



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

Institutions Promoting or Countering Deliberate Ignorance

Doron Teichman, Eric Talley, Stefanie Egidy, Christoph Engel,
Krishna P. Gummadi, Kristin Hagel, Stephan Lewandowsky,
Robert J. MacCoun, Sonja Utz, and Eyal Zamir

Abstract

This chapter examines the institutional implications associated with facilitating or combatting deliberate ignorance, and explores concrete institutional mechanisms that could serve to limit, distort, or otherwise structure peoples' informational environment. It examines the basic building block that individuals might use to achieve their goals—contracts—and highlights the advantages and problems associated with consensual mechanisms that could be used in this regard. The chapter further analyzes how organizational structures and mechanisms (e.g., corporations) may be utilized to compartmentalize information and construct the informational environment. Finally, it introduces a new institutional frontier—technology—and shows how developments in the areas of artificial intelligence and machine learning can promote the goals discussed throughout the chapter.

Introduction

Other chapters in this book have established that under certain conditions, individuals might hold a preference to remain ignorant. Such preferences raise tricky normative concerns: whereas in some situations it might be preferable for people to fulfil their preferences for ignorance, in other settings, it might be

Group photos (top left to bottom right) Doron Teichman, Eric Talley, Sonja Utz and Stefanie Egidy, Kristin Hagel, Christoph Engel, Doron Teichman, Robert MacCoun, Krishna Gummadi, Eric Talley, Eyal Zamir, Sonja Utz, Kristin Hagel, Stephan Lewandowsky, Eyal Zamir, Doron Teichman, Robert MacCoun, Stefanie Egidy, Christoph Engel, Stefanie Egidy, Krishna Gummadi, Stephan Lewandowsky

better to discourage or completely deter the ability of people to remain ignorant. In this chapter, we evaluate the institutional implications stemming from these insights, exploring which policy tools should be used in conjunction with deliberate ignorance.

Given the broad scope of the discussion, we begin by clarifying the types of situations that constitute the focus of this chapter and distinguish between different units of analysis germane to our discussion. A first category of cases concerns a *lone actor*, someone who seeks to perpetuate ignorance as to a personal matter for which others are unaffected or indifferent. A second category consists of *actor versus actor* scenarios, where the decision not to inform oneself entails externalities as to other persons. In *principal-agent* scenarios, for example, one actor (the agent) carries out actions on behalf of another (the principal). In such cases the principal may arrange for the agent to be kept ignorant (e.g., as when an editor blinds the identity of a manuscript author to reviewers), or the agent may act based on information the principal affirmatively does not wish to know (e.g., for efficiency, or in cases of “plausible deniability,” as a means of avoiding responsibility). More complex cases involve *distributed agency*, where individuals operate as members of larger collectives such as corporations or government agencies. Here, it may be meaningful to describe the collective entity as deliberately ignorant even though some members of the collective are not ignorant; alternatively, the collective entity might be said to “know” something even though each individual member is partially blind to the whole. Cutting across all these scenarios, there are also *actor versus audience* issues where third parties have an attenuated personal stake in a decision, but may nonetheless care about how it is handled (e.g., citizens may demand transparency as a matter of good governance, or officials may opt for transparency to promote a sense of legitimacy).

Three distinct strategies could be deployed in these situations to design the preferred informational environment. First, one could focus on the elimination of unwanted information. Such a goal might be achieved *ex ante* by preventing the creation of the information in the first instance (e.g., do not collect data on race), or *ex post* by destroying existing information (e.g., burn the Stasi files; see Ellerbrock and Hertwig, this volume). Second, one could take the existence of information as given and attempt to shield the person from the information. This goal could be attained either by quarantining the person (e.g., software that filters information out of a person’s environment) or the information itself (e.g., information escrows). Finally, if individuals have been exposed to the unwanted information, one could still attempt to limit its impact by adopting a decision rule that requires actors to ignore it.¹

¹ Existing studies in the area of judicial decision making suggest that people often cannot ignore relevant information (see Zamir and Teichman 2018). However, specific case studies do show that this strategy could prove effective (Rachlinski et al. 2011).

That said, to the extent deliberate ignorance is *undesirable*, the polar opposite policies are then warranted. Policy makers could mandate the creation of information by requiring the collection of data, barriers to the free flow of information could be removed, and decision makers could be obligated to incorporate certain information into their choices.

In this chapter, we examine these potential institutional responses. Before delving into the details, however, two preliminary remarks bear emphasis. First, we adopt the definition of deliberate ignorance as presented by Hertwig and Engel (this volume, 2016) according to which deliberate ignorance is defined as “the conscious individual or collective choice not to seek or use information.” We fully acknowledge that there are borderline cases which test the boundaries of this definition, and thus we have limited our discussion to what could be dubbed as the “easy cases” of deliberate ignorance.

Second, any institutional reaction to deliberate ignorance presupposes a normative judgment regarding the “all-things-considered” desirability of ignorance in the context under consideration. Whereas in some situations deliberate ignorance might be desirable behavior that should be facilitated, in other settings deliberate ignorance might reflect problematic behavior that should be discouraged. In still other settings, the relevant decision makers may not be confident about the normative desirability of ignorance. Here we limit our analysis to mapping the potential institutional tools that are geared toward ignorance, assuming the policy goal is prespecified and clearly understood.

Below, we explore concrete institutional mechanisms that could serve to limit, distort, or otherwise structure peoples’ informational environment. We begin by examining the basic building block that individuals might use to achieve their goals—contracts—highlighting the different consensual mechanisms that could be used and exploring whether they should be regulated. We then discuss the role of organizations by examining how organizational structures and mechanisms (e.g., corporations) construct the informational environment and how different mechanisms might be utilized to further or limit peoples’ ability to compartmentalize information into different organizational units. Our discussion of “deliberate opacity” progresses up the societal ladder and looks at the role of ignorance within the state, and highlights limits of transparency. We then introduce a new institutional frontier—technology—and explore how developments in artificial intelligence and machine learning might be used to promote the goals discussed throughout the chapter. In closing, we offer a roadmap to chart future research on the topic.

Contracting for Ignorance

Individuals may have a preference to ignore a piece of information or information within a defined category. A key institution that enables individuals to act on such preferences is through the making of enforceable promises or

contracts. Contracts that effect ignorance could, in principle, be narrow agreements that are limited to the informational assets (e.g., confidentiality/nondisclosure agreements). They might also be bundled into a broader transactional framework in which assent to the ignorance component is not independently elicited (e.g., the terms and conditions of a website). Here we examine the extent to which contracts can effectuate allocations of information and explore some of the potential challenges associated with such contracts.

