



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

Normative Implications of Deliberate Ignorance

Joachim I. Krueger, Ulrike Hahn, Dagmar Ellerbrock,
Simon Gächter, Ralph Hertwig, Lewis A. Kornhauser,
Christina Leuker, Nora Szech, and Michael R. Waldmann

Abstract

In this chapter the phenomenon of *deliberate ignorance* is submitted to a normative analysis. Going beyond definitions and taxonomies, normative frameworks allow us to analyze the implications of individual and collective choices for ignorance across various contexts. This chapter outlines first steps toward such an analysis. Starting with the claim that deliberate ignorance is categorically bad by the lights of morality and rationality, a suite of criteria is considered that afford a more nuanced understanding and identify challenges for future research.

Introduction

“The game is all taped. Germany won,” announced Grandma Harriet as she deliberately thwarted JIK’s attempt to remain deliberately ignorant of the game’s outcome during the 1998 Football World Cup, so that he could simulate a live experience later through the taped footage. Indeed, Germany beat Mexico 2:1, but the experience of watching the game live, with the associated mounting tension, surprise, and elation, was completely undone by her pronouncement.

The pursuit of knowledge is a fundamental mandate in Western philosophy and the sciences that have grown from it. The Socratic paradox, “I know that I know nothing,” reflects the idea that true knowledge, though difficult to attain, must be sought. Francis Bacon equated knowledge with power, in the

Group photos (top left to bottom right) Joachim Krueger, Ulrike Hahn, Dagmar Ellerbrock, Lewis Kornhauser, Nora Szech, Ralph Hertwig, Simon Gächter, Christina Leuker, Michael Waldmann, Joachim Krueger, Lewis Kornhauser, Ralph Hertwig, Ulrike Hahn, Nora Szech, Dagmar Ellerbrock, Christina Leuker, Simon Gächter, Joachim Krueger and Dagmar Ellerbrock, Ralph Hertwig, Nora Szech, Michael Waldmann

sense that knowledge has the epistemic power to change common assumptions. From this perspective, knowledge is good, and ignorance must therefore be bad—deliberate ignorance doubly so. During the Age of Enlightenment, the attainment of knowledge became accepted as a core value. Philosophers sought to emancipate knowledge from the shackles of religion, thereby asserting the human capacity and obligation to seek understanding. Kant (1784), for example, held that ignorance hinders rational reflection and thus compromises ethical behavior. Deliberate ignorance, therefore, is incompatible with the spirit of the Enlightenment.¹

Yet, many pre- and postmodern social systems embrace or even demand ignorance and its deliberate cultivation. In certain instances, they demand it, for instance, by imposing taboos or limiting the flow of information, presumably to maintain social stability by preventing individuals from gathering information that could be dangerous for the collective or the ruling class (Simmel 1906). Religions, like all systems of social control, make ample use of information-limiting taboos, which are often conveyed as cautionary tales. In everyday life, there are countless reasons for cultivating deliberate ignorance, as exemplified by the example in the epigraph: it is reasonable to want to preserve the suspense of an action by delaying outcome knowledge, yet this can easily be thwarted (Ely et al. 2015).

Recently, scholars from various disciplines have begun to explore the conditions under which deliberate ignorance occurs and may be defensible on moral, psychological, or rational grounds (Gigerenzer and Garcia-Retamero 2017; Golman et al. 2017; Hertwig and Engel, this volume, 2016). There is a certain educational irony to this project, as these scholars appear to be saying “we shall not remain ignorant about deliberate ignorance.” Then again, this project is not entirely new, but rather a re-enlivenment of an ancient one. The pre-Socratic Greeks understood the potential dangers of too much knowledge and the wisdom of carefully chosen ignorance. Aeschylus described Cassandra as cursed with painful foresight, and Prometheus, after granting humans knowledge of the future, was forced by stronger gods to leave us with “blind hope.” Whereas the accretion of knowledge is deemed beneficial in most contexts, the prospect of omniscience is unsettling. In his story of *The Golden Man*, Philip Dick (1980) gave us a thought experiment on the consequences of perfect foreknowledge. The golden man (a mutant in a postapocalyptic world) can see the future perfectly, including his own actions and those of others. When ordinary humans temporarily catch and examine him (which he knew would happen), they find his frontal lobes atrophied. His sole psychological capacity is perception, rendering cognition superfluous. *The Golden Man* is a cautionary tale,

¹ The present exploration proceeds along post-Kantian lines in the contemporary Western context; however, different understandings of deliberate ignorance hold in traditional, postmodern, or indigenous frameworks (e.g., Fahrmeir and Imhausen 2013; Fried and Stolleis 2009; Joas 1997).

illustrating the dangers of having too much of a good thing, be it food, sex, or information. But how do humans know when to forego knowledge that is freely available? Answers to this question will shed light not only on the moral and rational status of deliberate ignorance, but on its wisdom.

Norms reflect whether certain behaviors, actions, choices, or procedures are defensible (i.e., permissible or explainable) or even mandated (i.e., expected or prescribed). Norms, and the expectations they raise, provide a backdrop against which human behavior can be evaluated and are of great importance in the social sciences (e.g., psychology, economics) and related fields (e.g., philosophy, legal theory) (Popitz 1980; Schäfers 2010). There are fundamental distinctions among norms of morality, norms of rationality, social norms, legal norms, cultural norms, religious norms, aesthetic and other norms across different fields of normative investigation (Möllers 2015). Here, we focus on norms of morality and norms of rationality without claiming this initial analysis to be exhaustive. Before beginning our exploration into the moral and rational dimensions of deliberate ignorance, we make some clarifications and disclaimers.

First, we focus on cases in which deliberate ignorance does not appear to be unambiguously good or bad in the moral or the rational sense. For this reason, we do not question the rationality and ethical legitimacy of, for instance, the blind audition policy used by major classical orchestras. This policy is a normative advance and has helped to increase female musicians' access to the most selective orchestras (Goldin and Rouse 2000). Once it has been determined that specific knowledge has a biasing effect, the question is no longer whether deliberate ignorance is moral or rational but whether the desire to obtain the biasing knowledge is immoral or irrational. We are interested in cases in which moral and rational dimensions collide. In some strategic contexts, for example, deliberate (strategic) ignorance can be rationalized along game-theoretic lines but remains ethically problematic (Schelling 1956). According to legal norms, knowledge of a crime or of the perpetrator's identity entails an obligation to intervene or to testify against the perpetrator. A strategic approach might be to sidestep this obligation by remaining deliberately ignorant: someone without the relevant knowledge cannot, arguably, be blamed for failing to intervene or testify. Deliberate ignorance thus enables rationalizations designed to avoid costly involvement. This type of strategic rationality does not necessarily meet moral criteria.

Second, and following the prevailing approach in Western philosophy and psychology, we see rationality and morality as separable domains; in other words, we assume that the norms applicable in one domain do not reduce to the norms applicable in the other (Fiske 2018; Heck and Krueger 2017). Likewise, we focus on psychological theories and paradigms anchored in methodological individualism; that is, the view that norms of behavior and choice operate at the level of individual agents. We acknowledge that other disciplines emphasize the social construction of norms and knowledge (Porter 2005; Schütz 1993; Wren 1990). We consider the social dimension of deliberate ignorance

wherever possible, but leave a full interdisciplinary treatment to others. We also touch on questions of changing moral norms and expectations as may be seen across political systems, historical periods, and socioeconomic settings (Joas 1997).

Third, related to the previous issue, many instructive case studies of deliberate ignorance (see Appendix 14.1) involve individual decision makers who bear the consequences of their own decisions. Clearly, however, the consequences of an individual's choice for or against deliberate ignorance often extend to other individuals and larger aggregates. Consider an example of someone who knows they are deep in credit card debt, but does not know whether they owe a large or a very large amount. Suppose they will never be able to service the debt, but would feel much worse if they knew the debt is really very large. Would deciding not to find out the size of the debt be normative in the sense of morally permissible? Who else might be affected by this decision? Relatedly, there is the difficult question of how to conceptualize a decision maker that is not an individual (see also Kornhauser, this volume), but a supra-individual.

Another important dimension in choice is time. How may deliberate ignorance be evaluated when a person's preferences change, and when a person—like Odysseus—predicts different mental states and preferences within himself in future contexts (Duckworth et al. 2016; Elster 2000)? The Odysseus of myth happened to be correct in his predictions about his future preferences. Other mortals may not be so wise or so lucky. Affective forecasting is fraught with prediction errors, chief among them the inability to appreciate one's ability to adapt to highly emotional events (Wilson and Gilbert 2005). We also consider variation over contexts. What is normative in one context may not be in another. We hypothesize that preference spaces can be “fractured” instead of held together by a unitary self. These complexities will make it harder to evaluate deliberate ignorance, but they will also afford a more nuanced understanding.

