

# Reflections on Deliberate Ignorance

Lewis A. Kornhauser

## Abstract

Many different definitions of “deliberate ignorance” may be derived from the ordinary usage of these two terms. “Ignorance” may refer to an absence of belief, to an unjustified belief, to disregard of a fact, or to use of a fact known to be false. “Deliberate” may refer to a direct decision not to know some fact  $F$  or an indirect decision to know  $F'$  rather than  $F$ . An individual may be deliberately ignorant but so may a group be. These different interpretations of deliberative ignorance raise different issues in different contexts. This essay develops a taxonomy of accounts of deliberate ignorance, suggests the criteria one might use to select among definitions, and identifies some normative questions that arise from them in a selection of contexts ranging from debates over individual rationality to questions in political philosophy.

## Introduction

Hertwig and Engel argue that psychology has largely ignored an important set of phenomena that they, ironically, call *deliberate* ignorance (Hertwig and Engel, this volume, 2016). They define deliberate ignorance as “the conscious individual or collective choice not to seek or use information (or knowledge)” and note that they are particularly interested in situations in which the marginal cost of knowledge acquisition is low and the expected benefits high. They then offer a functional taxonomy of deliberate ignorance, discuss why it might be normatively desirable, and suggest modeling strategies.

In this essay, I offer a conceptual rather than functional taxonomy of deliberate ignorance.<sup>1</sup> This perspective sets deliberate ignorance within a more general framework that focuses on the distribution of knowledge and information. This framework raises questions concerning the appropriate scope of a concept

---

<sup>1</sup> I develop this taxonomy without regard to the net benefits of information acquisition. As the subsequent discussion suggests, the requirement of low marginal costs of information acquisition does not fit well with Hertwig and Engel’s interest in instances of collective deliberate ignorance.

of deliberate ignorance in the study of psychological and social phenomena. It also has implications for our understanding of both rationality and normative questions in moral and political philosophy.

The argument relies on the literature on extended cognition and the extended mind, which argues that individual knowledge does not rest solely in the mind of the individual but also in the minds of others and in other artifacts. This argument thus lessens the gap between individual and collective knowledge. Questions of “deliberate ignorance” shift into questions about the distribution of knowledge and decision-making authority.

The discussion begins with an analysis of the concept of deliberate ignorance. It then assesses its implications for understanding the norms of rationality and discusses issues pertaining to political and moral philosophy.

### What Is Deliberate Ignorance?

In the paradigmatic case of deliberate ignorance, an individual, Liza, consciously chooses not to know or to learn some fact: Liza’s relative is diagnosed with Huntington’s chorea; Liza decides not to discover whether she has the gene responsible for the disease. Liza, one may clearly say, is deliberately ignorant. Here, “deliberate” means the decision was at least intentional but possibly, and more strongly, reasoned. “Ignorance” apparently means “ignorance of the state of the world” (or perhaps of some set of states of the world or some aspect of a state of the world).

Even at this individual level, however, neither the characterization of “deliberate” nor of “ignorance” captures all the phenomena that ordinary language may encompass or the set of phenomena that might engage the psychologist or decision theorist. When deliberate ignorance is considered in the context of collectives, the inadequacy of these characterizations becomes more glaring.

Below, I suggest that there are three distinct senses that one might attach to “ignorance” and two ways to understand “deliberate.” I thus identify a  $3 \times 2$  taxonomy of “deliberate ignorance.” The senses of “ignorance” and “deliberate” derive directly from common usage of the two terms.

### What Is Ignorance?

The three senses of “ignorance” elaborated here reflect the standard usage of the noun “ignorance” and the verb “to ignore.” I begin with a definition suggested by the verb. The Oxford English Dictionary defines “to ignore” as “to disregard intentionally.”<sup>2</sup>

<sup>2</sup> This refers to the third definition; the first, identified as obsolete, defines “ignore” as “not to know.” Similarly, Merriam-Webster defines “ignore” as “to refuse to take notice of”; <https://www.merriam-webster.com/dictionary/ignore> (accessed Jan. 12, 2020).

Ignorance as “disregard” may come in a weak or a strong form. “Weak disregard” means simply that agents pay no attention to some of their knowledge. This occurs, for example, in the take-the-best heuristic (e.g., Gigerenzer and Goldstein 1996). This heuristic is deployed when the agent must choose among a set of options. For each option  $x$ , the agent distinguishes  $n$  characteristics  $(x_1, x_2, \dots, x_n)$  and chooses based on a lexical ordering. As between any two options,  $x$  and  $y$ , the agent first compares the options on dimension 1. If  $x_1 > y_1$ , the agent chooses option  $x$ . If  $x_1 < y_1$ , option  $y$  is chosen. Otherwise the agent proceeds to compare the two options along the second dimension. When the options differ in the first dimension, the agent disregards its knowledge of the other  $n - 1$  characteristics. If the first  $k$  elements of the two options agree but differ at the  $k + 1^{\text{st}}$ , the agent chooses on the basis of the  $k + 1^{\text{st}}$  element and thus disregards the final  $n - k$  elements.

A similar form of disregard occurs in impartial decision making when the decision maker is instructed to disregard decision-irrelevant information. For instance, juries are often instructed to disregard certain testimony that, though heard, is deemed inadmissible or irrelevant. Antidiscrimination law in the United States instructs employers to disregard race, gender, ethnicity, national origin, religion, and age (over 40) when making decisions on hiring, pay, and promotion. One might understand the introduction of double-blind procedures as a replacement for the instruction to ignore the information of which treatment the patient has received, largely because evidence suggests that conscious attempts to ignore do not fully filter out the effects of the information. Similarly, some orchestras have moved to blind auditions to reduce unconscious bias or use of tainted information.<sup>3</sup>

As noted, one may disregard information in more or less radical ways. The take-the-best heuristic simply puts the disregarded information to one side; the agent makes no use of this information even when it is, at least superficially, relevant. The action chosen is the same regardless of the specific content of the disregarded piece of information. Let us call this “weak disregard.”

One may “disregard” information more assertively as in the construction of a model. Here, an analyst may not simply disregard knowledge that she has, she may assume its negation. In studying the acceleration of a ball down a plane, the model may ignore friction, in the sense of assuming it does not exist. Let us call this form of disregard of knowledge “strong disregard.” The presence of false assumptions in models has provoked controversy over how models explain the phenomena under study.<sup>4</sup>

<sup>3</sup> Both examples appear to fall within the functional taxonomy offered by Hertwig and Engel (this volume, 2016). The take-the-best-heuristic may fall within the category of performance-enhancing devices while the structure of moral rules clearly falls within their category of impartiality and fairness devices.