The Feasibility of Contracting for Ignorance

The core of a contract for ignorance is a promise to enable the party or parties to structure the informational environment they wish to have. At times, the informational preference might be a by-product of a broader underlying contractual relationship. This occurs, for example, when patients have preferences regarding the type of genetic information they wish to receive. Alternatively, the contract might focus on information created by third parties. For instance, an Internet platform might promise its users to shield them from certain types of undesirable information (e.g., violence, pornography, hate speech).

Contractual arrangements can both eliminate unwanted information and quarantine it. As to the former, many commercial agreements include provisions that require the destruction of information. For example, Section 5a of the American Bar Association (ABA) Model Confidentiality Agreement, associated with corporate acquisitions, includes the following information provision mandating *ex post* destruction of information:

If either party to this letter of agreement determines that it does not wish to proceed with a Transaction, it will promptly inform the other party of that determination. In that case, or at any time upon the request of the Discloser for any reason, the Recipient will promptly, and in any event no later than 30 days after the request, deliver to the Discloser or, at the Recipient's option, destroy all Evaluation Material (and all copies, extracts, or other reproductions thereof), whether in paper, electronic, or other form or media, furnished to the Recipient or its Representatives by or on behalf of the Discloser pursuant to this letter agreement. In the event of such a determination or request, all Evaluation Material prepared by the Recipient or its Representatives shall be destroyed within such 30-day period and no copy, extract, or other reproduction thereof shall be retained, whether in paper, electronic, or other form or media.

Similarly, even if a contract does not go as far as to order the destruction of information, many such agreements may instead require that information be held by some sort of escrow agent, such as an external law firm, or limited only to certain divisions within an organization, such as the general counsel's office (ABA Model Confidentiality Agreement 2011, commentary at 356).

To be sure, contracts cannot realistically guarantee individuals the precise informational environment they desire. For one, informational barriers often impede the possibility of specifying all of the information from which the

promisee wishes to be shielded. Consequently, a contract calling for ignorance might not be able to capture all of the benefits associated with trade. For example, a patient might not realize that a benign piece of information obtained from a routine procedure (e.g., blood type) could result in unintended consequences and distress (e.g., misattributed parenthood). In addition, practical constraints limit the extent to which a promisee can be shielded from certain types of information. For instance, healthcare experts exposed to such information may have difficulties disregarding it, even though they are contractually instructed to do so, and in subsequent interactions with a patient, they might unwittingly reveal its content. Moreover, information is also obtained through real-world encounters (e.g., a billboard viewed every day, television ads) and might be difficult for a promisee to overlook. For contracts to serve as an adequate tool through which individuals limit their access to information and achieve deliberate ignorance, these limitations must be understood and addressed.

Regulating the Ignorance Contract

While contracts might help individuals design the informational environment they seek to create, numerous imperfections might lead to a need to regulate aspects of the contractual relationship. Most immediately, legal institutions can help facilitate the task of ignorance contracting when it is desirable via default rules, interpretive conventions, and remedies for breach. In addition, legal regulations may sometimes limit ignorance contracting when it appears undesirable for normative reasons.

Reviewing the entire range of possible regulatory interventions is beyond the scope of this chapter. Such interventions could range, however, from procedural rules that help facilitate the transfer of information between parties during the formation of the contract, to substantive rules that aim to alter the content of the contract (through both default and mandatory rules). In addition, the legal system could decide not to enforce certain contracts, if they are in conflict with public policy concerns. The precise legal tool depends on both the contractual context and the regulatory environment in each jurisdiction.

An initial set of problems that might merit a regulatory intervention in the ignorance contract stem from *internalities*; namely, the need to protect the interests of one of the contracting parties given the imperfections of the contracting process. Even perfectly rational parties make suboptimal contracting decisions due to a host of market failures. First among these is the lack of information available to customers about the content of their contractual obligations: most contractual clauses in most current contracts (i.e., in standard form contracts) are practically invisible to customers because they do not (and cannot reasonably be expected to) read them. This, along with a host of other market failures, allows suppliers to include provisions in the contract that appear to reflect the customer's wish to remain ignorant of some information, or to receive other types of information. When such clauses are the product

of suppliers' exploitation of customers' lack of information or other market failures, regulating them might rest on efficiency considerations (and often on non-efficiency grounds as well).

To the extent that contractual parties are not perfectly rational value maximizers, the case for a regulatory intervention grows more compelling. The complexity of the ignorance contract, coupled with issues such as bounded rationality, illiteracy, and limited education, might all cause promisees to agree to contracts that do not serve their interests.

The doctrine of informed consent, and the inability of patients to contract around it in certain instances, serves as a case in point. When patients receive medical treatment, the law usually requires medical professionals to explain the potential risks and side effects of the treatment. Only then are patients able to give informed consent. A patient, however, might not want to hear any of this information and choose instead to remain deliberately ignorant (see also Zamir and Yair, this volume). This might occur when the patient experiences anticipated anxiety. In addition, knowledge of side effects might increase the probability that they will actually occur. Evidence of the latter has been obtained in randomized clinical trials, where it was shown that if a patient forms a negative expectation regarding certain side effects of a medication, the patient may experience the anticipated effects even when treated with a placebo—a *nocebo effect* (Häuser et al. 2012).

Notwithstanding this potential preference for ignorance, some restrictions, at least in the German legal system, seem to be derived from a paternalistic idea of protecting the patient's autonomy and right not to be informed. As there is a dearth of case law, the exact extent of these restrictions is, of course, subject to debate. In essence, there seems to be agreement that patients are required to have an overall understanding about the medical process and the general risk level in order to be able to consent to a procedure. That is, German law removes from peoples' choice set the option to remain ignorant. For experimental treatments, the patient's right not to know seems to be even more constrained.

A second set of issues that might merit regulatory intervention in the ignorance contract arises from the effect of ignorance-creating contracts on third parties. Information can often be useful for many people. As a result, a contract that eliminates or quarantines information might entail significant negative externalities, since it prevents parties who are not part of the contract from utilizing and benefiting from the information at hand. Contracting parties may agree to destroy information that could be beneficial to third parties, or they may agree that one of them would be shielded from some kind of information, although such ignorance might adversely affect people interacting with that person. To protect the interests of such third parties, a regulatory response might be required.