Finally, we consider the role of deliberate ignorance in the context of our own craft (see also MacCoun, this volume): research and scholarship. At the social and societal level, government policies and cultural values tend to bias the direction of research. Certain projects are favored while others languish for lack of funding (Lander et al. 2019). Such biases implicate the actions of a collective agent, be it represented by a single person or a small group of powerful decision makers who choose deliberate ignorance with potentially far-reaching consequences for societies, local groups, and individuals. Scientific communities and individual scientists manifest deliberate ignorance when investigating some problems while neglecting others. The increasing availability of huge data sets creates ever more situations in which scientists can choose not to look at specific data and not to ask particular questions. By the same token, scientists gathering data over time face the decision of whether to perform sequential analyses. Current conventions discourage such statistical

previews, reflecting a norm designed to protect the scientific community from “p-hacking,” the practice of testing data periodically and stopping data collection once a desired significance threshold has been reached (Simmons et al. 2011). In other words, these norms positively demand deliberate ignorance from the researcher.

Taxonomic Issues: The Individual and the Collective

Normative analysis can be facilitated by, first, clarifying who chooses or rejects deliberate ignorance (who has agency) and who is affected by its consequences (who is impacted) and, second, distinguishing between the individual and the collective level, where the latter is the more ambiguous (Passoth et al. 2012). A collective may be a small group of interacting individuals or a large aggregate, crowd, network, or society. These distinctions suggest the two-by-two taxonomy shown in Figure 14.1.

We depart from the pragmatic starting point of an individual choosing or rejecting deliberate ignorance whose consequences are limited to the self (top left quadrant). In reality, of course, pure cases of this kind are likely to be few and far between. People are embedded in social networks, communities, and cultures, and whatever affects them tends to affect others as well. But let us set aside small and unintended side effects on others for the sake of this analysis and consider the case of genetic testing for an incurable disease. The paradigmatic case here is a person, Ian, who has been tested, perhaps due to a family history of the disease, and can now pick up the results. Assuming that no protective or preventive action can be taken, that no further reproductive choices

		Impact	
		<i>Individual</i>	<i>Collective</i>
Agency	<i>Individual</i>	<ul style="list-style-type: none"> • Genetic testing • Hedonic self-management (e.g., suspense regulation) • Not viewing one’s own Stasi file 	<ul style="list-style-type: none"> • Testing for communicable diseases (e.g., HIV) • Not wanting to know what others think (e.g., teaching evaluations) • AntiVax
	<i>Collective</i>	<ul style="list-style-type: none"> • Conducting blind peer reviews, blind auditions • Granting amnesty to junta members • Public and private agencies not accessing prominent individuals’ Stasi files 	<ul style="list-style-type: none"> • Sins of the past: Deciding not to find out about the acts of a previous regime • Not collecting data on certain personal characteristics (e.g., via social media)

Figure 14.1 A heuristic taxonomy of types of deliberate ignorance by agency/impact and individual collective levels.

will be made, and that the negative emotional response to a positive test result will be stronger than the positive emotional response to a negative test result, Ian may opt for deliberate ignorance. In so doing, he reveals a preference for a state of uncertainty over a state of near certainty. Yet how will his loved ones feel about this (Figure 14.1, top right quadrant)? How might Ian take their expected responses into account? Will Ian be correct in predicting not only his own responses but also the responses of others?

This example suggests that when the agent is an individual, normative analysis can, and perhaps should, consider both the individual-to-individual scenario and the individual-to-collective scenario. The outcome of the analysis may differ: a strict within-person scenario will return a judgment that deliberate ignorance is (non)normative but an individual-to-collective scenario will return the opposite conclusion.

Considering a collective as the agent (Figure 14.1, bottom two quadrants) raises a conceptually thornier problem. A collective could be a governmental office (e.g., the U.S. Environmental Protection Agency) or an individual representing a larger body (e.g., a government official or a CEO). One pragmatic option is to look at a collective *as if* it were a person. Yet, when considering collective agents such as co-acting crowds, one should be cognizant of the group-mind fallacy; that is, the idea that groups have minds that are functionally equivalent to individual minds.² The next question is how to distinguish cases in which an individual is affected from cases in which a collective is affected. When a collective agent chooses or rejects deliberate ignorance, many individuals are typically affected. A particular set of circumstances may be required for the decision of a collective agent to affect a single individual. For instance, how might a journal's decision to switch to blind peer review affect an individual academic? How does a government's decision to grant amnesty affect a member of a now dissolved junta? And how does it affect the victims? A full analysis of the normative status of deliberate ignorance requires a review of its projected effects on a class of individuals.

The taxonomic distinction between individuals and collectives matters because dissociations can swiftly emerge. For example, the same individual act of deliberate ignorance might be normatively defensible if the consequences are limited to the individual, but non-normative if they affect the collective (see Appendix 14.1). Furthermore, the individual–collective distinction intersects with the issue of social dilemmas. In a social dilemma, individuals choose strategies, such as cooperation or defection, but they cannot determine final outcomes, which are co-determined by the choices of others (Dawes 1980; Fischbacher and Gächter 2010). The decision to engage in deliberate ignorance may amount to cooperation or competition, and it may benefit or harm the self or others. In modified dictator games, for example, many participants fail to

² For arguments in favor of the group-mind concept, see Allport (1924) and Ellwood (1920); for arguments against it, see Krueger et al. (2006).

access information about their partners' payoffs if this deliberate ignorance allows them to claim more for themselves while maintaining a self-image of fair-mindedness (Dana et al. 2007). In trust games, by contrast, many players fail to inspect the payoffs of others and thus forego opportunities to protect themselves from betrayal (Evans and Krueger 2011). In prisoner's dilemmas, and other "noncooperative games," ignorance of the strategies of others enhances cooperation to the extent that players project their own preferences onto these others (Krueger 2013). Here, deliberate ignorance yields material benefits to the individual player as well as to the collective. If players learned the actual strategies of others, their incentives to defect would be greater and the sum of the payoffs would be smaller. In "anti-coordination games" such as chicken (Rapoport and Chammah 1966) or the volunteer's dilemma (Krueger 2019), knowledge of others' strategies is beneficial and deliberate ignorance is detrimental. The player who knows the other's move can select the best response, whereas the ignorant player must resort to the probabilistic and less efficient strategy of betting on the probabilistic Nash equilibrium. Social dilemmas such as these, as well as other manifestations of strategic decision making, complicate the normative evaluation of deliberative ignorance.

Norms of Morality and Rationality

Whereas some philosophical schools see morality and rationality as closely related, others note important distinctions, by highlighting cases of behavior that are both rational and immoral (Adorno and Horkheimer 2002; Dawes 1988a; Krueger and Massey 2009). A comprehensive review must consider all intersections between (ir)rationality and (im)morality. Much as most contemporary scholarship maintains that morality and rationality are distinct, if related, domains, ordinary social perception tends to map them onto a two-dimensional space, with judgments of morality and rationality being separable and, under certain conditions, orthogonal (Abele and Wojciszke 2014; Fiske 2018) or even compensatory; that is, negatively related (Kervyn et al. 2010).

Questions of morality can be approached from two major perspectives: deontology and consequentialism. From a deontological perspective, where the focus is on the morality of an action rather than on its consequences, deliberate ignorance seems wrong. Akin to lying, which according to Kant is always wrong, not taking potentially useful information into account is also wrong. Much like there is a moral duty to speak the truth, there is a duty to access relevant and accessible information. In contrast, consequentialism focuses on the potential outcomes of an action or failure to act and asks whether deliberate ignorance increases total happiness. Although consequentialism and deontology use different axioms, there are attempts at reconciliation (Hooker 2000; Parfit 2013/2017). For example, if from a deontological perspective the principles of nonmaleficence and autonomy need to be traded off (e.g., when deciding

whether or not to undergo genetic testing), consequentialist considerations may help to make such tradeoffs in a coherent and principled way.

As to rationality, the major perspectives address the coherence, correspondence, and functionality of judgments and choices. Coherence rationality asks whether deliberate ignorance introduces contradictions within the belief systems of individuals or collectives (Dawes 1988b; Krueger 2012); correspondence rationality asks whether deliberate ignorance threatens the accuracy of people's beliefs (Hammond 2000); and functional rationality asks about threats to an individual's or a collective's ultimate interests and goals, such as survival and reproduction (Haselton et al. 2009). These types of rationality are neither mutually exclusive nor do they entail one another (Arkes et al. 2016).