<sup>4</sup> The literature in philosophy of science on false assumptions is voluminous. Friedman (1953)

Let us now turn to “ignorance,” defined by Merriam-Webster as the “lack of knowledge.”<sup>5</sup> Philosophers analyze knowledge<sup>6</sup> in significant part as “justified true belief.” On this account, an individual could be ignorant in at least three different ways: (a) she might have an *unjustified* true belief, (b) her belief might not be true, or (c) she might not have a belief at all. These different paths to lack of knowledge suggest different forms of ignorance.

An unjustified belief might arise in at least two different ways. First, the agent may lack evidence but nonetheless hold a firm belief. A Bayesian would say that she had an “improper” prior, though it is not clear that priors can be improper. In any case, it is hard to see how this type of lack of justification could be deliberate. Second, the agent might have erred in the process of forming her belief. She may not have processed the evidence properly. A formation error, of course, could be deliberate.

Now consider agent Liza with a false belief. She too lacks knowledge and hence, on the dictionary definition, is ignorant. Yet, as discussed below, in some contexts, her false belief might be deliberate.

Finally, agent Liza might not have a belief. The absence of belief might refer to a number of different conditions. Suppose, for instance, that Liza has a parent who has Huntington’s chorea. Liza knows that she has the genetic predisposition for the disease. There are two states of the world:  $H$ , in which she has the disease, and  $\sim H$ , in which she does not have the disease. Presumably, prior to being tested, she believes she is equally likely to have the disease as not. She is clearly ignorant in the sense that her beliefs about the state of the world are in perfect equipoise. She has no belief in the sense that each relevant state of the world is equally likely.

Ignorance, however, clearly extends beyond situations where beliefs about the state of the world are in perfect equipoise. Suppose Liza learns that her brother Freddy carries the sickle cell trait. Prior to any test, she should believe that she too carries the trait with a probability greater than .5.<sup>7</sup> In this instance,

---

advocated the irrelevance of the falsity of assumptions. Cartwright (1983) argued, as her title suggests, that the laws of physics are false. Batterman (2006), Elgin (2004), Rice (2015, 2018), and Strevens (2019) are more recent contributions to the literature.

<sup>5</sup> The Oxford English Dictionary more quaintly defines ignorance as “a want of knowledge.”

<sup>6</sup> As noted above, Hertwig and Engel define ignorance as a lack of *information* or knowledge but then state that they equate “knowledge” and “information.” But information and knowledge are not identical. I confine my attention to “knowledge” as the meaning of “information” differs across various disciplines (e.g., computer science, cognitive science, linguistics, logic, and semantics).

<sup>7</sup> The calculation is a bit more complex. She knows that at least one of her parents carries the trait. If only one parent carries the trait, then there is a probability of .5 that she carries it. If both parents carry the trait, the probability rises to .75. The actual probability depends, however, on the ancestry of each parent. In any case, her belief that she has the trait should be greater than .5. Moreover, she must condition her prior on the evidence that neither she nor her brother has sickle cell anemia, a

we would surely say that Liza does not know whether she has the sickle cell trait; more strongly, we might say that she is ignorant of her state. But it is unlikely that we would say that Liza has no belief.

This example illustrates the simplest case in which the probabilities of all states of the world are known and objective. At the other extreme, an agent may face Knightian uncertainty under which there is no quantifiable information about the world; indeed, Liza might not even know what all or some of the possible states of the world are.

The possibility of radical uncertainty of this type may underlie a sense that the term “deliberate ignorance” is oxymoronic and the actual state paradoxical. This sense, however, is mistaken. Consider the world in 1492: many people believe that the world is flat, yet Christopher Columbus believes that the world is round. When Columbus sails west to reach China, what is his belief that he will, in fact, land on Hispaniola? More precisely, does he have a belief that corresponds to the actual map of the Western Hemisphere?<sup>8</sup> Presumably not. Yet, even in this instance, one might think that deliberate ignorance is possible. Columbus believes that China lies to the west of Spain. He might ask himself whether the water between Spain and China is empty. He could remain deliberately ignorant of whether the western seas were empty by refusing to sail west. Even if the world is in a state that Columbus has not even imagined (e.g., shaped like a torus), deliberate ignorance would still arise if he intentionally chose not to explore these waters.<sup>9</sup>

Finally, not all knowledge may require beliefs or propositional knowledge. Distinguish *knowing that* (propositional knowledge) from *knowing how* (an ability to do something, which I shall call practical knowledge).<sup>10</sup> On some accounts, *knowing how* is distinct from *knowing that*; the former does not involve any propositional knowledge. Riding a bicycle, for example, does not rely, at least superficially, on any propositional knowledge. It requires an ability or a disposition but not obviously any beliefs. An individual who knows how to read a book or play the piano has an ability that extends beyond propositional knowledge and belief.

Deliberate ignorance of practical knowledge may seem straightforward. Liza decides not to learn to play the piano or to ride a bicycle. She may have

---

fact which should affect her belief that both parents carry the trait.

<sup>8</sup> Columbus must have had some beliefs about the water to the west. He must have believed, for instance, that there was not a wall, barrier, or land mass that would prevent him from reaching China by sailing west.

<sup>9</sup> The required intention is difficult to state precisely. Columbus must have decided not to sail to avoid learning about the extent of the emptiness of the western seas.

<sup>10</sup> The distinction seems to be due to Ryle (1946). It is controversial whether practical knowledge is reducible to propositional knowledge. For a survey of the controversy over the reducibility of *knowing how* to *knowing that* see Bengson and Moffett (2012).

good or bad reasons for her decision but it hardly seems to raise interesting behavioral or philosophical issues.

Practical knowledge, however, is pervasive and some of it does present more complex questions. Technology, for example, involves *knowing how*. Some of that knowledge is embodied in the technical means by which it is applied, but some is possessed by individuals. When technology changes, both the physically embodied and individually possessed knowledge may be lost. Socrates noted this in the *Phaedrus* when he suggested that the advent of writing would undermine philosophical argument. He claimed not that individuals would lose propositional knowledge but rather their dialectical ability.

Other, more recent examples abound. The advent of the electronic calculator arguably has eroded the arithmetic facility of everyone. Similarly, when Liza downloads Google Maps or Waze to her smart phone, she initiates a process that may lead to her losing her navigation abilities. When her ability to navigate atrophies, she does not obviously lose propositional knowledge. Her capacity to move from point A to point B simply deteriorates.

One might understand the effects of adopting these technologies in a way that implicates *knowing that* as well as *knowing how*. One might understand these phenomena as “externalizing knowledge.” Knowledge may be externalized either physically or semantically (for a discussion on the extended mind, see Clark and Chalmers 1998; for essays on extended cognition, see Carter et al. 2018). In both instances, externalization is apparently the result of the specialization of labor and knowledge.