One example of this conflict is the case of sperm donation. To facilitate the creation of a child from a sperm donation, different agreements have to be

concluded. The actors involved are the sperm donor, the fertility clinic, and the recipient (e.g., the birth mother). The sperm donor often has an interest in remaining anonymous and ignorant (also in the future) of any biological children created, and might want to protect this interest within the sperm donation contract with the fertility clinic. The contract between the fertility clinic and the sperm donor will thus often include a clause that ensures the sperm donor's privacy, and possibly the deletion of his identifying information. This agreement will also be reflected in the contract between the fertility clinic and the recipients of the sperm donation. This legal construction of the relationship means that any child born through a sperm donation will be unable to gain access to information about the identity of his or her biological father for any of the following three reasons: the information will already have been deleted by the time the child is old enough to inquire, the child will not know whom exactly to ask, or the fertility clinic will deny access to the information because of its contractual obligations.

In Germany, regulators have decided that the risk of negative externalities merits regulatory intervention. Under the *Grundgesetz* (the constitutionally mandated Basic Law), individuals have a fundamental right to know their biological origin, as part of the *allgemeines Persönlichkeitsrecht* (provision that delineates general personal rights) enshrined in Articles 1 and 2. The contractual agreement described above concerning sperm donation would limit, however, an individual's ability to exercise this right. One potential solution to this problem could be not to enforce such secrecy agreements. Yet given the possibility that information could be destroyed before an individual pursues its right, this response may not suffice.

Against this backdrop, Germany has recently enacted a statutory regime that mandates the creation and preservation of information. The *Samenspenderegistergesetz* (law governing the registration of sperm donors) establishes a sperm donor registry. The law obliges fertility clinics to collect and transfer personal information about sperm donors (name, place and date of birth, nationality, address), recipients (name, place and date of birth, address, number and birth dates of any children born), and the fertilization process (time of use, successful conception, calculated due date) to the registry. Both donor and recipient must be fully informed about the process and agree to this informational component of sperm donation. The sperm donor can also include a personal message in the register, in which he can state his (unenforceable) wish not to be contacted in the future. The data is kept for 110 years (i.e., the maximal life expectancy according to German legislation). The statute allows persons who suspect they are donor offspring to access any relevant information from the registry. Once a request has been made, the law demands that the sperm donor be informed about the request four weeks before the data is handed out, thus creating awareness of any potential future contact initiated by the donor's offspring. This legal mechanism eliminates the possibility of deliberate ignorance among sperm donors.

Finally, it is worthwhile flagging a more recent context in which broad societal considerations might play a role in determining the desired institutional structure: the context of public and political discourse. In recent years, a significant part of public and political discourse has shifted to the digital city square. Platforms such as Facebook, Twitter, and the like have become the central point of political campaigns and public debates. Although these debates have been subverted by *fake news* (Allcott and Gentzkow 2017), the business model of platforms such as Facebook currently profits from attention and “likes,” thus reducing their incentive to facilitate deliberate ignorance among users (see also Krueger et al., this volume).

While constitutional provisions that safeguard freedom of speech generally protect the creation of information, they are imperfect when it comes to enabling individuals to quarantine information that they deliberately wish to ignore. As ever more information arrives to us via digital channels, those digital channels can be designed to further personal preferences regarding information. People could choose, for example, to remove information from their digital environment that relates to opposing political parties, social movements that they find repugnant, and so forth (we discuss this further below in the section on “The Role of Technology”). In fact, social media algorithms often reinforce deliberate ignorance, for example, by not exposing people to opposing opinions and creating filter bubbles and echo chambers (Pariser 2011). Technically, it is possible to program algorithms that expose people to the full range of arguments on an issue or that filter/flag fake news. Nonetheless, the decision to regulate such private platforms hinges on thorny normative questions that define what is the proper shape of public discourse in modern democratic societies.

Beyond Private Ordering

Markets hold promise as entrepreneurs hope to profit from addressing the preferences of potential clients. Entrepreneurs have an incentive to understand these preferences as precisely as possible, and to design solutions that exactly match these preferences. In principle, markets are powerful because buyers are protected by competition. If one provider does not satisfy them, they can stop purchasing their services and trade with a competitor. Yet competition in markets for content is notoriously precarious. Many content markets are in the hands of a very small number of providers, if not a single one. This is also the case for many commercial platforms. The main economic reason is network externalities: the value of the service grows nonlinearly with the number of customers. In such markets, the only competitive pressure results from the possibility that one content provider or one platform is superseded by a superior (or merely more popular) new player. The less credible the competitive threat, the more it is likely that desirable information is withheld from a customer. Such market failures can justify regulatory oversight.

In addition, markets by their very nature cater to preferences that are backed by an ability to pay. As a result, the power to design a person's informational environment might be limited only to those able to purchase this service. To the extent that societies care about an egalitarian distribution of the right not to know, this would imply that the state might need to regulate the provision of this service or turn to providing it via nonmarket mechanisms.

The Role of Organizations

Much of human activity is conducted within organizations. States, public and private firms, labor unions, and the like all play a central role in modern lives. Given this crucial function, we focus here on the interplay between organizational structures and deliberate ignorance. Organizational structures add another level of complexity: What does deliberate ignorance mean on a collective level? After a brief overview of the way in which information is produced, transferred, and stored within organizations, we explore how different liability regimes influence the knowledge acquired by firms and individuals within them, and highlight the ability of organizations to quarantine information into a defined domain, thus facilitating deliberate ignorance.

Institutional Knowledge

The literature on knowledge within institutions has focused primarily on knowledge sharing (rather than on deliberate ignorance) and how various knowledge management tools or practices can stimulate it. In this body of work, different perspectives on organizational knowledge can be distinguished. These perspectives influence which knowledge management strategy is selected by a company, but they could also be helpful for the discussion of deliberate ignorance.