If the immorality and irrationality of deliberate ignorance were foregone conclusions, we would simply need to work out how the major normative frameworks justify this conclusion. Yet, empirical cases cast doubt on the idea that deliberate ignorance is *necessarily* irrational. Returning to the example of Ian, let us suppose that he decided not to pick up the test results, which would reveal whether or not he carries the gene for Huntington disease. One might wonder why he took the test in the first place. Does Ian's behavior indicate a reversal of preference and, if so, might such a reversal be regarded as irrational? What if the test results become incidentally available? Consulting them would create a state of knowledge even if there is no necessary call for action. The assumption of no implications for action is crucial here because testing positive might affect reproductive choices (Oster et al. 2013). Implications for action are just one set of consequences. Another set is affective. Ian might choose deliberate ignorance because a positive test result would take a heavy emotional toll, particularly as there is nothing he can do to avoid or mitigate the onset of the disease (Schweizer and Szech 2018). When deliberate ignorance is deployed in the service of anticipatory regret regulation (Ellerbrock and Hertwig, this volume; Gigerenzer and Garcia-Retamero 2017), it may be adaptive and thus rational in the functional sense.

From a moral perspective, knowledge of a positive test result might take an emotional toll on Ian's loved ones. Is there a moral obligation to anticipate these negative emotions and prevent emotional harm by not getting tested or keeping the results secret? Do family members have a moral obligation to get tested as well so that they can respond in the best interest of their families? These are difficult questions, in part because a *right not to know* has been asserted in the field of genetic testing (Berkman, this volume; Wehling 2019).

Moral Principles

Is deliberate ignorance morally good or bad, neutral, or ambivalent? How might prevalent metatheories of ethics and morality be applied to deliberate ignorance (Waldmann et al. 2012)? As mentioned above, the principal candidate

theories here are deontology and consequentialism (specifically, utilitarianism). In the context of deliberate ignorance, deontology asks whether there is a moral obligation to consider freely available information regardless of the possible consequences, or whether it is permissible to remain ignorant. Any obligation to access such information will likely require the commission of an act, whereas the permission not to access information may result in acts of omission.

From the perspective of consequentialism, the morality of retrieving or not retrieving information depends on the totality of the consequences. In all but the simplest contexts, this perspective makes unrealistic psychological demands on a person's ability to foresee future states of the world. Consequentialism must be constrained by empirically sound assumptions about psychological capacities, lest human morality is judged on the basis of unrealistic standards.

Whereas utility theories describe human welfare as resting on subjective preferences and their satisfaction, some moral philosophers propose the existence of objective goods, which individuals may not recognize or use in their decision making (Rice 2013). In addition to health, freedom, and social connectedness, such objective goods may include knowledge, and thus information. This would imply that information should be retrieved whenever possible. It seems that objective list theories of well-being, though grounded in objectively desirable consequences (e.g., having attained knowledge), can be regarded as a variant of deontology. Relatedly, a case can be made for goods that have no direct or measurable consequences for humans. A healthy ecosystem in a remote location, for example, may be considered desirable even if it has no direct effects on human consumption or happiness (Sen 1987). Failing to acquire relevant information (or banning the acquisition thereof) may then be regarded as unethical.

"Information consequentialism" can be distinguished from "action consequentialism" and "rule consequentialism." According to information consequentialism, if information can be used to advance human welfare, such information must be acquired. An interesting problem arises when those who decide on behalf of a collective ignore the collective's preferences. For example, capital punishment may be deemed categorically wrong even if a majority of the population is in favor of it. A deontological analysis must ask not only whether decision makers should ignore popular will, but also whether they are even morally obliged *not to find out what that will is*. Deliberate ignorance may be warranted because a popular vote may create unwanted pressure to pass laws that undermine human dignity, human rights, or, for example, constitutionally enshrined rights for minorities (Anter 2004; Jellinek 1898).³

³ It is important to note that a consequentialist social planner, who makes decisions for collectives or a society, may exercise deliberate ignorance because efficiency is an attractive benchmark. Thus, a planner may choose a policy solution that better satisfies individual preferences, the more important they are for individuals. This benchmark can be invoked even if these preferences are nonstandard (e.g., are decidedly other-regarding rather than self-regarding).

There is no “one size fits all” moral principle or definition of human welfare by which to judge deliberate ignorance. A tractable project would be to explore what deontology and consequentialism have to say about specific cases, thereby gaining a deeper understanding of deliberate ignorance itself as well as the conditions under which it promotes or hinders human welfare. With the geometric growth of genetic data, for example, incidental information provides opportunities for incidental knowledge, some of which has unforeseeable consequences (Berkman, this volume). Do individuals or collectives have the right not to know this information? Cast as a permission, the right not to know is deontological. From the consequentialist perspective, agents may choose to remain ignorant if they perceive negative consequences outweighing positive ones.⁴

In some contexts, the growing quantities of “potential information” yielded by technological advances present challenges to personal identity. Algorithms that harvest social media data can predict sexual, religious, or financial preferences more accurately than friends and peers (Kosinski et al. 2015). Here, does an individual also have a right not to know to shield the individual from discovering potentially unsettling and formerly hidden sides of the self? What is the normative force of the Socratic injunction to know thyself? More generally, do people have the right not to know things about themselves that they either do not want to know or that challenge their beliefs? It is unclear whether the greater accuracy of such big-data inferences about personal traits limits the individual right not to know, but it would be odd if it did not. Arguably, the right not to know becomes *stronger* with the accuracy of algorithmic prediction, thereby bringing morality into greater conflict with certain types of rationality.

Moral Principles and Moral Intuitions

One important and rich research question in the context of deliberate ignorance concerns if, when, and why moral principles and moral folk intuitions conform or fail to conform (see also Heck and Meyer 2019). For instance, some consumers fail to obtain information on the conditions under which goods were produced, knowing that it might reveal practices such as child labor or factory farming. Research has shown that such “willfully ignorant” consumers also denigrate others who seek such information (Zane et al. 2016). This

Outcomes, however, may become repugnant if this efficiency-based approach permits any kind of preferences. Take, for instance, “happiness.” The satisfaction of happiness can justify nearly any kind of intervention. To avoid this problem, policy makers may choose not to find out about specific preferences or turn a blind eye if they happen to know or anticipate them (for more on this “laundering” of policy-relevant preferences by public policy makers, see Bierbrauer, this volume).

⁴ Interestingly, observers tend to regard those who engage in medical deliberate ignorance as less moral than those who do not, regardless of whether the information is actionable and regardless of information valence (Heck and Meyer 2019).

suggests that moral intuitive judgment is shaped by an individual's own choice of deliberate ignorance. However, principles and intuitions can also conform, while behavior diverges. Serra-Garcia and Szech (2018) experimentally created moral wiggle room by giving participants a choice between \$2.50 in cash and a sealed envelope that potentially contained \$10 for a worthy cause. Many respondents left the envelope unopened and took the money. This sort of self-serving reasoning is condemned by both normative models and folk judgment (Kunda 1990).

Moral judgment may be sensitive to past behavior that has no differential impact on future consequences or mental states. Let us return to the case of Ian who has undergone genetic testing but chooses not to pick up the results. Now consider an otherwise identical person, Niall, who never got tested. In this case, Ian has moved closer to obtaining the test results by having had, at some point, the goal to find out his genetic status. He is thus likely to be judged more harshly because it is easier to imagine him obtaining the results (Miller and Kahneman 1986). For Ian, only one event (retrieving results) needs to be changed (reversed); for Niall, two events (getting tested and retrieving results) are involved. In other words, reversing a previous decision against deliberate ignorance may seem particularly blameworthy. More generally, the distinction between acts and omissions is important both in formal deontological models (Callahan 1989) and in folk judgment. A person may achieve a state of deliberate ignorance by acting to block information or by omitting to retrieve it. In general, when the likely (and identical) consequences are negative, acts are evaluated more harshly than omissions, and presumably so when intentions are held constant (Haidt and Baron 1996).

Many decisions are made with goals in mind (Higgins 1997). The actors intend to produce certain consequences, and these intentions are also relevant to moral judgment (Malle et al. 2014). Someone who elects not to retrieve medical test results or to discover hidden sides of their personality (from their social media footprint) to avoid emotional distress may be judged less harshly than someone who remains ignorant to keep others in the dark (although informing others may involve a second and separate decision or the anticipation of emotional leakage that would reveal the result).

These questions suggest that there is a rich set of issues that pertain to moral folk intuitions about deliberate ignorance. In the wild, moral judgment may shift from deontological to consequentialist concern without any change in the consequences, as demonstrated by the trolley problem (Greene 2016). Why these shifts occur is a psychological question. One interesting argument is that the causal framing of the scenario determines the moral lens through which it is seen (Waldmann and Dieterich 2007). As shown by the different cases listed in the Appendix 14.1, both deontological and consequentialist perspectives are instructive for a normative and psychological understanding of when (not) and why normative benchmarks and intuitive judgments consider it to be ethical or unethical.