Many, if not most, of our concepts are deferential. When the pipes in my home are clogged, I know to call the plumber but I am innocent of knowledge of plumbing in general. Indeed, my (and others’) understanding of most of current technology is minimal at best. I am thus to a large extent ignorant about the content of most of these concepts.

The earlier discussion of *knowing how* and technology illustrates how knowledge might be externalized physically in the technological aids that an agent uses to understand her world or to act. The invention of writing externalized memory; the invention of moveable type and then digital storage accelerated that externalization. Similarly, counting boards, abaci, and then the computer (partially) externalized computation. Maps and street signs facilitated navigation.

These physical objects reduce, in some sense, each individual’s knowledge; at the very least, it permits Liza to deepen her knowledge about some domain *D* at the cost of a shallower knowledge of other domains. Collective knowledge clearly grows. This dispersion of knowledge throughout the population, however, blurs the distinction between individual and collective ignorance and, as I suggest later, renders assessment of “individual” rationality more complex.

Deferential concepts are instances of *knowing that*. As noted above, although they are pervasive, they render knowledge inherently communal as

each individual essentially relies on some knowledge that is mastered by someone else. Deferential concepts thus emphasize the distribution of information within a collectivity. Whether an agent *A* is ignorant of *X* thus becomes a complex question as the agent may, in fact, “know” *X* deferentially but nonetheless be unable to act or to reason effectively with the knowledge not located in her brain. These issues will be discussed further below (see section on collective ignorance).

### **When Is Ignorance Deliberate?**

Broadly, a deliberate action might either be intentional or, more strongly, reasoned. If deliberate ignorance requires a reasoned decision, very few instances of deliberate ignorance will exist. Moreover, if they exist, they are apt to be rational for the individual.

If deliberate ignorance entails only an intentional decision to ignore, then there will be many more instances of deliberate ignorance. How many will depend on the nature of the required intentionality. The agent Liza may directly choose to ignore *X* or she may indirectly choose to ignore *X* by intentionally choosing *Y*. This phenomenon may arise in different contexts, each of which presents difficulty.

Consider first how passive ignorance puts pressure on the idea of reasoned deliberation. Suppose Liza is a college student designing her curriculum. She is, in essence, deciding what to know. She develops a curriculum around decision making, choosing courses in psychology, economics, and political science. Has Liza chosen to be ignorant of physics? Does the determination of whether her ignorance of physics is deliberate depend on whether she explicitly considered taking physics? Suppose she simply disregarded all language courses? Does this amount to deliberate ignorance of the relevant languages?

Another perplexing instance arises when ignorance results from an unanticipated consequence of an intentional act. Suppose Liza buys a smartphone. She downloads Google Maps and uses it extensively when driving. Consequently, her navigational abilities deteriorate. Similarly, she actively uses a search engine to “remind” herself of various facts. Her memory thus deteriorates. Her explicit propositional knowledge declines though the scope of her knowledge expands (and the cost of accessing that external knowledge falls). Are the changes in her propositional and practical knowledge instances of deliberate ignorance?

Ignorance as disregard also raises some questions. Suppose Liza chooses between *X* and *Y* on the basis of only two out of ten criteria relevant to her choice. She thus disregards eight of the relevant criteria. Does it matter if she did this intentionally? Or does deliberate ignorance encompass decision protocols that evolved or are otherwise adaptive, independent of whether the agent acted intentionally or not?

The collective context highlights the importance of the distribution of information and the distribution of decision making within the collectivity.

Consider a medical trial of a new drug: the relevant group is the experimenter, the doctors administering the drug, the patients receiving it, and the data analysts. The experimenter adopts a double-blind protocol for the drug trial. The experimenter makes a reasoned choice to withhold knowledge from doctors, patients, and data analysts. In this case, the decision and the consequence of the decision—the ignorance—are separated. What matters is the distribution of information.

What would justify choosing ignorance over knowledge? Consider first *efficiency* concerns. It may be costly to acquire the information or to process the information. Moreover, it may be that decision protocols which ignore information prove to be more accurate. These types of reason underlie the category of performance-enhancing functions noted by Hertwig and Engel (this volume, 2016).

There are, however, other reasons to choose ignorance. First, one might want to filter out information that distorts or interferes with successful processing of information. This logic underlies the use of double-blind, randomized controlled trials and, on at least one account, Rawls's "veil of ignorance."<sup>11</sup>

Second, and relatedly, there may be normative reasons to suppress, or at least disregard, certain information. Rawls (1971) justifies the veil of ignorance as a fair procedure used to identify the considerations that an agent should weigh in choosing a set of institutions. Institutional choice should not, he argues, depend on an individual's actual position in society, but only on the distribution of social and economic advantages. Hertwig and Engel (this volume, 2016) also note this reason.

Civility provides a third reason for deliberate ignorance. Social life requires individuals to restrain their aggression. In many instances, too much information may provoke conflict. Often "manners" require an individual to disregard information that she has acquired or to withhold it.

Whether someone has deliberately chosen to be ignorant is, in many instances, straightforward. The agent explicitly averts her eyes or weighs considerations in favor or against acquiring the relevant knowledge. Problems arise, as will be discussed below, in the context of forgetting and remembering. Here, one might consider a deliberate decision to remember *X* as at least an implicit decision to forget *Y*.

### **Deliberate Ignorance within the Collective Frame**

As discussed above, the concept of individual ignorance cannot clearly be distinguished from the idea of collective ignorance. But collective ignorance raises perplexing problems. Deliberate collective ignorance doubles the perplexity, as both ignorance and deliberation apparently implicate mental states.

---

<sup>11</sup> Note how the randomized control trial transforms the injunction to disregard information inherent in the veil of ignorance into a true lack of knowledge on the part of the agents. If the experimenter is one of the doctors in a drug trial, she has effectively tied her hands.



An attribution of mental states to collective entities, however, seems simply to be a metaphor. Several issues arise:

- What is collective knowledge?
- By what means should it be defined?
- Does the collectivity know  $X$  if some individual within the collective knows it nondeferentially? If so, then the collectivity could know (non-deferentially) more than any individual within it.
- Is it possible for the collectivity to know something that no one in the collectivity knows?
- Conversely, can everyone in the collectivity know something that the collectivity does not know?

If we conceive of collective knowledge simply as the distribution of beliefs within the collectivity, then we would answer these questions negatively. If, however, we attend to the social processes within the collectivity, the answer could change.