According to Wasko and Faraj (2000), knowledge can be viewed as an object, as embedded in the individual, or as embedded in the community. When knowledge is viewed as an object, it is assumed that knowledge can easily be codified and that employees can easily store their knowledge in a repository. Accordingly, organizational knowledge is the aggregate of all knowledge pieces in an organization, and management provides a knowledge repository with search facilities and motivates sharing with financial incentives. Conversely, the knowledge-as-embedded-in-individuals' perspective argues that it is not as easy to separate knowledge from people, as not all knowledge can easily be codified. As a consequence, knowledge is often lost when an expert leaves the organization. Knowledge management should thus help identify the relevant experts and motivate them through recognition or status. The knowledge-as-embedded-in-community perspective goes one step further and argues that knowledge emerges through shared practices and routines. It is thus

more than the sum of individual pieces of knowledge; it is collectively owned and collectively produced in discussions and routines (for a more detailed discussion, see Wasko and Faraj 2000).

Cognitive psychology literature, on the other hand, has dedicated significant attention to the way in which information is dispersed within organizations. The concept of “executive ignorance” (Turvey 1977) refers to the notion that as a matter of efficiency, and perhaps necessity, the conscious component at the top of a hierarchical cognitive system will not have access to or knowledge of the details of lower-level processing. The term has been adopted in organizational behavior literature to refer to the notion that superiors in a hierarchy should delegate authority to subordinates and should not attempt to “micromanage” them; that is, the executive’s time and attention is better occupied by higher-level goals with a longer time horizon.

Liability Rules and Deliberate Ignorance

In many areas of law, corporations are the nominal defendant in either criminal or civil litigation. Nearly all of that litigation is fault based rather than strict liability based. Consequently, to prevail, the plaintiff/state must demonstrate by an appropriate standard of proof that the defendant acted with a requisite state of mind. In the United States, securities fraud, for example, usually requires a type of recklessness associated with a material misstatement or omission, tort cases usually involve showing negligence, and criminal cases usually require either extreme recklessness or willfulness to secure liability. Consequently, courts are routinely required to assess the information flow within the corporation and to examine whether this knowledge can be attributed to the corporation.

To take a concrete case, suppose a middle-level manager for an oil and gas company (Camile) makes an impromptu public statement about the company’s excellent proven reserves and its superb financial condition, and this statement gets investors excited and causes trading markets to respond upward. However, one of the company’s on-site oil-field managers (Emiliano) recently discovered that the company’s proven reserves are almost fully depleted. Meanwhile, the CFO (Tamika), a member of the company’s board, also recently discovered a material weakness in the company’s financial records. It is later revealed that Camile’s statements are both completely false, and the stock price crashes back to (or even below) its initial level. Given this market correction, private investors or the government (or both) are likely to sue the corporation, alleging securities fraud. Under U.S. law, an element of their case is that the plaintiff/the state must prove that “the corporation” acted recklessly in making the false statement; without delving into the nuances of legal doctrine, this requires some type of awareness on part of the corporation. Consequently, corporations are often incentivized to construct their institutional knowledge

so that it limits their legal liability—a task that often involves creating and perpetuating ignorance.

Jurisdictions have developed different legal rules that define the conditions under which one can attribute knowledge to a corporation. By most accounts, there are three prevailing “tests” for examining the corporate state of mind: (a) the common-law “bad actor” test; (b) the “collective scienter” test; and (b) the “puppet-master” test. Let us now review these rules and evaluate the ways in which they influence the incentives of corporate actors to engage in deliberate ignorance.

Under the common-law “bad actor” approach, the fact finder must inquire into the state of mind of the individual corporate official who actually acted or made the false or misleading statement.² In our running hypothetical, this would require the determination of whether Camile (the corporate speaker) knew, or was reckless in not knowing, that her statements were false. This approach focuses on the knowledge-embedded-in-an-individual perspective described above. In this case, proving Camile’s willfulness/recklessness would be difficult based on what is known, since it appears that she was not at the hub of information transmission about the company’s proven reserves or financial condition. Indeed, if the oil company knows it is going to be subject to the bad actor rule, it will plausibly organize itself to ensure Camile does not have that access since ignorance will shield the company from liability.

According to the collective scienter approach, corporate state of mind boils down to determining whether the totality of the officers’, directors’, and employees’ knowledge—if all are aggregated and collected (hypothetically) within the mind of a single person making the statement—meets the required level of sufficient knowledge to assign liability.³ Similar to the knowledge-as-an-object perspective, the collective scienter approach considers the aggregate of single pieces of information as the corporate knowledge. In our example, the collective scienter test effectively ascribes any knowledge that Emiliano and Tamika have to Camile (even if she did not, in fact, know it), and then determines whether (according to that ascribed knowledge) Camile acted recklessly. Given the facts outlined above, it seems almost certain that there would be liability.

Finally, the “puppet-master” test asks whether any of the company’s senior officers/directors had the requisite state of mind, regardless of whether they were the ones engaging in the assertive act.⁴ In our running hypothetical, the critical person is Tamika, who clearly knows about the financial weakness of the company, but seemingly does not know (yet) about the lack of proven

² See, e.g., *Southland Securities v. INSpire Ins. Solutions*, 365 F.3d 353 (5th Cir. 2004); *Phillips v. Scientific-Atlanta, Inc.*, 374 F.3d 1015 (11th Cir. 2004).

³ See, e.g., *U.S. v. Bank of New England*, 821 F.2d 844 (1st Cir. 1987, criminal case); *Monroe Employees Retirement Sys. v. Bridgestone* 387 F.3d 468 (6th Cir. 2004).

⁴ See *Glazier Capital Manage, LP v. Magistri*, 549 F.3d 736 (9th Cir. 2008).

reserves. Thus, under this test, one might imagine that the plaintiffs/government would be able to prevail only in their fraud claim as to the financial health of the corporation, but not as to the misstatement about the company's proven reserves. Much like the case of the bad actor test, the "puppet-master" test might also create strategic incentives for deliberate ignorance.

Organizations as Facilitators of Deliberate Ignorance

Thus far our analysis has focused on the structure of information within a single firm. Yet since firms and organizational structures create clear boundaries between assets, they can be utilized to partition knowledge as well. Consequently, organizational mechanisms can serve to quarantine knowledge, thus assuring that certain people remain ignorant. Isolating information into a discrete organizational tool can serve numerous goals: incentivizing the creation of knowledge, addressing the conflict of interests, and strategically avoiding liability. We explore these three functions in turn.