Rational Principles

The Standard Model of Rational Choice and Its Shortcomings

The standard model of rational choice⁵ places Bayesian inference in the service of expected utility maximization. In this model, observations, or data, reduce uncertainty, or at least do not increase it (Good 1950; Oaksford and Chater 2007; Savage 1972). Thus, there is no rationale for not seeking or not using easily available information, particularly information with large potential benefits (Howard 1966; for a critical discussion, see Crupi et al. 2018). Indeed, inferences tend to become more accurate as more information is used unless that information is systematically biased. Statistical hypothesis testing recognizes the value of observations such that larger samples are more likely to yield “significant” results if there is indeed an effect (Krueger and Heck 2017). Why then should people sample less data when the costs of sampling are negligible? Bayesian and frequentist (significance testing) models of statistical evaluation often become metaphors of mind (Gigerenzer and Murray 1987). The assumption is that lay people reason (or rather should reason) much like scientists (Nisbett and Ross 1980), preferring more information over less, lest their inferences suffer (Tversky and Kahneman 1971).

Descriptive and experimental studies on human reasoning show many departures from this rational ideal. Selectivity in information search is a common finding (Fischer and Greitemeyer 2010), and this selectivity is often interpreted to be part of a motivated bias to confirm existing beliefs or expectations (Nickerson 1998). The experience of disappointing initial observations can also prompt people to truncate information search (Denrell 2005; Prager et al. 2018). Cast as a variant of selective information search, deliberate ignorance thus appears to be irrational from the perspective of the standard model. Yet, inferences based on limited information can also be superior to inferences based on full information, thus making simplistic verdicts of irrationality problematic. We turn now to some examples of this counterintuitive finding (from the point of view of the standard model) and consider implications for deliberate ignorance, before discussing other shortcomings of the standard model.

When less information yields better inferences. Kareev et al. (1997) discovered that small samples are particularly sensitive to true correlations because they are likely to amplify them (but see Juslin and Olsson 2005). Perceivers with low working-memory capacity can thus be more accurate in detecting a true correlation than perceivers with a large capacity. The latter might therefore

⁵ Here the “standard model” of rational choice represents the dominant view in economics, psychology, and philosophy. In the humanities, by contrast, there is greater emphasis on the cultural relativity of epistemic practices and thus a greater reluctance to ascribe general normative force to any particular framework.

choose deliberate ignorance to do as well as the former. In a context that allows strategic behavior, Kareev and Avrahami (2007) found that when employers rewarded good performance with a bonus and without monitoring workers closely (thus exercising deliberate ignorance), both strong and weak workers exerted more effort, which improved the accuracy of assessment and increased productivity. Exploring the conditions under which small samples benefit inference can have far-reaching implications for the rationality of deliberate ignorance and the limitations of standard assumptions (Fiedler and Juslin 2006; Hahn 2014). Note that the sample information must be valid in the sense that it comprises observations from the latent population that the decision maker seeks to understand. In other words, it is critical to distinguish between information that is essentially relevant and information that is not.

Negative affect can restore rationality. Affect regulation is a variant of motivated reasoning. Individuals may choose deliberate ignorance if they suspect that retrieved information would be distressing. Many people see keeping up with news and world events as a moral obligation, but chafe under the relentlessly negative focus and tone of the coverage. Moreover, most news stories do not call viewers to action (Dobelli 2013). Hence, there is a case for rationally curtailing one's news intake. Foregone knowledge need not worsen a person's epistemic state, but the anticipated negative effect of bad but unactionable news may be factored into their expected and experienced utility. Likewise, individuals who choose not to read files compiled on them by the Stasi (the secret police of the former East German government) may be engaging in rational affect management (see Appendix 14.1 as well as Ellerbrock and Hertwig, this volume). Some of the negative affect triggered by avoidable information can have epistemic significance. The information may raise more questions than it answers, thereby deepening the person's unpleasant feeling of ignorance. In other words, actual and experienced ignorance can be inversely related. Bringing this dissociation to light was the devilish charm of the Socratic method. This dynamic has recently seen a renaissance in research on the so-called illusion of explanatory depth (Rozenblit and Keil 2002; Vitriol and Marsh 2018). In learning more, people come to realize how little they know, possibly blunting the mood at least for the epistemically ambitious. Choosing deliberate ignorance early on can protect them from this experience and thus be a strategy of effective emotion regulation.

In the standard model, however, Bayesian expectations about the possible outcomes of obtaining information already factor in all outcomes, including affective ones (Weiss 2005). A person should be neutral toward receiving information that requires no action and instead seek instrumental information (e.g., medical testing may have actionable implications for getting treatment, making career choices, family planning, and revising saving plans). In short, the possibility that information retrieval will cause negative affect fails to justify deliberate ignorance from the perspective of the standard model of rational

choice. In this context, it is interesting to note that humans with lesions to a range of brain regions implicated in the processing of emotions counterintuitively are more likely to conform to norms of the standard model than “normal” individuals (Hertwig and Volz 2013).

The role of transformative and disruptive information. In some extreme cases of deliberate ignorance, information is ignored because it might upset or transform a person’s set of preferences. This possibility applies to Ian and the genetic test for Huntington disease, or to anyone with reason to fear that their spouse was a Stasi informant (Ellerbrock and Hertwig, this volume). It is difficult to predict or appreciate how much damage receiving unwelcome knowledge of this type might do to a person’s inner world. It is thus also difficult to see how adequate decisions can be made on the basis of current preferences alone. In other words, these cases of “transformative experience” and decisions have been suggested to profoundly challenge the standard model of rational choice (Paul 2014), but these challenges are difficult to model quantitatively.

Deliberate ignorance can be an adaptation to prohibitive information costs. In the modern world, deliberate ignorance may occur because environmental change outpaces the evolution of the human mind (Higginson et al. 2012). Ancestral environments afforded certain costly learning opportunities that are no longer pervasive or relevant. An early hominid could learn to discriminate between lethal and harmless species of snakes and spiders by being bitten (i.e., by receiving costly and dangerous information). The relevant information motivating avoidance of all snakes and spiders is now efficiently acquired through cultural transmission (Larrick and Feiler 2015). The modern world offers opportunities to approach snakes and spiders in safe environments (e.g., zoos), yet many people respond like early hominids, showing strong aversion to such creatures, even behind glass. Although the cost of obtaining information about the animals is now low, many people opt for deliberate ignorance. In this sense, some modern manifestations of deliberate ignorance can be seen as being rooted in adaptations to risky worlds that no longer exist.

*The Standard Model of Rational Choice and Deliberate Ignorance:
What Gives?*

It would seem rash to consider the case for rational deliberate ignorance closed for at least two reasons. First, some of the points raised above suggest that decision theory is too limited. Second, it seems odd that the standard model of rational choice, by disregarding human constraints, pronounces even behaviors where people “do the best they can” as irrational. Constraint-sensitive notions of bounded rationality are unlikely to replace

the unbounded standard, but we wish to retain the prospect of normative guidance in an area where it is *possible* for the agent to act. Potential solutions to such problems may lend themselves to more differentiated treatments of deliberate ignorance once an adequate formal machinery is in place. Therefore, it appears worthwhile to examine extensions to decision theories as well as their ability to provide a subtler normative treatment of deliberate ignorance.

Bounded and Ecological Rationality and Bounded Optimality

The standard model of rational choice fails to take search and acquisition costs or processing costs into account. Extensions to the model that respond to these limitations have interesting implications for deliberate ignorance. For example, theories of bounded (Simon 1997) or ecological rationality (Gigerenzer and Selten 2001; Hertwig and Herzog 2009) map evolved mental capacities, ask what types of judgment or decision task they will be able to contend with, and posit parsimonious rules or heuristics as a solution. This research paradigm shows that a surplus of information or cues raises the danger of overfitted forecasts. Adding valid information (e.g., parameters) increases the fit between a model and the data used to build it. Future data bring new uncertainty, however, not only from random sources but also from systematic ones, such as features of the environment or agents' preferences. New factors that are unable to be foreseen cannot be part of the model. A model that uses just a few valid predictors is likely to be more robust; that is, it will perform better than a fully parameterized model (Dana 2008; Dawes 1979). Being willing and able to deliberately ignore information to avoid overfitting is the ecological decision maker's secret weapon (Katsikopoulos et al. 2010).

