

Deliberate Ignorance and the Law

Eyal Zamir and Roi Yair

Abstract

This chapter offers a bird's-eye view of existing legal doctrines and institutions that overcome or foster deliberate ignorance, critically assesses these doctrines and institutions, and considers extensions thereof. It begins by focusing on three legal means of discouraging deliberate ignorance: subjecting people who *could* have acquired the relevant information to the same treatment as those who acted knowingly, imposing positive duties to acquire information, and rendering information more conspicuous, thereby making it more difficult to ignore. It also touches upon the issue of collective ignorance. Thereafter it discusses instances in which the law encourages deliberate ignorance to facilitate better decision making and promote other values. It starts from the basic issue of designing the system of government and constitutional protection of human rights using veils of ignorance and then moves on to more specific legal topics: inadmissibility and other evidence rules, anonymity and omitted details of candidates to overcome the biases and prejudices of decision makers, expungement of criminal records, and the right to be forgotten.

Introduction

In this chapter, we present a bird's-eye view of existing legal doctrines and institutions that foster or overcome deliberate ignorance, assess these doctrines and institutions, and consider extensions thereof. As part of a series of studies on deliberate ignorance, we do not discuss the conceptual, philosophical, psychological, or economic aspects of deliberate ignorance, as these are considered by others in this volume. In particular, we do not delve into the question of whether knowing something is an all-or-nothing matter or (more realistically) a matter of degree of belief (Buchak 2014).

For the purpose of our discussion, we liberally use Hertwig and Engel's working definition of deliberate ignorance (this volume, 2016): "the conscious individual or collective choice not to seek or use information." Without delving

into the definitional question, a preliminary comment about the boundary between so-called rational and irrational ignorance is in order. After offering the above definition, Hertwig and Engel note that their focus is on “situations where the marginal acquisition costs are negligible and the potential benefits potentially large.” But what about instances in which the gap between the costs and benefits is narrow or nil? While Hertwig and Engel do not squarely exclude such cases (sometimes dubbed as *rational ignorance* or *rational inattention*) from their discussion, they implicitly distinguish between various motivations for not seeking or using information. Specifically, one may or may not wish to include instances in which the individual or societal costs of attaining and processing information exceed their benefits, as well as instances in which disregarding certain types of information can improve decision making.

The first group includes the rational choice of citizens to remain uninformed about governmental and political issues that do not significantly affect their lives. It is also closely connected to the debates about disclosure duties that pervade the law (e.g., in the spheres of consumer transactions and medical malpractice law). Inasmuch as the costs of grasping, understanding, and using the disclosed information by the disclosees are prohibitive, these duties are arguably futile (Ben-Shahar and Schneider 2014; Zamir and Teichman 2018:171–177, 314–318). The issues of a citizen’s choice to remain uninformed about political issues, as well as the relationship between deliberate ignorance and mandated disclosure, are both worth exploring, but due to space limitations we will not discuss them in any detail.

In terms of the second group, it may be rational (in the sense of cost-benefit analysis) to disregard certain types of information when this will likely *increase* the accuracy of fact finding and decision making. This is, for example, the common rationale for some (but not all) inadmissibility rules in the law of evidence, as will be discussed below.

This chapter consists of two sections. The first discusses cases in which the law strives to discourage deliberate ignorance, and the second turns to instances in which the law strives to ensure or encourage such ignorance.

Overcoming Deliberate Ignorance

Both retributive justice and deterrence presumably require distinguishing between a behavior that is done with full knowledge of the relevant facts and its expected consequences, and the same behavior that occurs without such awareness. From a retributivist perspective, knowingly harming other people (or not coming to their aid) is more culpable than harming through negligence or by accident. From a deterrence standpoint, subjecting negligent or accidental (as opposed to knowing) conduct to criminal liability creates an incentive to refrain from lawful activities to avoid the risk of erroneous punishment, and therefore warrants a cautious approach (Posner 2014:279). Laws based

on these distinctions, however, create an incentive for deliberate ignorance to avoid fault-based liability. The primary means by which the law can eliminate or weaken this undesirable incentive is by subjecting people who *could* have acquired the relevant information to the same treatment as those who acted (or failed to act) knowingly. A second technique is to impose positive duties to acquire information, while a third one is to render information more conspicuous and salient, thereby making it more difficult to ignore. We discuss each of these techniques in turn and then briefly address the issue of *collective ignorance*.