Governments hold a vast amount of information regarding individuals. This information could serve many competing legitimate goals the government might wish to promote. Nonetheless, individuals might be reluctant to transfer information to the government if that information can be freely used to promote any end the government sees fit. For example, while citizens might agree to transfer biometric information willfully to the government to obtain a passport, they might not agree to allow the government to make any other use of this information. Thus, if the government wishes to promote the transfer of information from its citizens, it might wish to credibly commit to limit the use of the information it receives. To this end, the government may adopt rules that prevent the free flow of useful information within it and partition the information it holds into separate organizational units. Viewed as a whole, the government in this setting can be said to be partially ignorant: while it holds the relevant information, certain sections of the government cannot gain access to this information and consequently remain ignorant. Similarly, this framework holds for scientific institutions, which often rely on the collection of public data to address scientific questions. According to the principle of the protection of data privacy and good scientific practice, the collection of data that can reveal an individual's identity (including IP addresses in online surveys) is either prohibited or subject to an anonymization process. In this sense, deliberate ignorance is highly desired to promote common knowledge.

Organizational partitioning of information may help address situations involving conflicts of interests. For example, an accounting firm might operate as both an auditor and as a business consultant for the same client. To fulfil both roles faithfully, the accounting firm might wish to conduct the two activities within two separate organizational frameworks. Once separate entities are created, different barriers to the free flow of information can be constructed (i.e., so-called Chinese walls/firewalls).

Finally, it should be noted that organizational structure can also be used strategically to avoid legal liability. As discussed above, legal liability often hinges on the ability to attribute knowledge to the organization. By adopting a complex organizational structure in which corporate knowledge is dispersed between numerous subsidiaries, the firm might be able to insulate itself from liability.

Numerous real-world cases suggest that although many corporations are relatively flat, others maintain a highly segmented corporate structure involving dozens (if not hundreds) of subsidiaries and just as many lower-generation subsidiaries. To take a particularly extreme example, at the time of the Deepwater Horizon oil spill in 2010, British Petroleum (BP plc) had 75 immediate subsidiaries, which in turn collectively had 90 second-generation subsidiaries, which in turn collectively had 54 third-generation subsidiaries, which in turn had 25 fourth-generation subsidiaries, which in turn had two fifth-generation subsidiaries (E. Talley, pers. comm.). Most of BP's subsidiaries were identified with either a unique subindustry (e.g., chemicals), a unique geographic region (e.g., Africa), or both.

Multiple reasons contribute to such segmented business structures. For many oil and gas companies, national regulatory requirements (both for energy and tax) often mandate that their in-country operations be separately incorporated. In the event of radical forms of regulation (such as nationalization), the subsidiary structure likely reduces the uncertainty associated with expropriation. In addition, the multi-subsidary structure confines other forms of liability risk, usually containing it within the operating subsidiary (so long as the corporate structure adheres to the formalities of its own separate structure), making it possible for a multinational to operate at scale without similarly magnifying their exposure. Yet because such corporate structures mandate formal governance "separateness" between parent and subsidiary firms, segmented business structures, such as BP's, also create potential organizational barriers within companies, particularly those horizontally connected to one another in the corporate hierarchy. Although informational separation is not a requirement for maintaining limited liability, when it is convenient or desirable for the organization to cultivate such separation, the subsidiary structure is amenable to it.

Deliberate Opacity

Democratic societies place a high value on transparency for well-known reasons: open and transparent procedures discourage corruption, facilitate improvement, and promote a well-informed citizenry. As Justice Louis D. Brandeis famously noted: "Publicity is justly commended as a remedy for social and industrial diseases. Sunlight is said to be the best of disinfectants; electric light the most efficient policeman." Nonetheless, there is value in opacity to many decision-making processes. Otto von Bismarck is reputed to

have suggested that “[l]aws are like sausages. It’s better not to see them being made.” More recently, MacCoun (2006) demonstrated that transparency itself is not unequivocally positive in its effects (see also Lewandowsky and Bishop 2016). To a certain degree, transparency is the opposite of deliberate ignorance, as it focuses on making information easily available to those who wish to use it. In this section, we elaborate on the connection between deliberate ignorance and transparency, and explore why and how the public might decide to shield itself from information regarding the activity of the state. It should be noted, however, that while this discussion focuses on the transparency of the state, many of the arguments apply to private entities as well.

As a descriptive matter, much of the operation of governments and organizations is not transparent. The government routinely classifies documents and restricts public access to information. A collective decision of a community to shield itself from certain information could reflect the limited value this information has to group members and the potential adverse effects of publicizing it. Citizens might rationally choose not to be aware of the intricate details associated with the national security of their country, simply because this information is inconsequential from their perspective, and highly valuable to hostile powers.

A perhaps less obvious reason for limiting governmental transparency arises from situations in which the desirable decision requires conducting *taboo trade-offs*. For example, a community might want governmental decision makers to conduct cost-benefit analysis with respect to investment in safety, a process that would require decision makers to put an explicit price tag on human life. However, such cost-benefit analysis might run against the taboo that human life should not be evaluated in monetary terms. To sustain the taboo, it might be valuable to shield the public from the cost-benefit analysis. Thus, limiting transparency can allow communities to have their incommensurable cake and eat it too. Against this backdrop, Fiske and Tetlock (1997) argue that societies often develop norms to avoid openly acknowledging necessary choices that inevitably trade off two or more deeply held values (for an elaborate discussion, see Calabresi and Bobbitt 1978).

Governments might also choose to limit transparency to aid complex collective decisions that require balancing between different constituencies (Brunsson 1989). By acting in a stealthy manner, the government might avoid some of the political friction that would have been caused had its actions been conducted transparently. For example, if mere knowledge of a government’s activity causes distress to a subset of the population (e.g., orthodox Jews who strongly believe that the state of Israel should not engage in public works such as road construction on the Jewish day of rest, *Shabbat*), then adapting a “don’t ask don’t tell” type of policy vis-à-vis the distressed community can enable the government to act in a way that reflects the preferences of the majority (i.e., pave roads and lay rail lines over the weekends), while avoiding psychic harm from being inflicted on some of the population.

Another aspect of the institutional design of deliberate opacity policies relates to temporal availability of information. Democratic societies often make a conscious decision to shield themselves from current information, while committing themselves to reveal this information in the future through a process of declassification. For instance, it has been suggested that members of the European Union consciously chose to shield information relating to the trade-offs involved in the creation of the Union, since revealing them might have undermined the process at the time (for additional examples, see Zamir and Yair, this volume).

Declassification could reflect a careful balance between transparency and deliberate opacity. If deliberations take place in private yet are recorded, then internal monitors (e.g., legal advisors, state comptroller office) can still provide some oversight. Furthermore, the prospect of future publicity might serve as a check on governmental power and assure that at least some of the benefits of transparency are realized over the long term.