Theories of bounded optimality take an alternative approach: the inferential "optimizing" of the standard model is retained but viewed as operating within a set of limitations or bounds. For example, the ideal observer analysis (Geisler 1989, 2011) attempts to understand human capacities by comparing them against those of an ideal agent. Where discrepancies are found, the ideal agent is given realistic constraints (e.g., on the nature of the input) until close alignment is achieved. In other words, optimal rationality is reduced to the mind's realistic boundedness. The intended result is an understanding of the mechanisms and processes designed to achieve what is possible given the available constraints (Griffiths et al. 2015; Howes et al. 2009; Lieder and Griffiths 2019).

The Attention Economy and the Strategic Rationality of Deliberate Ignorance

Consumers of goods and services, much like consumers of news, must contend with the growing power of online actors to misdirect their attention. Simon (1971:40–41), the father of bounded rationality, anticipated an "information-rich world," a dystopia in which a "wealth of information creates

a poverty of attention and a need to allocate that attention efficiently among the overabundance of information that might consume it.” Compared with today’s reality, Simon’s vision seems quaint. As a psychological resource, attention is more precious than ever, and deliberate ignorance can help preserve it. Many agents (e.g., companies, advertisers, media, and policy makers) design “hyperpalatable mental stimuli” to engineer preferences and erode autonomy (Crawford 2015). In the same way that obesogenic environments are replete with food products designed to hit the consumers’ bliss point (i.e., the concentration of sugar, fat, or salt at which sensory pleasure is maximized), informationally fattening environments reduce consumers’ control over their information intake.

Many human-interest stories are attention bait masquerading as information. They have opportunity costs that people often fail to notice. In this kind of information ecology, deliberate ignorance can support individuals’ agency and autonomy; it may even qualify as a psychological competence of rational decision making. This task remains admittedly difficult as the methods for the cultural production of ignorance evolve, a development studied under the label *agnotology* (Proctor and Schiebinger 2008). Consider, for instance, the phenomenon of “flooding” in the coverage of news. On August 4, 2014, an earthquake in China’s Yunnan province killed hundreds and injured thousands. Within hours of the earthquake, Chinese media were saturated with coverage of an Internet celebrity’s alleged confession of having engaged in gambling and prostitution. News of the earthquake was thus not censored but rather crowded out. The flooding of the media with reports of a trivial scandal reflected a concerted government effort to distract the public from the devastating effects of the earthquake, as objective coverage would have revealed severe weaknesses in the government’s readiness for and response to natural disasters (King et al. 2017; Roberts 2018).

In treating attention as an essentially unlimited resource, the standard model of rational choice overlooks the dangers of the attention economy. This blind spot, perhaps ironically, may be seen as an element of deliberate ignorance built into the standard model itself. It remains to be seen how such costs or indeed errors resulting from junk stimulation can be modeled. Adding capacity constraints to the model need not be difficult; the question is how limits in attention can best be captured. Woodford’s (2009) model, for example, includes entropy-based information costs. Of course, the aim of the standard framework is to be as simple as possible. Thus, the question is: In which contexts are additions to the framework important enough to deserve coverage?

Expanding the Machinery of the Standard Model

Several approaches seek to expand the scope of the standard model itself. Utilities, the currency of the standard model, may not be stable but state-dependent. Beliefs, classically treated as distinct from utilities, may be

partially constitutive of them (Loewenstein and Molnar 2018). People often avoid information they suspect will challenge cherished beliefs—beliefs that have high utility for them and their sense of identity (Abelson 1986; Brown and Walasek, this volume; Tetlock 2002). Other revisionist models incorporate strategic unawareness (Golman et al. 2017). Whereas the standard decision model assumes that decision makers can describe all possible contingencies, all possible actions, and all possible consequences, unawareness models relax this assumption by allowing decision makers to describe their world in terms of subsets of objectively possible contingencies/actions/consequences and allowing that awareness to change over time. If decision makers are unaware of their unawareness, behavioral predictions are not fundamentally different from those of the standard model. In particular, decision makers who are unaware of their unawareness have no incentive to choose deliberate ignorance. The question of how to model decisions under awareness of unawareness awaits further attention (see Trimmer et al., this volume).

Rational Collectives

Thus far we have focused on individual agents, partly because various strategic contexts (e.g., the choice of ignorance) can be rationalized along game theoretic lines (Schelling 1956) and thus the issue of its rationality prompts less debate. Let us now briefly consider collectives. Do they present specific challenges for an analysis of the rationality of deliberate ignorance? The question of whether organizations can be thought of as having mental states is thorny, reaching beyond the empirical realm and into the metaphysical. Lacking a compelling normative answer, researchers have investigated the conditions under which lay perceivers attribute mental states to groups or organizations (Cooley et al. 2017; Jenkins et al. 2014). It seems prudent to say that organizations or institutions should not be treated holistically or anthropomorphized, nor should they be treated as mere aggregations of individuals, where knowledge, foresight, and intentions can be attributed only to each individual separately. For some purposes, governmental branches may be assumed to have knowledge and intentions, as some philosophers have argued (Pettit 2003).

Any evaluative standard, such as correspondence or coherence, can make conflicting demands at the individual versus collective level, and thus potentially justify deliberate ignorance. In social choice (as opposed to social agency), the aggregated judgment is more important than the judgments of individuals (Paldam and Nannestad 2000). Here, information beneficial to the individual may harm the performance of the collective, even in nonstrategic contexts. Variants of Condorcet's jury theorem (e.g., Ladha 1992) show that collective accuracy (i.e., the probability that the majority vote on a binary proposition captures the true state of affairs) depends both on the mean individual accuracy of the group members as well as on their degree of independence (Hastie and Kameda 2005). The same dynamic holds for "wisdom of

the crowd” scenarios concerned with estimation for discussion on the diversity prediction theorem (see Page 2008). Here, collective error is an increasing function of individual error and a decreasing function of the diversity (variance) of the individual estimates. As a result, additional information given to individuals (e.g., through interagent communication) may be detrimental to collective accuracy even though it potentially increases individual accuracy (Goodin 2004). This is not merely a theoretical possibility; it has been demonstrated in experimental investigations (Lorenz et al. 2011; cf. Becker et al. 2017 and Jonsson et al. 2015) and simulations (Hahn et al. 2019). Collectives and individuals may place conflicting demands on information acquisition and, hence, on the rationality of deliberate ignorance. Where collective accuracy is paramount, avoiding communication (and hence information uptake) may improve performance.

Intertemporal Choice and Multiple Selves

In the standard model of rational choice, updated preferences are consistent with initial preferences, a feature known as dynamic consistency (Sprenger 2015). A dynamically consistent decision maker is indifferent between committing to a course of action conditional on receiving a later signal and taking the preferred action once the signal is received. Many experiments have shown systematic violations of dynamic consistency due to temptation, self-control problems, or updating of multiple priors. Suppose a decision maker chooses from a set of available actions, each of which is optimal in a different state of the world. Evaluating the utility derived from the chosen action relative to the utilities of the foregone alternatives may cause feelings of regret. Even if a decision maker can expand awareness of available actions at no cost, this expanded awareness can cause more potential regret. Thus, decision makers may not aspire to consider a wider range of options.

Intertemporal choices often show dynamic inconsistencies. People making choices often ignore consequences that will occur only in the distant future. This discounting of the future is particularly likely in the presence of temptation; that is, when attractive rewards in the near future have a high risk of adverse consequences in the more distant future. The pleasure of each cigarette smoked is immediate, whereas the risks of disease or untimely death are faraway and uncertain. The standard model assumes that people exponentially discount streams of utility over time such that preferences are consistent with or independent of time. The relative preference for well-being at an earlier date over a later date is thought to be the same regardless of whether the earlier of the two dates is near or remote. With exponential discounting of the future, such preferences are rational in that they are coherent. The empirical evidence, however, shows that the near future is discounted more steeply than the distant future (Ainslie and Haslam 1992). For example, when presented with a choice between doing seven hours of house cleaning on December 1 or eight hours

on December 15, most people (asked on October 1) prefer the seven hours on December 1. When faced with the same choice on December 1, most chose to put off the chore until December 15. This preference reversal is also known as present bias (Jackson and Yariv 2014). When considering trade-offs between two moments in the future, present bias puts greater weight on the earlier date as it draws closer.

In the struggle between the pursuit of short-term and long-term preference, deliberate ignorance can be detrimental in the long term. Take the example of smoking. By one estimate, each cigarette smoked reduces life expectancy by about 15 minutes or half a microlife; that is, by about the time it takes to smoke another two cigarettes (Spiegelhalter 2012). Smokers who enjoy the nicotine buzz and would rather not worry about the future health risks choose to ignore relevant information and remain trapped in a cycle of self-damaging behavior. One useful psychological perspective on this phenomenon is that of “multiple selves” (Jamison and Wegner 2010). While the present self enjoys the act of smoking or its direct physiological effects, it would nevertheless like its future self to get informed and quit smoking if long-term, detrimental consequences are probable and severe. Once this future self becomes the present self, however, it too will yield to the temptation of the present and postpone seeking information on health risks.