Willful Blindness, Constructive Knowledge, and Strict Liability

Equalizing the treatment of deliberately ignorant and fully cognizant actors may be attained in two primary ways (Turner 2009:360–362):

1. Treat deliberately ignorant actors as though they were aware of the relevant information.
2. Lower the threshold of liability by dispensing with the requirement of actual knowledge in favor of an objective standard of conduct in light of the attainable information, such as negligence, or even imposing strict liability.

The most direct technique the law uses to overcome deliberate ignorance is to treat it as a substitute for actual knowledge. Thus, the common-law doctrine of *willful blindness* considers the requirement of knowledge to be satisfied when the person in question suspects the existence of a fact but deliberately avoids looking into it. For example, many states criminalize the knowing exposure of sexual partners to the risk of HIV transmission. Absent the doctrine of willful blindness, individuals who suspect that they are at risk of carrying the virus have an incentive not to get tested, to avoid such liability (Ruby 1999:330–331). Sometimes, the law explicitly treats even *reckless* disregard of information as tantamount to *deliberate* ignorance.¹ While the willful blindness doctrine is primarily known for its application in criminal law²—as a means of satisfying the requirement of *mens rea* of knowledge (Robbins 1990)—it may also be used in civil proceedings where an element of actual (as opposed to constructive) knowledge is required. For instance, in some legal systems and under certain conditions, the doctrine of *market overt* protects the bona fide purchaser of stolen goods. When this protection is granted only to a buyer who is unaware that the goods were stolen, willful blindness amounts to actual awareness (Zamir 1990:109–112).

From the standpoint of deterrence, the justification for the doctrine of willful blindness is straightforward: it deters people from circumventing criminal (or civil) liability through deliberate ignorance (Kozlov-Davis 2001:483).

¹ See, e.g., U.S. False Claims Act, 31 U.S.C. § 3729(b)(1)(A).

² See, e.g., §2.02(7) of the U.S. Model Penal Code.

Arguably, the doctrine may also enhance deterrence by serving as a means to overcome the evidentiary challenge of proving that a person actually knew something (Charlow 1992:1359–1360). However, for the remainder of our discussion, we assume that the defendant was deliberately ignorant, rather than actually knew the relevant facts.

From a retributivist perspective, the picture is more nuanced. The point of departure is that someone who acts in ignorance is less blameworthy than one who acts knowingly—but when the ignorance is deliberate, the acts are equally culpable. This *equal culpability* thesis has been subject to some criticism (Charlow 1992; Husak and Callender 1994; Sarch 2014:1052–1071). Specifically, it has been argued that different kinds of deliberate ignorance may involve different levels of culpability. For example, someone who buys stolen goods from his/her friend at a suspiciously low price is more culpable if the reason for not asking any questions about the provenance of the goods was to avoid criminal liability, rather than not to embarrass the friend by questioning the legality and morality of the latter's conduct (Charlow 1992–1413). People who contrive their ignorance so as to allow themselves to act in an unlawful manner are possibly even more culpable than mere knowing actors and are similar to purposeful offenders (Luban 1999:968–969; see also Aquinas 1265–1274, q. 76 art. 4).

Another question that may be of importance in determining culpability is: What would the person have done had s/he known the incriminating facts? David Luban has argued that while deliberately ignorant people, who would have acted in the same way had they known the incriminating fact, are as culpable as knowing offenders, such equivalence does not hold for people who would have refrained from the same action had they been informed (Luban 1999:973–976). One could, however, argue that such counterfactual thinking may be of relevance to the evaluation of the moral status of the agent, but not the culpability of the action under the actual circumstances. Furthermore, it is possible that people who would have refrained from acting, had they known, are *more* blameworthy, since in a sense they could have easily prevented the crime by confirming their suspicion.

Another doctrine that appears to be relevant is that of *collective knowledge* in corporate criminal law. This doctrine allows courts to find that a corporation acted knowingly if the *aggregate* of its employees' knowledge was sufficient for the purpose, even if no single individual within the corporation possessed that knowledge. Thus, for example, if one employee makes a report, and another knows that some of the reported facts are false but is unaware that a report has been made, the corporation may be convicted for knowingly making a false report, even though no single employee had such knowledge (Colvin 1995:18–19). This doctrine is especially useful in preventing corporations from avoiding criminal liability by cultivating a regime whereby no single employee is aware of the incriminating facts (for further discussion, see Teichman et al., this volume). To illustrate, consider the above scenario, but now add the

fact that the manager of both employees suspects that the first employee may know that some of the details in the report are false, so he makes certain that the latter is unaware that the report has been submitted. While the doctrine does not criminalize the conduct of the manager (although he may sometimes bear personal criminal liability, possibly under the willful blindness doctrine discussed above), it does deter the use of deliberate ignorance within the corporation. In fact, two commentators have suggested that the doctrine only applies when there is an element of deliberate ignorance in play (Hagemann and Grinstein 1997).