Finally, limiting transparency might also reflect a conscious choice to enable a decision-making process that is insulated from external pressures. Recent experience shows that transparency comes at a considerable cost: When private emails between climate scientists were hacked and distributed on the Internet a decade ago, this enabled political operatives to construct a narrative about alleged corruption and misconduct among climate scientists that arguably delayed policy action and reduced public commitment to environmental policies. In fact, the scientists involved were exonerated in nine independent investigations in the United States and the United Kingdom (Lewandowsky 2014). This is not an isolated incident; it illustrates the unavoidable implication of unlimited transparency when it is not counterbalanced by privacy considerations. Freedom-of-information requests for scientists' emails have become a common weapon in the arsenal of political operatives who seek to undermine scientific findings they oppose, which in turn has arguably led to self-censorship among scientists in their exchanges with colleagues, thereby compromising the rigor of scientific debate.

In conclusion, our point is not to minimize the risks and costs of limiting transparency, but simply to suggest that transparency should not automatically block the option of deliberate opacity. The relative costs and benefits of deliberate opacity and transparency need to be evaluated on a case-by-case basis.

The Role of Technology

Here we turn to a topic that lies at the frontier of research on deliberate ignorance and examine the way in which the emergence of technologies such as artificial intelligence and machine learning impacts the design of peoples' informational environment. As the discussion will show, computerized decision making interacts with many of the institutional questions reviewed throughout this chapter. To

a large degree, our ability to control the decision-making environment in which computers operate, including the possibility of instructing computers to disregard information they were exposed to, suggests that computers could be an effective tool through which information can be deliberately ignored.

We begin with a brief introduction to key terms in the area of computer science, before exploring whether machines can be blinded. Thereafter, we turn to compare human and machine decision processes, and highlight their comparative advantages.

Principles of Machine Learning

Algorithms are increasingly being used in decision-making tasks that affect human lives. These tasks vary from predicting the risk of a defendant recidivating to estimating the creditworthiness of a person seeking a loan. By algorithms we mean computer programs that outline precise and detailed step-by-step instructions for how information given as input to the programs (e.g., features of a defendant or a loan applicant) should be processed to obtain decision outcomes (e.g., recidivism or credit-risk predictions). Intuitively, algorithms can be thought of as detailed recipes for turning information inputs (ingredients) into decision outcomes (dishes).

Traditionally, algorithms were designed (programmed or coded) by humans with considerable planning, effort, and time. When outlining an algorithm (i.e., the precise step-by-step plan for processing input information), human programmers would carefully consider what information would be accepted as input and in what form or representation.

In contrast to traditional human-driven algorithm design, there has been a rise in recent years in machine-driven algorithm design. This is more commonly referred to as *machine learning*. The key idea behind machine learning is to learn automatically the decision-making algorithm from a set of example decisions. For example, to learn the algorithm that decides who is credit worthy, a bank could simply compile a data set of past credit assessments made by their employees and derive the algorithm from this sampling. The data used to learn algorithms is referred to as *training data* and the increasing availability of large training data sets in various decision-making scenarios from banking to predictive policing has catalyzed the adoption of machine learning to design algorithmic decision-making systems.

There are many significant pros and cons to machine-driven algorithm design compared to human-driven algorithm design. On one hand, machine learning saves considerable human effort when designing algorithms; utilizing advances in computer hardware and learning techniques (e.g., deep learning), machine learning can potentially identify complex, yet useful patterns in large-scale training data that are beyond human cognitive abilities. On the other hand, the lack of human supervision over the algorithm design raises concerns about the ability to explain the algorithm's outcomes and points to the vulnerability

of the algorithm design to any form of bias in the training data. In particular, many recent studies have raised concerns about the fairness of learning-based algorithm designs, and fair machine learning has emerged as a topic of considerable interest within the machine-learning research community.

The explored approaches to fair machine learning center around deliberately ignoring certain types of information at different stages of the learning and decision-making process. In the next section, we describe the different approaches and explore whether the approaches are normatively equivalent.

Blinding Machines

One can distinguish three different stages in the design of learning-based algorithmic decision-making systems:

- Stage 1: Prepare the training data
- Stage 2: Learn from the training data
- Stage 3: Deploy the learned algorithm in practice

Let us first focus on Stage 2 to understand how discrimination might arise when learning algorithms from data. The traditional goal of learning models is to design an algorithm that achieves maximum accuracy in its predictions (e.g., risk assessments) over the entire population. Unfortunately, this goal does not guarantee that different socially salient groups of users in the population (e.g., based on race or gender) would achieve similar prediction accuracy. Suppose you have a population where 90% of people belong to race 1 and 10% belong to race 2. Suppose further that you have two predictors, p_1 and p_2 , where p_1 achieves 100% accuracy for people of race 1 and 0% accuracy for people of race 2, and p_2 achieves 85% accuracy for people of both races. When learning, traditional learning models would prioritize learning p_1 over p_2 because p_1 has a higher overall accuracy (90% across both races) compared to p_2 (85% across both races). However, human designers of algorithms might justify selecting p_2 over p_1 due to considerations of nondiscrimination; that is, p_2 has considerably lower disparity (inequality) in accuracy across both races in the population compared to p_1 .

In response to concerns about discrimination, a number of recent works in machine learning have explored new methods to train nondiscriminatory algorithms. At a high level, the methods can be categorized into three main classes, corresponding to the three different stages of the design of algorithmic decision systems at which they are applied.

The class of methods that are applied at the learning stage (Stage 2) are referred to as *in-processing methods*. These methods modify the traditional goal of obtaining an algorithm that maximizes prediction accuracy across the entire population. Specifically, the goal of nondiscriminatory learning is to find an algorithm that maximizes prediction accuracy across the entire population, but subject to the constraint (i.e., bounded by the requirement) that the inequality

in accuracy for different salient social groups in the population is kept low (ideally zero). To achieve these goals, in-processing methods need to use information about social group membership of people during the learning stage (Stage 2).

The class of methods that are applied at the deployment stage (Stage 3) are referred to as *post-processing methods*. These methods do not alter the traditional goal of designing an algorithm that maximizes prediction accuracy across the entire population. Instead, they track the decision outcomes of the potentially discriminatory algorithm in practice (i.e., measure inequality in its prediction accuracy across different salient social groups) and alter the predicted outcomes in a manner that lowers (ideally equalizes) the inequality in the prediction accuracy across social groups. To achieve this goal, post-processing methods need to use information about social group membership in the deployment stage (Stage 3) but not during the learning stage (Stage 2).