Deliberate ignorance is irrational when it contributes to time-inconsistent preferences and self-destructive behaviors. In the case of some medical treatments, however, deliberate ignorance may turn out to be the wise choice. Consider the case of a highly effective drug that also happens to have dreadful but highly improbable side effects. People tend to overweight such low probability and might therefore forfeit the restoration of their health. Those who elect not to review these side effects have a better prognosis (Carrillo and Mariotti 2000; Mariotti et al. 2018). Deliberate ignorance may also protect a person from the danger of certain medical interventions of low utility and a risk of overdiagnosis. The U.S. Preventive Service Task Force recommends that men 70 years of age and older should not submit to the PSA test to screen for prostate cancer: “Many men with prostate cancer never experience symptoms and, without screening, would never know that they have the disease” (Grossman et al. 2018:1901). Choosing not to take the PSA test can protect older men from the psychological harm associated with false-positive results (distress and worry) as well as from the harmful effects of invasive treatment (e.g., incontinence or erectile dysfunction).

We have explored the implications of dynamic inconsistencies for (ir)rationality, but there is also a moral dimension. From the deontological perspective, different temporal selves can lay claim to their own unique rights and obligations. Much like the present generation must grant rights (and obligations) to future generations (Gosseries 2008), present individual selves must be mindful of their own future incarnations. An analysis along consequentialist lines suggests the same conclusion.

The Case of Science

Decisions about which research projects to pursue imply decisions about which areas we, individually and as a society, want to learn about and which we wish to ignore. An exploration of the dynamics of knowledge production in science and scholarship may go beyond the strict definition of deliberate ignorance, but it is instructive with regard to the broader impacts of choosing for or against knowledge. Scientists do not gather information that is lying around ready for the taking; they operate at the interfaces of discovery, knowledge production, and knowledge construction. Hence, the core question driving the exploration of deliberate ignorance remains: How should deliberate ignorance be managed when foregoing knowledge has potentially large, though uncertain, impacts?

Here, we focus on the high- and mid-level strata at which deliberate ignorance can affect science and research in ways that may be questioned on normative grounds (for related phenomena, see Proctor and Schiebinger 2008). As a potent example, consider the decision of the U.S. federal government not to fund research on gun violence or policies that might mitigate its effects (i.e., the 1996 Dickey Amendment). This decision amounts to an attempt to keep a population of stakeholders ignorant, and it was soon criticized as a strategy to protect the gun industry (Jamieson 2013). Supporters of the policy argue that research would endanger rights guaranteed under the Second Amendment of the U.S. Constitution. Focusing on a trade-off between values, they make a deontological argument in favor of deliberate ignorance. Other examples of policy-based deliberate ignorance have a weaker deontological grounding, such as the historical suppression of research on the harmful effects of tobacco or, more recently, sugar.

There are also cases where there is a collective demand for deliberate ignorance without reference to commercial interests or competing political values. Research on the hydrogen bomb is such a case, at least in hindsight. Robert Oppenheimer, himself instrumental in the development of the atomic bomb, warned in vain against research on a hydrogen bomb. The future will tell whether contemporary research on artificial intelligence will be judged similarly. Some serious risks are currently being discussed (Tegmark 2017), and it is not clear whether deliberate ignorance will be achievable in this area. The implications of artificial intelligence outstripping human intelligence are by definition unpredictable, and we will not know until it is too late (Hawking et al. 2014). Similarly, there is room for debate on whether deliberate ignorance is advisable, ethical, and feasible in the context of biological research on deadly viruses or human cloning (i.e., not doing such research).

Considering the distinction between individual and collective agents is also instructive in the context of science. The examples presented thus far highlight collective decisions. But individual researchers may also decide not to perform particular types of work. On one hand, choosing one research

topic inevitably has opportunity costs: it takes away time from other projects. A scientist may therefore have to choose *where* to engage in deliberate ignorance, not *whether* to do so. On the other, strategic considerations come into play. A researcher may decide that a particular project cannot be done in good conscience. Yet, they may reasonably suspect that others have fewer scruples. The researcher is now caught in a prisoner's dilemma where opting for deliberate ignorance is the cooperative choice and rejecting it is an act of selfish defection.

In their definition of deliberate ignorance, Hertwig and Engel (this volume, 2016) specified that information-acquisition costs incurred are zero or negligible. Clearly, the financial costs of pursuing research are not negligible, at least for society, yet science may nevertheless be considered in this context. Contemporary Western societies have committed resources to doing science, such that the question of which issues to pursue is not determined primarily by cost. This locates the case of science within the potential space of deliberate ignorance, at least for those cases where costs play a negligible role in the choice of what to study. Some scientists may reject deliberate ignorance because they seek to enhance their reputations (Falk and Szech 2019; Loewenstein 1999). For example, many physicists and engineers saw involvement in the Manhattan project to develop the first nuclear weapons as a historic opportunity (Mårtensson-Pendrill 2006). Using the replacement logic of the prisoner's dilemma, these individual scientists could reject deliberate ignorance by arguing if they did not do the work, someone else would (Falk and Szech 2013). At the same time, involvement in a collaborative research project provides opportunities to diffuse responsibility and blame if outcomes turn out to be more damaging than desired (El Zein et al. 2019; Fischer et al. 2011; Rothenhäusler et al. 2018).

The benefits and harms of scientifically attained knowledge are not always predictable. Whether basic discoveries prove to be relevant or applicable is often a matter of time. This unpredictability is inevitable. If research outcomes and their consequences were already clear, there would be no need to do the research in the first place. Biological research may lead to medical advances, but it may also create new toxins or diseases. Uncertainty extends into the future. It is impossible to tell how future generations will evaluate what now appears to be scientific progress. As preferences can change within individuals, they also often change across generations.

This sketch of deliberate ignorance in the context of science points to larger issues beyond the scope of this preliminary exploration. Science—as a personal, group, or social project—represents a wager on an uncertain future. Much of its yield, and the decisions underlying it, will be comprehensible only in hindsight. To say that science is basically rational and morally neutral is perhaps a useful normative starting point. Once deliberate ignorance is recognized as one of the forces shaping the direction of scientific work, this set of assumptions will require continual reevaluation.

Open Questions

This exploration of the normative issues raised by deliberate ignorance began with taxonomic questions, continued with concerns about morality and rationality, and ended with questions about intertemporal choice and scientific work. Some questions found preliminary answers, others need to be addressed in future work. At this stage, it seems that established frameworks for moral and ethical judgment cover most manifestations of deliberate ignorance. The same general principles (deontology, consequentialism) that apply to other types of action (or failures to act) seem sufficient for the normative evaluation of deliberate ignorance. This does not, however, rule out the possibility that the application of those principles to deliberate ignorance may raise unique issues. Only a continued normative analysis of specific examples will tell. With respect to rationality norms, however, an extension of the standard model of rational choice seems necessary to accommodate deliberate ignorance. Some extensions already exist (see Brown and Walasek, this volume), and they suggest that some forms of deliberate ignorance are irrational only by the lights of the standard model. Continued refinements, or radical alternatives, will contribute to a nuanced assessment of this intriguing phenomenon. One particular challenge for future research is the modeling of collective decisions (see Trimmer et al., this volume).

We have repeatedly seen that a comprehensive understanding of deliberate ignorance in normative terms requires a careful analysis of specific cases, instances, and episodes. In this chapter we have been able to consider only a limited subset of cases. There is no way to tell how badly this sample might be biased. Indeed, there is as yet no sense of whether a true but latent population of deliberate ignorance cases can even exist. To stimulate further debate and analysis, we have assembled a set of cases (see Appendix 14.1) to provide broader coverage of the domain. We now conclude our exploration with a few final observations.