Consider first how the collectivity might have knowledge of how to do something that none of its individual members has, as in market processes. Each consumer knows how much they are willing to pay to purchase a good while each producer knows how much they are willing to accept to bring goods to the market. No one knows, however, what price will clear the market; the market process produces this knowledge. Of course, the market has no belief about this knowledge. A striking example, raised by Tom Seeley (2001), involves the choice of a new hive by a swarm of honeybees. Although individual bees are unable to reach this decision on their own, some know how to dance and are thereby able to communicate pertinent bits of information that leads the swarm to its decision. Here, *knowing how* is not propositional knowledge as bees presumably do not have beliefs.

One might argue that the market price does not truly represent collective knowledge because the collectivity cannot reason with that knowledge. One might say, in response, that the collectivity knows how to allocate the resource through the market. Alternatively, one might point to bureaucratic organizations as collective entities that do, in fact, process collective knowledge. A bureaucracy segments decision-making authority as well as access to information. This fragmentation implies that each individual has significantly less knowledge than the bureaucracy as a whole.<sup>12</sup>

To understand how everyone in a collectivity might know something without the collectivity knowing it, one must also consider the nature of the social processes within the collectivity. As noted, collective knowledge might be understood as the distribution and nature of the beliefs within society. Yet

---

<sup>12</sup> Arguably we should understand the bureaucracies' knowledge as externalized in its files just as historians' knowledge of some events are externalized in their computer files that archive relevant facts.

characterizing collective knowledge simply as knowledge that each individual member has does not acknowledge the collective nature of the knowledge. Each individual may also need to have some higher-order belief about the knowledge of  $X$  by others. That is, collective knowledge of  $X$  might require mutual knowledge, understood as the belief by each person that (every) others knows  $X$ . Alternatively, it might require common knowledge, understood as the infinite cascade of beliefs, that each knows that the others know that each knows that..., or any finite cascade in between. If members of the collectivity do not have mutual knowledge, then the collectivity cannot act collectively on that knowledge, although each individual could.

This conception of collective knowledge would seem to have broad application across different types of collectivities, which may range from a mass of diffuse, unorganized individuals to highly organized and structured groups such as the modern state. We often attribute beliefs, preferences, and other attitudes to these structured groups but the procedure and the demands of rationality that we place on these structured groups is contested. When the attitude of the group depends on the attitudes of the individuals that comprise it, serious logical difficulties must be confronted (Arrow 1963; List and Pettit 2013). Organic accounts of these groups are equally problematic.

Both ignorance and *deliberate* ignorance would take different forms on the different accounts. Suppose that one thinks collective knowledge consists of a set of beliefs at least as strong as mutual knowledge: deliberate ignorance, then, might take the form of interventions that block the formation of the higher belief that the agent knows that others know. Organizations create bureaucracies to accomplish such knowledge segmentation. Banks that provide investment advice to firms as well as to investors construct “Chinese walls” to prevent knowledge from crossing the divide between the two divisions. States that seek to control social media and the Internet arguably pursue this strategy to prevent the coalescence of opposition to the regime.<sup>13</sup> By contrast, the form that deliberate ignorance would take on aggregate accounts of belief in structured groups is unclear. In some instances, the group might require the subgroup making a decision to disregard certain information, as the rules of evidence do in jury trials (see Zamir and Yair, this volume)

Because information and action are distributed across individuals within a group, the attribution of deliberate ignorance to the group raises issues not present at the individual level. Suppose Henry is the CEO of a corporation, Liza its CFO, and Freddy a sales manager. Suppose Freddy submits fraudulent sales reports to Liza, who uses these reports to prepare a financial report for

---

<sup>13</sup> This example illustrates a second difficulty with the application of collective ignorance to collectivities. On one account, we can view the state as a collectivity itself so that the censor is influencing the knowledge of members of the collectivity. Alternatively, one might treat the state as the “government” so that its actions censor individuals external to the group. While the former has the structure of the deliberate ignorance of the collectivity, the latter does not.

Henry. Henry then makes statements based on the report. These statements, of course, are fraudulent. Was the corporation ignorant of the fraud? We might ask the same question of Henry and Liza. In both the collective and individual cases, the attribution of legal or moral responsibility will depend on what knowledge is imputed to each entity/individual, as no single agent in the corporation has all the information necessary for the attribution of responsibility.

### **Defining Deliberate Ignorance**

The prior discussion suggests several different definitions of deliberate ignorance:

1. An intentional lack of knowledge
2. An intentional lack of knowledge or weak disregard
3. An intentional lack of knowledge or disregard, either weak or strong
4. An actively or passively intentional lack of knowledge
5. An actively or passively intentional lack of knowledge or weak disregard
6. An actively or passively intentional lack of knowledge or disregard, either weak or strong

One could construct other definitions, for instance, that restrict “ignorance” to “disregard” or look solely at “lack of knowledge.” These alternative definitions, however, exclude some of the core instances of deliberate ignorance. The choice of definition does not depend on a proper analysis of a folk concept of deliberate ignorance but rather rests on the questions one wishes to ask, as well as on the answer and the intellectual fertility of the definition.

Hertwig and Engel (this volume, 2016) suggest the first definition, with the added condition that the net benefits of knowledge acquisition be large. This definition fits well with an inquiry into individual psychology. If the aim is to identify conditions under which deliberate ignorance constitutes a rational response to a decision problem, however, the motivation for the restriction to situations of low costs of acquisitions is unclear. When costs are high, of course, the rational motivation for deliberate ignorance may be straightforward, but the question and answer seem to fall well within the general research question asked.

Hertwig and Engel occasionally seem to suggest the second definition. As noted earlier, the functional categories of performance enhancement and fairness suggest that disregard falls within the scope of (deliberate) ignorance. Examples of collective deliberate ignorance suggest a further broadening to include the fifth definition, which includes ignorance of *Y* (that arises from a choice to know *X* instead), disregard, and intentional ignorance.

On what basis should a definition of deliberate ignorance be selected? There are at least two different approaches. The first simply explicates the “folk” concept of deliberate ignorance. Though I have presented a taxonomy in this

section based on the ordinary usage of the terms “ignorance” and “deliberate,” I have not offered an analysis of the folk concept. (It is not even clear that a folk concept of “deliberate ignorance” exists.) Rather, my aim was to identify some nuances in the meanings of the two terms “ignorance” and “deliberate,” to distinguish more finely among various situations. This process of refinement might help with the second approach to defining deliberate ignorance; namely, to ask what definition is most intellectually fertile and helpful. As Hertwig and Engel may have adopted their definition based on investigations into individual decision making, rationality, and psychology, the second approach might be the most appropriate. Nonetheless, this approach seems too narrow to encompass investigations of collective decision making as well as to address some normative questions in moral and political philosophy.