The final class of methods that are applied at the training data preparation stage (Stage 1) are referred to as *preprocessing methods*. These methods do not alter the traditional goal of creating an algorithm that maximizes prediction accuracy across the entire population, nor do they change the predicted outcomes of the algorithm in deployment. Instead, they transform the training data in a manner that lowers the chance of generating algorithms that might yield disparate prediction accuracies across different social groups. Examples of such transformations include sampling training data where people belonging to different races are equally represented in the training data or learning new representations of training data (i.e., via dimensionality reduction or expansion techniques) such that people belonging to different races are desegregated in the new feature space. To carry out such data transformations, the preprocessing methods only need to use information about social group membership of people during the data preparation stage (Stage 1) and can ignore the information in learning and deployment stages (Stages 2 and 3).

To date, the machine-learning community has compared the different approaches based on the performance trade-offs they offer between overall and group-level accuracy, as well as practical considerations, such as privacy concerns with needing information about social group membership at deployment stage (for post-processing) versus learning stage (for in-processing) versus data preparation stage (for preprocessing). However, it is an open question whether the three classes of approaches are normatively equivalent. While we cannot present a definitive answer to this question, there are several indications that as a positive matter, people might view preprocessing methods as more palatable, even if their end results are similar to those of the other methods. Specifically, Ritov and Zamir (2014) examined people's assessment of various processes of recruiting students and employees with a view to achieve affirmative action or other goals regarding the composition of the recruited group. They found that processes in which it was theoretically possible to identify people who were tentatively selected,

but ultimately excluded, from the group of selected students/employees, were judged considerably less acceptable than processes in which no such theoretical possibility existed. The authors attributed these results to the psychological phenomena of identifiability (people's different responses to identified, or even identifiable, people, compared to unidentified ones) and loss aversion (the phenomenon that losses loom much larger than unattained gains). Although the study did not address artificial intelligence or mechanical learning, it appears that the same factors may drive people to prefer the use of preprocessing measures in machine learning.

Considerations of public acceptance and trust might also favor preprocessing methods. While in-processing and post-processing measures are inherently unclear and opaque to most people, some of the available preprocessing measures are relatively simple and straightforward. Removing certain types of data from the machine's information set can be easily grasped, even by those who are not trained in computer science. As a result, it could very well be the case that public acceptance of preprocessing measures will be relatively higher. One should note, however, that once more complex preprocessing methods are used, it is unclear whether this hypothesis holds.

Machines versus People

Having discussed the nature of machine-learning processes and how they can be interfered with to achieve desirable goals in contexts such as selection of people in a nondiscriminatory way, let us now examine the pros and cons of using human or machine-learning processes to make decisions. At the outset, we note that the dichotomy between machine and human decision making is to some extent false, for two reasons. First, humans are inevitably involved in practically all stages of the machine-learning processes (and make the normative choices whether and how to use the products of those processes). Second, there is a continuum between primarily human and primarily machine-based decision processes. Thus, for example, whenever people use the Internet to gather information that would feed into their decisions, the information they get is determined to some extent by the machine-learning process used by the search engine. For the sake of the discussion, we nevertheless consider prototypical human versus machine-based decision processes.

Reviewing all of the considerations relevant to the comparison between human and machine-based decision processes lies well beyond the scope of this chapter. Our aim here is to map some of the considerations that should be addressed when deciding whether to entrust humans or machines with various types of decisions, with a focus on the issues related to deliberate ignorance.

A first consideration stems from the discretion afforded to the decision maker. Decision processes may use precise algorithms, employ a well-defined set of variables, give predetermined weight to each variable, and so on. Decision processes may also set more or less abstract goals or values and leave

the decision maker with broader or narrower discretion as to what weight to give to what factors in each particular case, paying heed to all the particular circumstances. In the realm of human decision processes, both possibilities (and any intermediate or hybrid forms) are commonly used in judicial and other contexts (in the legal context, the choice is often framed as a choice between rules and standards). Machine-based decision processes are intuitively associated with precise, algorithmic processes. But is this characteristic necessary? Is it possible now, or is it expected to be possible in the future, to create fuzzier, probabilistic machine-based decision processes that would be more flexible, indeterminate, and standard-like?

A second dimension of comparison relates to explainability, transparency, and accountability. Machine-based processes are in principle more amenable to *ex post* examination. Such examination (or reverse engineering) could let us know which factors were actually addressed, what weight was given to any one of them, and so forth. Human decision making—primarily when made by a single decision maker, but often also when the process is collective—is frequently less tractable and explainable.

At times we are interested in maximizing transparency and accountability of decision making; however, as noted in our discussion on intentional opacity, society might find it preferable to make decisions in less explicit and transparent ways. While it appears that both human and machine-based decisions may be more or less transparent and explainable, there may also be characteristic differences between the two (and in any subcategories thereof) that one might want to consider. One advantage that machine-based processes might have is that they may facilitate “acoustic separation” (see Dan-Cohen 1984) between the public and decision makers, and perhaps even among different parts of the decision-making process.

A concrete example that might illuminate this abstract discussion can be found in the context of antidiscrimination and affirmative action. When zooming in on the issues of antidiscrimination and affirmative action—as major fairness concerns in selection processes—the potential use of machine-based processes is associated with several noteworthy issues.

First, the underlying motivation of discrimination might be relevant to the analysis. There are those who argue that only prejudicial, animus-based discrimination is unacceptable, whereas statistical discrimination is not. Accordingly, as long as the decision maker (e.g., an employer, an education institute, or a landlord) strives for the most accurate decision and uses such attributes as gender and ethnicity merely as indicators of legitimate variables (assuming that the alleged correlation or association does exist, and that making decisions without them would be much costlier), there is nothing wrong with such *statistical discrimination*. If this is so, it is easy to design machine-based selection processes that would avoid animus-based discrimination. However, if one believes that statistical discrimination is unacceptable, then using machine-based processes that rely on correlations and

associations may not only fail to avoid discrimination, but may actually perpetuate it, because they reflect associations established in the past. Careful measures (of the type discussed above) may then be necessary to avoid the discriminatory results.