Expanding the Domain

We recommend a graduated consideration of cases of deliberate ignorance, as illustrated in Figure 14.2, beginning with the strict criterion of “no [use of] knowledge,” proceeding to “delaying knowledge” and “disregarding knowledge,” and ending with “negating knowledge.” This nested hierarchy offers progressively more open definitions of deliberate ignorance. It covers instances in which an agent chooses to delay accessing information or to disregard information that is already known. The broadest definition includes instances in which the agent acts on information known to be false. We are not committed to the idea that a conceptual expansion of this type is necessary. The core definition restricts deliberate ignorance to situations in which an agent chooses not to learn a knowable fact that may, in principle, offer large benefits

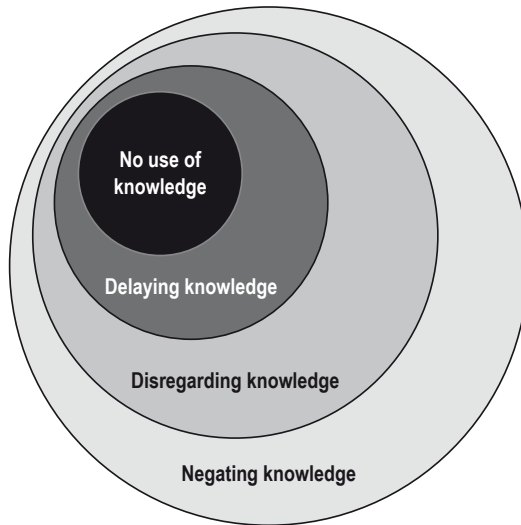


Figure 14.2 The nested structure of kinds of deliberate ignorance: no use of knowledge, delaying knowledge, disregarding knowledge, negating knowledge.

(Hertwig and Engel, this volume, 2016; see also Schwartz et al., this volume). This definition is attractive because of the sharp demarcation line it draws, thus facilitating rigorous analysis, yet it is limited in that it restricts the degree to which analytical or empirical results can be generalized. A tight definition may also run into problems in the context of collective agents when states of “ignorance” and “knowledge” are differentially distributed across members of the group, making it difficult to assess the ethical legitimacy and consequences of a collective decision to exercise deliberate ignorance.

Individual Autonomy

There are at least two interpretations of the value of autonomy and their assessments of deliberate ignorance. One account views autonomy as the free exercise of an agent’s will. From this perspective, instances of deliberate ignorance will be assessed on the basis of their consequences. Some instances may be acceptable, others not. The second account sees autonomy as the free exercise of an agent’s will as informed by all relevant reasons. In other words, full autonomy can, by definition, be exercised only if all relevant knowledge is available. This account appears to condemn every instance of deliberate ignorance because the deliberate choice not to know is, by definition, anathema to a fully autonomous individual.

Norm Conflict

In transitional or transformational societies (Ellerbrock and Hertwig, this volume), conflicting and changing norms raise difficult issues. A society can pursue, in the process of recovering from a troubled past, many different but ethically justifiable goals and principles, including repairing the social fabric and fostering justice, transparency, and peace. Such goals individually have pragmatic and moral force, but they may be mutually incompatible. On what normative or empirical grounds can a preference for one goal over others be justified? One empirical approach is to conduct field experiments to implement policies serving different goals and to measure their long-term effects on various outcome indicators (Campbell 1969; Staub 2014; Staub et al. 2005). Although this evidence-based approach is important, it does not provide a full answer to the question of how norm conflicts can be resolved: Individuals may place different weights on an outcome's indicators depending on the normative goals they prioritize. It may prove beneficial to take a closer look at how norm conflicts are approached in law (e.g., in human rights cases).

From Analysis to Empirical Testing

Exploration of the normative contexts and subtexts of deliberate ignorance is only just beginning, but some empirically tractable questions suggest themselves. To what extent, for example, would the general public agree with the various ethical principles presented here? How can research probe folk intuitions? To what extent do actual decision makers meet the (new) assumptions made in this chapter? A decision maker may still be described within a rationality framework if assumptions are added or adjusted. How then can we test whether these assumptions are descriptively accurate? And to what extent are decision makers rational in their choice of deliberate ignorance? For instance, a decision maker considering getting tested for Huntington disease may think about the consequences, or utilities, of this knowledge in narrow terms (e.g., how it will affect my health decisions) or in much broader terms (e.g., how it will affect my social, financial, and professional decisions and, by extension, my family's well-being).

Deliberate Ignorance in the Context of Political and Normative Change

Our normative analysis did not assume any specific temporal coordinates, thus acting as if norms were stable in time. This is, of course, a simplifying assumption. Norms are subject to variation across time and cultural space. This kind of normative change is salient in times of political upheaval. Given these temporal dynamics, how might deliberate ignorance be deployed to consolidate

sociopolitical change? Modern democracies tend to legitimize norms by grounding them in human rights and freedoms. The right to know and the right to have access to all information collected on oneself by governmental agencies has been central to modern statehood and notions of liberty since the French Revolution. The right *not* to know has recently complemented our understanding of modern democratic theory, although the relation between these two types of rights continues to be debated.

Future work will need to explore how instrumental deliberate ignorance can be in the development, consolidation, or erosion of social norms. Conversely, how do norms shape the practices, opportunities, and (moral) outcomes of deliberate ignorance? As the analysis of deliberate ignorance in the context of the Stasi files demonstrates (Ellerbrock and Hertwig, this volume), deliberate ignorance has an ambivalent normative quality in transformational societies: it can stabilize norms or erode them. Interestingly, the genesis and formation of normative orders has only recently emerged as the subject of historical study, emphasizing the interdependency of formation of knowledge and morals (Frevort 2019; Knoch and Möckel 2017). Deliberate ignorance is not (yet) an object of research in this context, but it is clearly relevant. In modern Western societies, normativity is negotiated and legitimized in the public discourse. Analysis of the prevalence of deliberate ignorance in the public discourse, its structure, and its role in situations of norm conflicts should therefore pique the interest of both historians and behavioral scientists. Similarly, the relatively recent reevaluation of deliberate ignorance in the context of political transformation deserves detailed analysis. Such analysis would offer new opportunities to understand when and under which conditions societies treat deliberate ignorance as an ethically legitimate or condemnable practice.

In Lieu of Closure

You are not obliged to complete the work, but neither are you free to desist from it.
—Rabbi Tarfon (Pirkei Avot, 2:16)

We have embarked on a journey toward the demystification of ignorance, and especially its deliberate variant. In Western societies, ignorance (not knowing) is associated with stigma. Babies are ignorant, and overcoming that ignorance is an essential part of growing up. How then can ignorance be a deliberate choice? We hope to have shown that some instances of deliberate ignorance are normatively defensible, but that depends on the confluence of the type of norm (moral or rational), the type of agent (individual or collective), and the type of person or group bearing the consequences. Choosing deliberate ignorance in a context in which such a choice is normatively defensible may be the mark of wisdom, and continued research efforts

are needed to enable people to choose wisely. Peeling back the veil of ignorance remains a powerful normative mandate: Overall, accepting ignorance is normatively less defensible than deliberately choosing it, but sometimes it must be chosen.

Appendix 14.1

To stimulate further debate and analysis, we have assembled ten real-life examples to illustrate the primary functions of deliberate ignorance, the actors affected, as well as the ethical (consequentialist, deontological ethics) and rationality principles that may be involved (e.g., expected utility maximization, game theory).

1) Huntington disease (also known as Huntington's chorea)

Background: This inherited autosomal disease progressively breaks down nerve cells in the brain, resulting in severe impairments in a person's ability to move, think, and reason. An affected person eventually requires help with all daily activities, although language comprehension and awareness of family or friends do not diminish. Although most people develop symptoms in their 30s or 40s, the rate of disease progression varies. Genetic testing provides reliable diagnosis at any age. Yet, even with a family history, some people deliberately choose not to take the test.

Function of deliberate ignorance: Since a positive test result augurs an early, agonizing death, individuals may choose to regulate their fear by remaining deliberately ignorant. This choice, however, also impacts others: family members will be unable to prepare for the role they may need to assume as the individual's health deteriorates or the trauma they will experience in witnessing a loved one's physical demise and early death.

Ethical principles: From a consequentialist perspective, if the test result is negative, the choice to know will undoubtedly bring about the best result for all. If, however, the test result is positive, it is not obvious whether a consequentialist perspective favors deliberate ignorance or knowledge. The emotional consequences are difficult to predict, and people generally are not very good at affective forecasting (e.g., Wilson and Gilbert 2005). Moreover, a positive test result may also have profound consequences for relatives who learn that they are likewise at risk.

Rationality principles: The choice of deliberate ignorance cannot be accommodated within an expected utility maximization or game theory framework. Additional assumptions, such as belief-based utility, are necessary to model it.

2) Genetic testing (23andMe)

Background: 23andMe is a genetic testing service that provides information on customers' ancestry composition and genetic predisposition to health risks. A person who does not get tested may not know that they are at an increased risk for a certain disease (e.g., cancer, cardiovascular disease), which may increase the likelihood of its manifestation (e.g., because no precautions are taken). Moreover, if one family member has their genes sequenced, other family members are able to infer that they likewise have an increased risk of certain diseases.

Function of deliberate ignorance: Results pointing to an increased risk of a certain disease may imply monetary, emotional, and other costs for the individual and others (e.g., partner, family). Not getting tested helps to regulate these emotions.