### **Rationality**

Broadly, analysts identify two distinct domains of application of the concept of rationality: decision outcomes (or choices) and decision processes.<sup>14</sup> To these domains they bring three distinct, but related, concepts of rationality: instrumental rationality, procedural rationality, and substantive rationality. Within each class of concepts, numerous distinct conceptions of rationality have been elaborated. In game theory, for instance, one might understand each proposed solution concept as a conception of instrumental rationality.

Instrumental rationality requires the agent to choose the best means to achieve her ends. Procedural rationality requires the agent to follow appropriate procedures when making her choice; she should use “right reason” to reach her decisions, form her beliefs, and argue successfully. Substantive rationality requires the agent to have appropriate ends.

Superficial reflection on theoretical rationality, or the rationality of belief, suggests that instrumental, procedural, and substantive rationality converge. In theoretical rationality, after all, the obvious substantive end is truth; substantive rationality is thus straightforward and trivial. Procedural rationality in deduction points to the rules of logic as both instrumental and normative as they are “truth-preserving.” They identify rational decision processes and reach from true premises rational outcomes. In induction, procedural rationality points to statistics and decision theory as correct processes for reaching correct outcomes. Together, these norms apparently condemn deliberate ignorance.

When one considers practical rationality, however, this convergence is less clear. Many different goals compete for the ends defined by substantive rationality. Procedural rationality might point to the same norms to govern decision processes but practical concerns (i.e., the costs of decisions and limited cognitive capacity) suggest that these procedures may not be instrumentally

<sup>14</sup> I include the processes of belief formation and adjustment and of argument in this category.

best, even if implementable. The rational assessment of deliberate ignorance requires a more nuanced analysis.

This section lays out a more nuanced assessment. It begins with discussions of procedural and instrumental rationality in the individual case, and concludes with a discussion of these issues in the collective context.

### **Procedural Rationality**

Deliberate ignorance seems irredeemably at odds with procedural rationality. This tension, however, is not as strong as it appears. In some circumstances, procedural rationality endorses deliberate ignorance.

Procedural rationality requires agents to use “right reason” in their deliberations. These accounts of rationality generally consider “right reason” to include a narrower or wider set of reasoning protocols that are normative for actual reasoning processes. Classical logic constitutes the core of these normative rules. Around the core sit the rules of probability theory, including Bayer’s Rule for revising one’s beliefs in light of new evidence. On the periphery, lies decision theory, generally subjective expected utility theory, and further out still, game theory. Decision theory assumes that agents seek to maximize their “expected utility” subject to constraints. Game theory analyzes strategic interactions among rational agents; each solution concept offers an account of rationality in these circumstances.

Over the last forty years, some psychologists claim to have demonstrated that individuals are procedurally irrational; in general, humans do not conform to the normative standards of any of these reasoning protocols. This discrepancy between actual reasoning processes and logic, probability theory, and decision theory has provoked a vigorous debate over the nature of rationality understood as right reason (Cohen 1981; Stanovich 2010; Stanovich and West 2000). These experimental results, of course, do not challenge instrumental conceptions of rationality directly. Indeed, it is not clear that these results challenge right reason conceptions of rationality at all.

There are two subclasses of conceptions of rationality as right reason. One position holds that the norms of reasoning are grounded outside of psychology. This position derives from the late nineteenth century, when Frege separated logic from psychology (Hanna 2006). Logic, for Frege, was not an empirical study but an *a priori*, necessary investigation of rules of inference. It was decidedly not the study of how people actually reasoned. A fully externalist account must then hold that probability theory and decision theory are also externally justified. Call this view “external procedural rationality.” On this account, it is superficially easy to see how the rationality of individual reasoning processes could diverge from the normative requirements of rationality: the agent’s reasoning processes simply do not always conform, as the experimental evidence shows, to rules of classical logic, to the logic of probability, or to the logic of decision theory.

A second position holds that the norms of reasoning are constituted by human psychology through a process of reflective equilibrium (Cohen 1981). The rules of logic, probability theory, and decision theory are only normative if they are endorsed by a process of reflective equilibrium that reconciles judgments about appropriate rules of inference and about successful instances of inference. Call this view “internal procedural rationality.”

The internalist argument often proceeds by analogy to Chomsky’s program, which distinguished between a syntactic *competence* and syntactic *performance*. On this account, humans are inherently rational; deviations from the norms of reasoning occur only because something interferes with the agent’s exercise of her innate rational competence and yields a performance error.

Much of the philosophical discussion has focused on the debate between the external and internal conceptions of procedural rationality, and much of this debate has, as noted, relied in some form on the performance/competence distinction. The externalist contends that the norms of reasoning do not derive from human reasoning competence. The experimental evidence thus does demonstrate human irrationality. The internalist, by contrast, argues that rationality is simply defined by human reasoning competence. Departure from the norms of rationality are attributable to performance errors, not to a rational incapacity.

Deliberate ignorance plays an ambiguous role in right reason concepts of rationality as it might either promote or undermine such accounts. Deliberate ignorance undermines all of the standard norms of reasoning to the extent that it discards or disregards decision-relevant information. The take-the-best heuristic, for example, may be described as deliberately ignoring decision-relevant facts. The heuristic, therefore, might counsel a course of action that the agent knows (or would know) is at odds with the conclusion drawn on the balance of *all* reasons. This situation underlies one of the paradoxes of authority.

On the other hand, deliberate ignorance seems to support or constitute procedural rationality in other ways. On both the internalist and externalist right reason accounts, deliberate ignorance would be a rational strategy when it filtered out stimuli that interfered with the exercise of right reason. During the Vietnam war, for example, the United States changed a draft system that granted local boards great discretion to one that, to some extent, relied on a lottery to determine who was drafted.<sup>15</sup> More interestingly, advocates of “resolute” choice in dynamic decision-making contexts argue for what may be a form of resolute ignorance (Gauthier 1986; McClennen 1990).<sup>16</sup>

Consider a slightly modified version of the classic difficulty faced by Odysseus, who wants to travel from Troy to Ithaca. There are two routes: one

<sup>15</sup> The system was supposed to produce a fairer allocation of the burden of conscription across income classes and races.

<sup>16</sup> Admittedly, one might construe resolute choice as a form of instrumental rather than procedural rationality.

passes by the island of the Sirens and the other does not. He faces a dynamic choice problem. At time 1 he chooses a route: either route *D*, which does not pass by the island of the Sirens, or route *S*, which does pass by it. If he chooses *S*, Odysseus faces a second choice when he passes the island at time 2: continue on to Ithaca or crash into the rocks by the island. He thus faces three plans: *D* which leads to a safe return to Ithaca (*I*), plan *SI* which leads to hearing the Sirens' song and a safe return to Ithaca (*SI*); and plan *SD* which leads to hearing the song of the Sirens and death (*SD*). At time 1, Odysseus most prefers *SI*, then *I*, then *SD*. At time 2, however, conditional on having chosen *SI* or *SD*, Odysseus prefers *SD* to *SI*.