Unlike criminal law, private law (including contract, tort, property, and unjust enrichment law) is generally more interested in facilitating fair and efficient behavior than in moral culpability. Consequently, it may well treat those who *should have known* something similarly to those who actually knew it, even though the latter are often more blameworthy. The determination as to whether or not someone should have known a given fact is usually made by asking whether refraining from obtaining the information was reasonable under the circumstances. A nice example may be found in the U.S. law of contractual duress, as encapsulated in Section 175 of the Restatement (Second) of Contracts. When a person enters a contract as a result of an improper threat that leaves him or her no reasonable alternative, s/he is entitled to annul the contract. Typically, the threat is made by the other contracting party, but occasionally by someone else. For example, a husband may threaten his wife that he will leave her unless she sells her jewels and gives him the proceeds, in which case he is not a party to the contract of sale between his wife and the jewelry buyer. In such cases, the victim (wife) can annul the contract, unless the other party (the jeweler) made the transaction “in good faith and without reason to know of the duress” (Section 175(2)). Thus, to remove the jeweler’s temptation to go ahead with a profitable transaction under suspicious circumstances, the law treats a person who had *reason to know* of the duress as one who actually knew about it. Similar rules apply to transactions where one party had reason to know (a) that the other party’s consent was induced by *misrepresentation* by a third party, (b) that the other party made the contract due to a fundamental *mistake*, (c) that the latter was unable to act in a reasonable manner in relation to the transaction due to *mental illness*, or (d) that s/he was unable to reasonably understand the meaning of the transaction due to *intoxication* (Regan 1999).³ While such rules cover both *deliberate* and *negligent* ignorance, they do meet the challenge posed by the former (and overcome the difficulties of proving that one’s ignorance was *willful*).

Imposing *strict liability* similarly discourages deliberate ignorance. Under such a regime, a defendant is subject to (criminal or civil) liability for his/her harmful conduct regardless of whether s/he knew, or even ought to have known, the relevant facts. Consequently, such a regime creates an incentive

³ See Restatement (Second) of Contracts §§ 164, 153(b), 15(1)(b), and 16, respectively.

to avoid the harmful conduct and to acquire the necessary information to that end (as long as the cost of acquiring the information and using it to avoid the harm is smaller than the expected sanction). In fact, Hamdani (2007) has argued that strict liability in criminal law is best explained by this rationale and is thus prevalent mostly in instances where there are strong market incentives to remaining ignorant, such as the sale of liquor to minors. Similarly, the general principle that *ignorantia juris non excusat* (ignorance of the law is not an excuse) is, in a sense, a standard of strict liability with respect to knowledge of the law and serves a similar function. Indeed, absent this principle, there is virtually no incentive for anyone to acquire information about the law (Hamdani 2007:448–449). Again, doing away with any requirement of awareness serves goals that are much broader than merely coping with deliberate ignorance, yet it serves the latter as well.

The rules concerning willful blindness, constructive knowledge, and strict liability may be used not only to deter harmful conduct, but also to encourage beneficial conduct, assuming one accepts the viability of this distinction (see Zamir 2015:177–199). However, given that the law is much more hesitant to impose duties to actively help others than to prohibit the active infliction of harm, the willingness to equate negligence or even willful blindness with actual knowledge is more limited in this context. For example, with respect to the crime of failure to prevent a felony, the Israeli Supreme Court has declined to apply the willful blindness doctrine as a substitute for actual knowledge of a plan to commit a crime.⁴