Second, the introduction of computerized decision making could influence the degree to which the analysis focuses on the issue of *disparate impact*. Some criteria may appear to be perfectly benign, yet using them in selection processes may result in extreme underrepresentation of some groups in the selected body of students, employees, and so forth. While it is a highly controversial issue, some believe that as long as the motivations and processes used by decision makers are fair, the end results are fair as well. In contrast, some believe that criteria and processes that lead to socially unacceptable results should be avoided. Here again, the use of machine-based processes may help to overcome some difficulties, but they may also exacerbate others.

In conclusion, machine-based decision making offers opportunities to operationalize peoples' preference for ignorance, but raises serious concerns at the same time. Many interesting questions still loom in this context. At the institutional level, who should make the above choices, and through which processes? What role should public discourse, academic discourse, legislative, judicial, and administrative deliberation play? On the practical regulatory level, assuming that regulation of machine-based decision processes is legitimate and may be needed, at what point in time are regulations necessary? Should computer experts be free to design whichever tools they wish, and only their use be regulated, or should the development of new tools be subject to regulation *ex ante*? Finally, in the long run, to what extent should decisions be made according to the current development of machine-based decision processes? While the evolution of human faculties is very slow, technological advancement proceeds at a very fast pace. It may therefore be advisable to make any choice contingent on the present or future developments of the technologies.

Conclusion

In reviewing a variety of institutional implications of deliberate ignorance, we have mapped several mechanisms that could be used to design the information environment consistent with a presumed informational goal: prevention of the creation of information, the destruction of information, quarantining information, quarantining users of information, and limiting limitation of the use of acquired information. We have also explored several instantiations of such mechanisms throughout society, from private contracts, organizational structures, to the state itself.

Choosing among the different institutional responses to deliberate ignorance is a herculean task, and constitutes to a large extent a critical research

question for future study. For almost any institutional design task, there is more than one option. As a result, selecting the ideal institutional mechanism requires incorporating all relevant factors: consequences, fairness, political constraints, legitimacy, as well as social norms and practices, to name but a few. Throughout this chapter, we have endeavored to delineate the most important dimensions that should be assessed.

A primary concern involves effectiveness. How well does the institution achieve its intended goal? Threatening a decision maker with a sanction, in the event of using protected information, is less safe than making sure that information is never generated. Destroying the information is a riskier option, but it suffices to check whether those who have ever seen the information can be trusted not to report on it. Effectiveness, however, is not confined to the immediate outcome (the information is withheld). An institution is also more effective if it can work reliably or be relatively easily adjusted when a task is incorrectly specified or the environment changes in unpredicted ways. An institution can also be regarded as more effective if it enables institutional learning, be it for the case at hand or analogous cases that present themselves in the future.

Another dimension of effectiveness results from the relationship between the institution (and those running it, if it is purposefully steered by individuals) and its addressees. Effectiveness may be good or bad for this purpose, depending on whether it is easy to understand and observe the channels at which the institution achieves its intended goal. Such institutional transparency can create trust, but it can also make normative conflicts patent, which would make it difficult for those whose lives are being governed to accept the intervention.

A secondary concern involves externalities. If they are not part of a hidden agenda, one could refer to this concern as unintended negative consequences: on other tasks that this same decision maker expects to face, on other (private or public) parties, on the social fabric, or on the evolution of the community. The spillover may also originate from the fact that the solution, in the guise of an institutional transplant, is picked up in different areas of life, or by different communities. What is good for one context can be detrimental for another.

Even if the institution does not have any adverse effects on third parties, it can amplify heterogeneity among its addressees. When this heterogeneity is pecuniary, it is usually referred to as distributional effects. Given the wide body of distributional theories along with the different legal responses (i.e., regulation vs. taxation), we limit ourselves to highlighting this issue as one in need of future study.

Numerous specific paths for future research emerge from the issues discussed in this chapter. For instance, our discussion of the firm focused on the relations within the firm, yet an entirely separate set of questions relates to the relationship between the firm and its shareholders. On this front, shareholders might wish to shield themselves from financial information that might cause

them to make biased decisions that run against their long-term interests. In addition, shareholders might wish to shield themselves from information alluding to the firm's actual activities to avoid tension between those activities and their views on issues, such as protecting the environment, distributive justice, and so forth. (Of course, it is debatable whether shareholders should be allowed to remain ignorant in such cases.) Similarly, our analysis of technology focused on a relatively small set of questions associated with machine learning. There are numerous other technological frontiers, such as the emerging use of blockchain technology, which could be used to design the informational environment.

There are also many open questions concerning the role of machine learning in fostering or mitigating deliberate ignorance. When it comes to fairness judgments (where deliberate ignorance of attributes such as sex or race is often desirable), research in the machine-learning community has compared *pre*-processing, *in*-processing, and *post*-processing approaches for accuracy and group disparity. These three classes of strategies are roughly equivalent, but research has not yet examined whether they differ in terms of trust or acceptance. Post-processing approaches might, for example, be perceived as unfair because a human at the end of the process changes the machine-learning output. People may prefer the preprocessing approaches since the variable that is protected by antidiscrimination laws is not used, or its use is penalized. One could also argue that in-processing approaches might be favored since reducing disparity is an explicit goal (next to accuracy) of this approach. Which approach is ultimately favored might depend on group membership as well as on beliefs regarding the perceived truth of a stereotype. Although the accuracy rates of the three approaches are similar, decisions on the individual level are not necessarily so. A particular woman might be hired when a preprocessing approach is used, but not when an in-processing approach is employed. That people's preference for equality versus equity, in distributive fairness, depends on whether they are better off is a well-known phenomenon (Messick and Sentis 1983); when it comes to machine learning, which is less easily understood, beliefs about one's chances under the use of different approaches come into play. Future studies should therefore not only compare the acceptance of the different approaches, but also examine the role of explainability. Closer cooperation between social scientists and machine learners may be useful to answer these questions.

Finally, it should be acknowledged that the discussion in this chapter does not purport to be a comprehensive overview of the institutional responses to deliberate ignorance. By its nature, this discussion was defined by the expertise of the participants in the research group. Thus, additional perspectives need to be incorporated into the analysis. For instance, the education system might be able to help train students to deliberately ignore information. To this end, familiarizing students with the tools that facilitate deliberate ignorance could become part of media literacy programs. By teaching young people when to

shield themselves from too much trivial or untrue information, policy makers could foster a more informed public discourse. Hopefully, this void will be filled over time.