On the standard analysis, Odysseus is either naive or sophisticated. The sophisticated Odysseus knows that the Sirens' song will cause him to choose *SD* over *SI* at time 2. He thus understands that his true choice is between *I* and *SD*.<sup>17</sup> Sophisticated Odysseus chooses *I*. Naive Odysseus believes he will choose *SI* at time 2 so he chooses *S* at time 1 and crashes into the rocks at time 2.

Fortunately, Odysseus receives expert advice from Circe, who offers him two strategies for choosing *SI* over *SD* at time 2. First, Odysseus may literally tie his hands to the mast to prevent himself from steering into the deadly shoals; that is, he *commits* not to steer onto the shoals. Second, Odysseus can fill his ears with wax to avoid hearing the Sirens' song. This second strategy is one of deliberate ignorance, as Odysseus chooses not to hear (learn) the Sirens' song.<sup>18</sup>

Advocates of resolute choice believe that there is a third way that fuses commitment and deliberate ignorance of a different type to get past the island. They suggest *resolution*, a mysterious psychological commitment not to choose *SD* at time 2. At time 1, Odysseus *resolves* to adopt plan *SI*. As noted, plan *SI* presents Odysseus with a choice at time 2 to continue to Ithaca or to crash into the rocks. Resolution means that Odysseus adheres to his plan *SI*; one might say that a resolute Odysseus decides to disregard, at time 2, the choice of *SD* that is available to him and, at time 2, preferred by him to *SI*. On this account, resolute choice would be a form of deliberate ignorance.

Subjective expected utility theory is, as noted, considered part of procedural rationality. On this account, rationality requires the agent to maximize her expected utility. Consequently, it treats an agent's beliefs about states of the world and her preferences over outcomes as mental states.<sup>19</sup> Much research in psychology shows that agents do not in fact maximize their expected utility.

<sup>17</sup> A technically better story would make the worse choice at time 2 subgame perfect. In the story of Odysseus, Odysseus knows that when he hears the Sirens' song, he will act irrationally and steer onto the shoals.

<sup>18</sup> In the *Odyssey*, of course, Odysseus fills his sailors' ears with wax and ties himself to the mast. Thus, he arrives home safely and hears the Sirens' song.

<sup>19</sup> The economics literature typically starts from the opposite direction; it assumes that agents have, as primitives, preferences over actions in an uncertain world. It then proves a representa-

One can, however, interpret subjective expected utility theory not as a theory of procedural rationality but as a theory about the aims of the agent. The analyst can treat the agent as someone who pursues the objective of maximizing subjective utility, though not necessarily explicitly. Subjective expected utility thus defines the goal of agents but does not describe their decision-making process. It serves as a benchmark that identifies what is optimal, not the road that the agent takes to achieve the optimal decision. On this account, subjective expected utility theory falls under the rubric of instrumental rationality.

### **Instrumental Rationality**

The concept of deliberate ignorance appears, at least superficially, to be neutral with respect to instrumental rationality. Whether deliberate ignorance best furthers one's ends will depend on the agent's ends and the relative costs and benefits of further acquisition of knowledge. It may also depend on the strategic situation in which the agent finds herself and the distribution of information among all interacting agents.

The third position on rationality denies that rationality concerns processes of reasoning at all. Rather, rationality is defined solely in terms of its success at achieving goals. Ecological rationality, therefore, identifies performance criteria for success on a decision-making task rather than assesses the process of reaching the decision against either external or internal reasoning norms (Arkes et al. 2016; Gigerenzer and Todd 2012; Schurz and Hertwig 2019).

Ecological rationality thus seems to have a strong appeal. It avoids the claims that individuals are irrational by rejecting the conceptual tie of rationality to reasoning processes. This rejection is compelling in many but not all contexts. An agent trying to solve a practical problem cares primarily about achieving success, not the reasons underlying her solution to the problem. From this perspective deliberate ignorance is unproblematic as long as it better promotes the agent's success than the pursuit of knowledge. Moreover, deliberate ignorance plays a significant role in some ecologically rational strategies. The take-the-best heuristic, for example, directs the agent, when choosing between two alternatives, to follow a lexical rule that dictates the agent to take the alternative which ranks higher on the first (ordered) criterion that discriminates between the two options. Similarly, the recognition heuristic directs an agent choosing, on the basis of some criterion, between two alternatives to choose the alternative she recognizes; if she does not recognize either or recognizes

---

tion theorem that identifies the conditions under which these preferences over actions can be represented by a set of beliefs over states of the world and preferences over outcomes such that the agent ranks action  $A$  over action  $A'$  if and only if the expected utility of  $A$  exceeds the expected utility of  $A'$ .



both, she should revert to another recognition heuristic: randomize (for further discussion and examples, see Gigerenzer and Brighton 2009).

The argument for ecological rationality faces at least two difficulties. First, how do we know that fast-and-frugal heuristics perform well? Evolutionary arguments might suggest that fast-and-frugal statistics perform well enough when evolutionary fitness is at issue, but the relation between most human objectives and evolutionary fitness is loose at best. Deliberate ignorance as a strategy for ecological rationality thus requires validation before it becomes rational to adopt it. But against which criterion do we validate it? What is the agent trying to accomplish? What is her objective function? And what constraints does she face? Given any heuristic, there is apt to be an objective function and a set of constraints for which the heuristic is optimal. There is also apt to be an objective function and a set of constraints for which the heuristic is *not* optimal.

The second difficulty is related to the first: What happens when the agent's objective function is "truth-seeking" or the justification of a theoretical or practical conclusion? What does success in a truth-finding project entail? Can one justify a truth claim? In particular, can it be justified on the basis of a process that deploys deliberate ignorance? In these contexts, right reason would seem to provide the appropriate objective function or, at least, the way to identify whether the heuristic has performed well or not.<sup>20</sup>

These questions are complex and difficult. Double-blind randomized control trials are justified in part by the deliberate ignorance embedded in the procedure. Here, deliberate ignorance promotes truth-finding by eliminating unconscious biases of the experimenter, experimental subjects, and those who implement the experiment. Randomized assignment of subjects to treatment and control groups insures that the experimenter does not unconsciously introduce selection bias whereas the double-blind feature insures that the behavior of subjects and implementers are not influenced simply by knowledge of their group assignment. A similar argument would support disregard if decision-irrelevant information was used in making moral choices.