Specific Duties to Acquire Information

Rather than, or in addition to, imposing liability for the harm caused by the actions of the deliberately ignorant, the law occasionally focuses on the information-acquisition phase by imposing a specific duty to acquire information, and even dictating the procedure for doing so. For example, the Financial Action Task Force's recommendations on anti-money-laundering laws require financial institutions and some nonfinancial business professionals to complete a process of customer due diligence. This includes acquiring information about the customer's identity (and, in the case of corporations, its ownership structure) as well as various details about the transaction or nature of the business in which the customer seeks to engage (FATF 2012/2019). If this information raises suspicions about the nature of the transaction, an obligation to report is triggered. This protocol prevents institutions from skirting the obligation to report by remaining ignorant about the nature of the transaction. Similarly, both international and national public law may require an *environmental impact assessment* to be conducted before approval is given to any enterprise that may adversely affect the environment. Meeting this procedural obligation—which

⁴ *Har-Shefi v. State of Israel* (2001) 55(ii) P.D. 735, 756–68.

involves gathering information about the impact of the proposed project, including explicit descriptions of potential alternatives, and identifying potential risks and uncertainties—ensures that a decision will, at the very least, be informed, if not necessarily wise or ethical (Craik 2008).

Although specific information-acquisition duties may impose significant costs (Gill and Taylor 2004), they may be desirable for several reasons. First, they may reframe the context of the process: when a banker asks a customer for additional information about a given transaction, it is not because s/he suspects misconduct, but because the law requires her/him to do so. This type of reframing makes the process more comfortable to both parties. Second, such duties, if accompanied by criminal or administrative sanctions, create additional deterrence against deliberate ignorance by capturing cases where it has not resulted in harm. This issue relates to a broader debate about liability that is risk based, rather than harm based. Third, setting a clear procedure makes it psychologically more difficult to avoid information. In particular, it reduces the vulnerability to *motivated reasoning* (see below), which in part is facilitated by distorting the information-acquisition process so as to promote a desired conclusion.

Provision of Information

Policy makers may reasonably conclude that, all things considered, there is insufficient basis for imposing legal sanctions against a given undesirable act or omission, yet wish to encourage desirable (and discourage undesirable) behavior noncoercively. When people engage in undesirable conduct (or fail to engage in desirable conduct) because they are (possibly deliberately) ignorant of the significance and ramifications of their behavior, providing them with the relevant information may induce them to change their practices. The state itself may provide such information or mandate others to do so.

A case in point is the provision of salient information about the hazards of cigarettes (including graphic warnings that seek to evoke visceral negative reactions) to discourage smoking (Hammond 2011). In a different sphere, where the goal is to reduce employment discrimination, institutions may inform recruiters about existing underrepresentation of women and minority groups in the workforce and educate decision makers about unconscious biases that affect hiring decisions, although the efficacy of these measures, in and of themselves, is rather limited (Kalev et al. 2006). Finally, it has been suggested that a reliable and uniform labeling system for products, which grade them in terms of the impact of their production process on human welfare (e.g., employees' working conditions), the environment, and whether they involve animal abuse, may encourage ethical consumption (Assaf 2016). This suggestion is supported by findings that while consumers deliberately avoid information about the ethical attributes of products, once it is forced upon them, they do use it (Ehrich and Irwin 2005).

While these and comparable instances of information provision do not target *deliberate* ignorance, the line between deliberate and nondeliberate ignorance is often unclear, both conceptually and empirically. Psychological mechanisms such as *motivated reasoning* (the acquisition and processing of information in a manner that leads to the sought-for conclusion; Kunda 1990) and *confirmation bias* (the tendency to seek and process information in ways that are partial to one's interests, beliefs, and expectations; Nickerson 1998) blur this line. People consciously and unconsciously look for confirmatory evidence and tend to ignore disproving evidence, deny its relevance, or give it less weight. Providing clear and conspicuous information about the medical hazards of smoking, the unequal representation of women and minorities in a given organization, or the conditions in which products are manufactured makes it more difficult for people to ignore the troubling information, and might induce them to change their behavior accordingly.

It has also been shown that when people are faced with a self-benefiting choice that might potentially harm someone else, they prefer not to know whether such harm would indeed ensue, so that they can make their choice in good conscience (Dana et al. 2007). Compelling offenders to meet with their victims and realize the consequences of their actions makes it impossible for them to ignore the harm they have caused. Fostering such accountability and responsibility taking is a key element in *restorative justice* programs (Van Ness and Strong 2015:81–96).

