer and Garcia-Retamero (2017) suggests older adults (in this case, age 51 yr and over) were more likely than younger adults to say they did not want to get information, such as when they would die. In addition, Hertwig, Woike, and Schupp (in preparation) observed that the strongest predictor of deliberate ignorance was chronological age (age 14 to > 80 yr): The older a person was, the more likely s/he exercised deliberate ignorance. Evidence of age differences in the components of deliberate ignorance, such as in affect regulation goals, might therefore lead to interesting age differences in the decision itself. To the extent that deliberate ignorance does vary by age, it raises interesting questions

about what mechanisms (e.g., learning, cognitive changes, motivational shifts) may contribute to the change. However, age differences in deliberate ignorance should not be presumed. There may be substantial age similarity in affect regulation processes (Livingstone and Isaacowitz 2019). That older individuals may not want to know when they will die may be taken as evidence that older age increases deliberate ignorance; however, a longer time window might reveal that the same individuals would have made the same decision on other health topics earlier in life as well.

A final, broader temporal scale involves evolution. Canonical examples of deliberate ignorance focus on individual decisions, but evolutionary considerations may be relevant for two reasons. First, biological and cultural evolution equip us with skills and dispositions that are relevant to individual examples of deliberate ignorance. For example, people often ignore otherwise costless information because knowing itself is aversive. There may be biological reasons for this. A simple, true item of information, if acquired, may set in motion a complex and effortful train of cognitive activity that is aversive. For example, learning that a spouse has been unfaithful may lead to a consideration of alternative actions in confronting (or not) the errant spouse, the spouse's reactions to the confrontation, and the reactions of friends, relatives, children, and legal authorities to the couple's joint actions. Some of these considerations will involve cultural norms. If a spousal transgression occurs in a culture of honor societies, and if the spouse knows, the knower is under obligation to beat or kill the spouse or the person with whom they have been unfaithful. On the other hand, among the Tupi speakers of South America, for example, women are quite free to have affairs with men who are not their social husbands, and any expression of jealousy will expose the husband to censure (Walker et al. 2012). In the former case, knowledge of infidelity will set into motion complex calculations; in the latter, perhaps, not so much.

Second, evolution may lead to adaptations or maladaptations that are analogous to deliberate ignorance, in the same way that Darwin coined the term "natural selection" by analogy with artificial selection. Animal breeders favored genetic variants that made livestock tame and unwary of humans. In effect, livestock were bred to ignore the exploitative intentions of their owners. Perhaps, Darwin thought, "natural selection" worked in a similar fashion.

What may be most interesting here is to look for cases in which natural selection, or analogous processes in cultural evolution, favored ignoring seemingly useful information. For example, Europeans react swiftly and decisively to the buzzing sound of a rattlesnake's rattle even though Europeans have only a very shallow evolutionary history of interaction with New World rattlesnakes. How could natural selection have favored Europeans reacting adaptively to rattlesnakes' rattling? They are not "deliberately" ignorant of the sound, though it seems they ought to be. As far as we are aware, this is an unsolved puzzle.

Cultural evolution generates cases that are reminiscent of deliberate ignorance. Lewandowsky (this volume) offers the example of a constructivist

populism in which a group of people seem to deny plain material facts as if the world is actually built according to how they would like to have it. Anthropogenic climate change denial is one such example. Cultural evolutionists have considered two models that might be helpful in understanding these types of cases. Boyd and Richerson (1985) considered a model of prestige-based bias. People often use apparent success as a cue to determine who they should imitate. Boyd and Richerson imagined a display trait, such as the size of yams that farmers bring to public celebrations on the island of Pohnpei in the Pacific. The farmer with the biggest yams might originally, and reasonably, have been judged to be the best farmer, and young farmers might have attempted to imitate his practices. This system, however, can run away maladaptively, as ever bigger yams become preferred. In fact, on Pohnpei, ceremonial yams evolved to weigh a hundred kilos and their growers were accorded much respect. The skills for growing them, however, became uncorrelated with ordinary farming success. Many exaggerated cultural practices evolve, such as elaborate, costly, nonfunctional rituals. The affected individuals seemingly remain ignorant of the costs of such practices. On the other hand, students of ritual often suggest that costly rituals actually do have functions in proportion to their costs (Rappaport 1979), much as apparent cases of deliberate ignorance at the individual level have, upon close examination, sensible functions.

Consider another example: symbolic markers of groups. Young children prefer to imitate adults who speak the same dialect as they do (Kinzler et al. 2009). McElreath et al. (2003) showed that neutral symbolic group markers, combined with a bias to imitate those like you on the symbolic marker, can evolve when groups differ in their norms for solving games of coordination. Other groups may be sources of useful information as well, but if imitating them is sufficiently likely to mis-coordinate you with your neighbors, it will be adaptive to remain ignorant of them. An interesting case is the one in which individuals choose to live in socially constructed realities that insulate them from learning adaptive facts, as in the current climate change denial situation. This could arise when a symbolic marker is not neutral. Climate change deniers refuse to accept actionable evidence about rising sea levels even when they own a valuable property on the seashore. Perhaps inadvertently, or due to malicious propaganda, denial has become a marker of belonging to an emotionally salient group, such as a resident of a noncoastal state in contemporary United States. Whether by mistake or by imposition, this case is analogous to error, which is outside the narrowly defined cases of individual deliberate ignorance.

It should be emphasized that each of these evolutionary or cultural examples is only analogically related to deliberate ignorance as precisely defined at the beginning of this chapter. The examples either lack deliberateness unambiguously, or their deliberateness is somewhat speculative. Thus, our discussion in this section is meant to be evocative, not decisive.

Conclusion

Throughout our discussions, we aimed to delineate and clarify the phenomenon of deliberate ignorance. We have been restrictive in our definition of deliberate ignorance, to make clear cases salient and push others that merely resemble these cases to the background. In doing so, we have tried to focus attention on the class of phenomena most in need of additional empirical and theoretical investigation. In relation to other aspects of deliberate ignorance being investigated, we hope this clarifies what phenomena modelers need to model, which normative implications need to be evaluated, and which institutional policy concerns need to be addressed. We say this in full recognition that there may be a great deal that deliberate ignorance, strictly defined, has in common with the other members of its “Wittgensteinian” family.

Going forward, as researchers continue to investigate deliberate ignorance and its effects, care must be given to distinguish the intended and expected consequences of deliberate ignorance from its actual consequences. If someone is deliberately ignorant for good reason but the effect backfires, we might treat this as an error in forecasting that needs correction, rather than a rebuke of deliberate ignorance per se.

Institutionally, Enlightenment ideology notwithstanding, we think it unwise to design programs to reduce deliberate ignorance indiscriminately. Instead, we urge the development of guidelines that help individuals and institutions judge accurately when deliberate ignorance will enhance welfare (see Bierbrauer, this volume) and when it will reduce it.