In other contexts, however, deliberate ignorance does not adequately support the truth-finding process. Consider, for example, formal dispute resolution by courts. Society aims to achieve accurate determinations of responsibility, but trial procedures, at least in common-law countries, do not conform to the standard probability calculus (Cohen 1977). Perhaps society does this to promote a more complex goal than truth-finding. Similarly, though the rules of evidence preclude the introduction of relevant but potentially prejudicial

---

<sup>20</sup> The situation here parallels arguments that identify utilitarianism as the standard against which decision criteria are assessed rather than as a decision procedure itself. This argument is offered to reconcile the observed use of "commonsense" morality, which is largely deontic, with the claim that utilitarianism is the correct moral theory.

evidence, it is unlikely that they would preclude the exclusion of the evidence ignored by many heuristics.

### **Rationality when Cognition Is Extended (and in Collective Entities)**

Thus far, this discussion of rationality has implicitly assumed that knowledge resides in the agent's head. The prior section, however, suggested that knowledge has, at least to some extent, been externalized. The presence of deferential concepts and extended cognition has a significant impact on an understanding of rationality and deliberate ignorance.<sup>21</sup>

At the outset, note that the agent can only act on the reasons she has immediately before her. Here, "immediately before her" excludes reasons embedded in externalized knowledge, thus either embodied in physical objects or possessed by other individuals. Ought the agent acquire this externalized knowledge? If acquisition was costless, then presumably the answer would be yes. As acquiring this external knowledge is not costless, however, it would thus seem that she should, under some conditions, disregard it.

Deferential concepts and extended cognition imply that the agent has some reasons that are not, in some sense, before her. For instance, the reasons may be in someone else's head. How should the agent proceed? Consider, for example, Liza, who must decide whether to undergo surgery or radiation to treat a cancer. She has deferential knowledge about her cancer and its treatment. To make her decision, however, she needs actual, not deferential, knowledge. Should she acquire that knowledge? Or should she defer to someone with the actual knowledge?

Arguments developing this line of thought have appeared in the literature on political authority. Note that the individual with extended knowledge faces the same problems as collective entities. Consider a decision maker in a corporation, say the CEO tasked with making an investment decision. The information needed for that decision is distributed throughout the organization. A good CEO will organize the flow and processing of that information in an effective fashion. The organizational structure will delegate parts of the decision-making task to different individuals, none of whom will act directly on all the reasons the organization has. The organizational structure thus dictates that each individual will be deliberately ignorant of some reasons, and the person who is ignorant will not be same person who decides who should be ignorant. Thus, the ignorance of some of its agents or members promotes the performance of the corporation. Moreover, in some instances, the corporation might have an interest in insuring that information dispersed throughout the organization is not in fact brought to the attention of the final

---

<sup>21</sup> For more extensive discussions of the sociality of reason, see Sloman and Fernbach (2017) and Mercier and Sperber (2017). These discussions are consistent with, but different from, the discussion here.

corporate decision maker. This structure thus deliberately creates ignorance in the corporation.

### **Deliberate Ignorance in Public Life**

Deliberate ignorance lies at the heart of several other issues that arise in political philosophy and public life, in general. Here, I briefly introduce three issues that raise complex normative issues.

#### **Knowledge about Collective Entities**

The earlier discussion suggested the difficulties in attributing knowledge or ignorance to collective entities. Problems also arise, however, when one asks not whether or how such entities know, but also when one examines the nature of our knowledge of or about such entities.

Typically, collective self-knowledge is indexed or statistical. These indices and statistics choose the aspects of a complex entity that it considers relevant and ignores, explicitly or implicitly, other aspects of the group. Consider, for example, any measure of income inequality. The measure obviously summarizes the distribution of income into a single number that conceals both positive and normative information. Specifically, each measure of inequality embeds judgments about the significance of different types of transfer of income. The relative mean deviation measure,<sup>22</sup> for example, is invariant to transfers between people below the mean (or between those above the mean) but not invariant to transfers across the mean. The choice of index thus implicitly determines which transfers are normatively important.

The choice of a particular index or statistic determines what data the analyst gathers and considers. In some contexts, such determinations are made more explicitly, though not always with a full understanding of the consequences of the choice.

France prohibits the state from collecting any data on the basis of race or ethnicity. This policy affects both national debates over discrimination and the formulation, implementation, and evaluation of policies, particularly ones intended to address any (imperfectly observed) discrimination. This policy of deliberate ignorance may further two distinct social goals. First, it may, as in the individual case, serve as a filter against the use of information morally irrelevant to social decisions. It is not clear, however,

---

<sup>22</sup> The relative mean deviation measure of inequality is, roughly, the sum of the absolute value of the deviation of each agent's income from the mean income in the population. Transfers of income between individuals on the same side of the mean thus have no effect on the measure, though we might believe that a transfer from a poor person to a poorer person, in fact, reduces inequality.

that a justification of the policy as a filter succeeds. The filter has no effect on decisions that involve individuals on a personal basis (where race and ethnicity may be evident), and it has a detrimental effect on the implementation of policies to remedy the effects of discrimination against racial and ethnic minorities because the state lacks the information to target remedies to groups that it refuses to “see.” Second, the policy may *express* a commitment to nondiscrimination. This justification is equally problematic. Though the policy of deliberate ignorance formally expresses the state’s commitment to nondiscrimination, society may be rife with discriminatory outcomes. This pattern of outcomes also has an expressive aspect: it shows violations of the nondiscrimination principle.

### Collective Memory

The problems of collective memory arise most starkly in the context of the societal transition from an oppressive regime to a less repressive one (see Ellerbrock and Hertwig, this volume). In these situations, the demands of collective memory may conflict with the benefits of collective forgetting. An oppressive regime may have committed multiple injustices; the more widespread and egregious these injustices, the greater the demand for an accounting of the wrongdoing. Such an accounting, of course, requires remembering the offenses.

To avoid a violent transition, the oppressive regime might require an amnesty prior to ceding power. Amnesty, as its etymology demonstrates, requires a forgetting of the wrongs previously done. Moreover, prospective concerns recommend not simply a legal forgetting but an actual forgetting. Forgetting may facilitate reconciliation between disparate groups and may help avoid a cycle of violence and recrimination.<sup>23</sup>

Truth and reconciliation commissions might be understood as an attempt to balance the two forces of justice and of reconciliation. A commission first seeks to disclose the wrongdoing of the prior administration and to identify individual wrongdoers. This task must be reconciled with criminal prosecution. As noted, individuals may be unwilling to testify about misdeeds when that testimony might subject them to prosecution. After acknowledging the past, reconciliation may require remedial actions, such as reparations of some sort. Successful reconciliation, however, will require that the victims of wrongdoing disregard, to some extent, the wrongdoing of the prior regime.