Ethical principles: From a consequentialist perspective, which actions bring about the best result will depend on how the actor and their loved ones respond if the actor indeed has an above-average propensity of developing a serious and potentially life-threatening disease. From a deontological perspective, the principle of nonmaleficence can be invoked: "Do not hurt the feelings of your family and loved ones."

Rationality principles: The choice of deliberate ignorance cannot be accommodated within standard rationality frameworks such as game theory or expected utility theory; additional assumptions such as belief-based utility would be necessary to model it.

3) Respecting privacy: Reading a family member's e-mails or diary

Background: A person has the opportunity to secretly read a family member's private correspondence (e.g., e-mails, love letters, diary).

Function of deliberate ignorance: The choice not to breach another's privacy maintains trust in significant social relationships. Accessing another's e-mail account is a breach of trust, irrespective of what the e-mails might contain. Both the immediate "victim" and others are likely to lose trust in the snooper.

Ethical principles: From a consequentialist perspective, the choice to respect privacy seems to bring about the best result, as this choice reduces the risk of collective mistrust and its downstream effects. From a deontological principle, people are obliged not to betray the trust of others and to respect their privacy.

Rationality principles: Interactions between family members can be understood as repeated games, and the decision not to know (i.e., the decision not to breach another's privacy) can thus be modeled as rational.

4) Bone marrow donation

Background: Bone marrow produces new, healthy blood cells (around 200 billion every day). Healthy people can become bone marrow donors for patients fighting life-threatening illnesses (e.g., some types of cancer). The donation is a surgical procedure in which liquid marrow is drawn from the donor's pelvic bone and transferred to the recipient. The blood type of donor and recipient must match. Donors may experience side effects such as headaches, dizziness, fatigue, muscle pain, and nausea.

Function of deliberate ignorance: Choosing to remove one's name from a bone marrow donor registry helps to eschew responsibility; the potential donor will never find out if there is a need for their tissue (Dana et al. 2007).

Ethical principles: From a consequentialist perspective, everybody being registered and getting notified in case of need would bring about the best result. The choice of removing one's name from the registry seems to be in conflict with the deontological principle of beneficence (helping others), because a potential recipient may die. At the same time, the principle of autonomy leaves the choice of whether or not to become a donor to the individual. There seems to be societal consensus that it is undesirable, but not unethical, to "opt out" of being a bone marrow donor or, more generally, an organ donor.

Rationality principles: The decision not to be on a registry can be rationally reconstructed if the choice is understood as one of strategic ignorance that allows the agent to eschew moral responsibility. Dana et al. (2007) refer to such a strategy as exploiting moral wiggle room.

5) Society's sins of the past

Background: Societies that undergo transformations from one political, knowledge, value, and social system to another (e.g., Germany after the defeat of the Third Reich) may decide not to ask, tell, or find out about the sins of its citizens under the old regime.

Function of deliberate ignorance: From a collective perspective, deliberate ignorance may help to maintain social cohesion and peace. From an individual perspective, choosing not to know can help regulate emotions (e.g., not having to grapple with the fact that one's grandparents may have been Nazis).

Ethical principles: From a consequentialist perspective, the choice of deliberate ignorance may bring the best result in terms of social cohesion and peace, especially if the number of victims is relatively small or the number of perpetrators relatively large. From a deontological perspective, not punishing past sins and failing to compensate victims is unethical.

Rationality principles: A game theoretical view might suggest that social welfare is better served under deliberate ignorance than under a knowledge regime (at least for certain periods of the transformation process).

6) Sexual orientation algorithm

Background: Wang and Kosinski (2018) developed an algorithm that, on the basis of five facial images per person, can detect a person's sexual orientation with 91% and 83% accuracy for men and women, respectively. This is higher accuracy than is achieved by humans. Should policy makers prevent the development of such algorithms?

Function of deliberate ignorance: The key function of preventing the algorithm's use is to maintain impartiality and fairness and to leave the choice of whether or not to disclose one's sexual orientation to the individual. The party exercising deliberate ignorance here is not necessarily the individual but the community (e.g., citizens, regulators).

Ethical principles: From a consequentialist perspective, the regulatory choice to ban such algorithms may bring the best result given that, relative to the meager benefits (e.g., targeted advertising), people's right to privacy is at risk (e.g., in the case of a person who has chosen not to make their homosexuality public). From a deontological perspective, this regulatory action is not in conflict with any key duty. On the contrary, in a society that respects the human right to freedom from discrimination based on sexual orientation, deontological principles (avoiding harm) would also be consistent with banning such algorithms.

Rationality principles: From a marketing perspective, not creating the algorithm may be irrational as it means foregoing profit. A choice of deliberate ignorance here can thus not be explained by rational principles unless ethical principles are invoked or added.

7) Governmental data collection

Background: The collection of any data on race, ethnicity, or religion is prohibited under French law.

Function of deliberate ignorance: The primary goal of the law is to maintain impartiality and fairness.

Ethical principles: France initially applied this form of deliberate ignorance to avoid various forms of discrimination. From a consequentialist perspective, it may indeed bring the best results; however, it depends on the measurable effect. A probably unanticipated consequence of not collecting these data is that it is harder to prove discrimination in the workplace based on race and ethnicity or to design and implement policies to counter its effects (e.g., affirmative

action). From a deontological perspective, this law is not in conflict with any key duty.

Rationality principles: Outside the realm of expected utility theory and game theory.

8) Blind auditioning

Background: In an attempt to overcome bias (e.g., gender, race, affiliation to certain teachers or musicians) in the hiring of musicians, most major U.S. orchestras moved in the 1970s to change their audition policies, both in terms of democratizing the decision-making process and in hiding a musician's identity from the selection committee. Holding auditions behind screens, at least in the early rounds, became the new standard. In 1970, fewer than 5% of the musicians in the top five orchestras in the United States were women; by 1980, the proportion increased to 10%; by 1997, to 25%. This increase has been attributed, at least in part, to the adoption of blind auditions (Goldin and Rouse 2000:715).

Function of deliberate ignorance: The primary goal of having musicians perform behind a screen is to ensure impartiality and fairness; otherwise, the selection committee may be biased (consciously or unconsciously) to select male candidates, candidates known to them, or individuals recommended by teachers or musicians from major conservatories or other orchestras.

Ethical principles: From a consequentialist perspective, shielding the selection committee from irrelevant but biasing information (gender, ethnicity) fosters the best result. From a deontological perspective, the selection committee's act of deliberate ignorance is not in conflict with any key duty.

Rationality principles: The fact that blind auditioning is even necessary cannot be explained; it entails that biases and discrimination exist in the first place. A rational decision maker should be able to consider only decision-relevant knowledge (i.e., the musical performance) and ignore irrelevant or misleading information (e.g., who studied with whom).

9) Teaching evaluations

Background: Teaching evaluations are systematic procedures for reviewing teacher performance by having, for example, students complete a questionnaire. Ideally, the teacher will use the resulting feedback to improve their teaching practice. Some teachers, however, may decide not to read students' teaching evaluations.

Function of deliberate ignorance: A teacher may avoid reading evaluations to avoid any negative emotions elicited by negative feedback. By remaining ignorant of the feedback, however, they will forego insights into how their teaching could be improved.

Ethical principles: From a consequentialist perspective, this act of deliberate ignorance does not produce collectively the best outcome. Not reading teaching evaluations makes it harder to improve one's teaching and thus maximize students' learning outcomes. From a deontological perspective, this choice may risk harming others, with harm being interpreted widely and relating to students' learning outcomes.

Rationality principles: Deciding not to read student evaluations can potentially be deemed rational to the extent that acting on them implies the disutility of additional work in the future, but only if we assume that it is impossible for a teacher to simply disregard critical evaluations.

10) Entrepreneurial success

Background: It is thought that 80–90% of all start-ups fail. An entrepreneur may decide not to explore the chances of their company's success (e.g., by analyzing the market or retrieving data on the success rates of start-ups in a comparable domain).

Function of deliberate ignorance: Not knowing that most new businesses fail can safeguard an entrepreneur's motivation to set up a new business and their belief in its success (and thus the motivation to invest significant time, cognitive, and financial resources).

Ethical principles: From a consequentialist perspective, individuals exercising this type of deliberate ignorance may dare to take entrepreneurial risks, to innovate, and potentially reap large rewards. From a deontological perspective, it may harm others to the extent that they invest their capital or labor in a business with a low chance of success. The financial well-being of the new business's employees or the entrepreneur's family may be adversely affected should it fail.

Rationality principles: From a utility perspective, not knowing cannot be rationalized; the probability of success should be found out and assessed.