Similar issues may arise in collective, commemorative decisions. States constantly erect statues, designate national landmarks or monuments, and

<sup>23</sup> The issues here are complex and may parallel some that arise in the context of individuals. The folk advice to “forgive and forget” exemplifies the tension noted above. Forgiveness requires remembrance but the efficacy of forgiveness may require forgetting after the forgiveness; the advice might better be phrased as “forgive, then forget.”

declare public holidays. These decisions create an official memory that may be at odds with the unofficial history preserved by individuals or social subgroups. This discrepancy may lead to political conflict.

The political construction of the past is most evident in the United States in the design of the history curriculum for primary and secondary education. Individual states identify a set of textbooks that local school districts may adopt to present what is, in effect, an official history. The overt politics that occurs in this arena may reflect the more hidden choices and biases of historical research in general. All history selects one narrative out of many potential narratives; correlation across the narratives may essentially “silence the past”<sup>24</sup> (Trouillot 1995).

### **Democratic Education**

The growth and specialization of knowledge implies that each individual has command over, at best, a very narrow segment of that knowledge. This growth and specialization results from and is sustained by deliberate ignorance in the broad definition (5) offered earlier. Though in the Renaissance, a single person might be able to know everything, now it is impossible to know even all of something. This widespread, general ignorance is especially problematic in a democratic society, as citizens should participate actively in policy making, or at least in the evaluation of the policies made by their elected representatives.

Policy formation in the modern world is extremely complex, requiring knowledge of complex scientific issues, social processes, and ethical questions. Consider, for example, the issues that must be investigated to form and evaluate public policy concerning climate change. Effective policy requires knowledge of the underlying science, both to assess the risks and to develop remedies. Understanding the science itself requires knowledge of statistics and decision making under uncertainty. Knowledge of the science, however, is not enough as any remedy must be both politically viable and socially efficacious. The policy maker must predict how individual behavior will respond to the policy. How will agents adjust their behavior in light of the policy? By how much would a carbon tax reduce emissions of greenhouse gases? How would

---

<sup>24</sup> This is Trouillot’s term, who identifies four moments at which silencing may occur: the moment at which (a) the “source” fact is created, (b) facts are assembled, (c) a narrative is chosen, and (d) the narrative is contested. Silencing may occur at any moment. Some agents may choose to leave fewer traces than others. Legend, for example, contrasts the reputed strategy of Samuel Adams to destroy the record of his actions (the source moment) with the strategy of his cousin, John Adams, to memorialize them. According to this legendary account, Samuel Adams deliberately sought to forestall historical inquiry into certain questions. Similarly, some sources are easier to accumulate to create an archive. Note that technological decisions can determine whether sources continue to exist or be accessible (e.g., when publishers switched from acid-free paper, they shortened the lives of books; when software developers upgrade a product, they may render certain prior files unreadable).

it affect employment and prices? Will public officials monitor behavior and enforce any mandates? Moreover, in designing the policy, citizens and policy makers must determine how the interests of different generations and different groups within generations should be weighted.

Assessment of each of these issues relies on knowledge not available to most citizens. Technocracy apparently offers a solution to this problem. The democratic policy determines the goals that the society should pursue (and how they are to be traded off), and experts then identify the optimal means for achieving these goals. This approach, however, faces at least two difficulties.

First, setting goals may require specialized knowledge. To set goals, an agent must determine what interests and values are at stake and then integrate them into an “all things considered” judgment. Often, this process of integration must resolve complex ethical issues, such as how future generations should be treated or how to integrate the interests of distinct individuals.

Second, experts do not typically agree on the optimal means to achieve any social goal. In terms of climate change, for example, although there is general agreement among scientists that global warming is occurring, there is disagreement about the rate of warming, the consequences of warming, and the best way to respond. Moreover, various interest groups in society will have vested interests in pursuing different remedial strategies. These interest groups will attempt to influence the technocrats to adopt strategies that the group favors. An “informed” citizenry might be needed to resolve these controversies and to monitor the conduct of the technocrats.

Democratic education must therefore determine what basic instruction it needs to give to each of its citizens. In the nineteenth century, reading, writing, and arithmetic might have sufficed. In the twenty-first century, however, the scale and complexity of technology as well as the policy challenges of large, modern societies require an education that trains citizens to engage in moderate policy discussions that include controversial questions well outside of the citizen’s domain of “active” knowledge; that is, the knowledge on which a citizen can intelligently act.

### Concluding Remarks

The understanding of deliberate ignorance rests on an understanding of cognition and the mind. Each understanding points to different meanings of deliberate ignorance, justifications for it, and methods for creating it.

On the classical conception of individual mind, knowledge resides exclusively in the brain of the knower. Even in this setting, deliberate ignorance can be understood in at least four different ways. “Ignorance” might refer to a lack of knowledge or to a disregard of knowledge possessed. “Deliberate” might refer to a direct or indirect intention; the agent, that is, might intentionally choose not to acquire the relevant knowledge, *K*, or might intentionally choose

to learn  $K'$  rather than  $K$ . Right reason accounts of rationality endorse deliberate ignorance as a direct choice not to know  $K$ ; deliberate ignorance serves to filter out information that may undermine the operation of rational faculties. Deliberate ignorance here requires effort to conceal or remove the knowledge, as in double-blind studies.

More strongly, deliberate ignorance may entail simply the disregard of information that the agent has. This sense of deliberate ignorance underlies some strategies recommended by ecological rationality. Deliberate ignorance here requires only averting one's eyes and is justified on the basis of its success. More strongly still, deliberate ignorance might be understood as endorsing a belief known to be false. Formal models in the natural and social sciences follow this explanatory strategy. Whether and how this practice explains the phenomenon, however, remains open.

When one understands cognition as extending beyond the brain of the individual, the understanding of deliberate ignorance shifts from what is known or not known to how knowledge is distributed across individuals. This distribution of knowledge presents a number of problems. In the context of formally organized groups, deliberate ignorance may result from the division of labor within the organization so that no individual has the relevant active knowledge to decide well (either rationally or morally). In democratic societies, this poses questions of how to structure education so that citizens can participate intelligently in the formulation of policy.

### **Acknowledgments**

The financial support of the Filomen d'Agostino and Max E. Greenberg Research Fund of the NYU School of Law is acknowledged. Comments on an earlier draft from Richard Brooks, Ralph Hertwig, Liam Murphy, and two anonymous referees greatly improved the discussion, as did the extensive discussion at the Forum.