Collective Ignorance

Thus far, we have discussed instances in which the law strives to overcome the deliberate ignorance of individuals, but ignorance (and prevalent misconceptions) sometimes characterizes entire societies, or segments thereof. In such cases, its ramifications may be greater than those of individual ignorance. Relevant examples include slavery in ancient Greece and eighteenth-century United States; racial and gender discrimination in many societies; and even active participation in genocide, or turning a blind eye to it, in Nazi Germany. Such practices are typically accompanied by deliberate ignorance or misconceptions of associated facts (e.g., about gender differences) and a blindness to their profound immorality (at least according to current moral values). In fact, customary practices, such as eating meat, may be condemned in the future just as much as the aforementioned practices are condemned by us (as, indeed, some critics already do). We cannot delve into the issue of whether members of such societies are *unable* to see the wrongness of such practices (and are therefore blameless); or should be held *fully accountable* for their active or passive involvement in them; or perhaps held accountable, subject to some kind of *cultural mitigating circumstances*. Importantly, those who support either of the latter two views tend to think that collective ignorance is, at least in part,

deliberate (Moody-Adams 1994; on the related issue of collective forgetting of past atrocities, see Ellerbrock and Hertwig, this volume).

Assuming that collective ignorance facilitates objectionable or even abhorrent practices, what legal means may be taken to fight it? In principle, all three types of measures discussed in this part could be used, *mutatis mutandis*, to counteract not only individual, but also collective ignorance. The difficulty, of course, is that when the government or hegemonic segments of society benefit from the collective ignorance, they are unlikely to try to counteract it, and may even try to frustrate attempts by minority groups to counteract it. Constitutional safeguards of free speech (in democratic countries) and supranational legal norms (that apply to all countries) may have some beneficial impact in this respect. In fact, issues of willful blindness have recurrently been discussed in international criminal proceedings (Van Der Vyver 2004:75–76). Lastly, transitional justice processes are sometimes designed to shed light on the truth of past atrocities and eradicate surrounding ignorance (be it deliberate or otherwise). Such processes may or may not be accompanied by legal sanctions against those who committed such atrocities. For example, truth and reconciliation commissions, like those established in postapartheid South Africa, may offer amnesty in exchange for testimony about the extent of the crimes committed, thereby creating an extensive historical record, and possibly preventing its recurrence (Simonovich 2004:351–352).

Facilitating Deliberate Ignorance

While deliberate ignorance may have negative consequences, it may also bring about positive outcomes. It may, for example, facilitate better decision making and promote other values in instances where full information might be detrimental. In this section, we review some such instances, starting from the general and basic issue of designing the system of government and constitutional protection of human rights, and then moving on to more specific legal topics.

Veils of Ignorance

People often make decisions that affect not only, or even primarily, their own interests, but those of others as well. These include decision and policy making by public officials, such as legislators and judges, CEOs, and other office holders in corporations, attorneys, and other fiduciaries. Such principal-agent relationships often involve a conflict of interest. More generally, people's interests are often misaligned with the overall social good (e.g., due to externalities). One technique, or rather group of techniques, for handling such conflicts of interest and misalignment is to keep decision makers ignorant of how their decisions may affect their interests. In Rawls's (1979) famous thought experiment,

this is achieved by decision makers not knowing what their abilities, tastes, and positions will be in the society whose social order they design. In the real world, people know their current characteristics and position, but a veil of ignorance may be approximated by creating uncertainty about how their decisions might affect them (Vermeule 2001).

One path that the law takes is to structure decision making such that decision makers do not know, or do not know precisely, where their interests will lie when their decisions are implemented. Thus, the separation between a constitution and ordinary legislation—and the supremacy of constitutional over ordinary norms—may be viewed as a means of establishing the fundamental norms of government and the protection of human rights before it is apparent who would benefit from those norms, and who would be adversely affected by them (Vermeule 2001). The same holds true for the ideal of separation between the enactment of general norms by the legislature and their implementation in specific cases by the law courts or the executive (Nzelibe 2011; Vermeule 2001). In the same vein, public and private bodies (including the legislature, the Cabinet, and faculty councils) may adopt two-stage decision procedures with the purpose of achieving similar goals. For example, the Cabinet might first decide on the size of a critical budgetary cut and only then allocate that reduction among the various ministries and agencies. Similarly, an academic department may adopt a long-term development program, in which its needs and aspirations are defined in the abstract, and only then make specific decisions about new recruits.

Another technique that may be viewed through this prism is deferred implementation of legislation and other decisions. It is easier to overcome sectorial opposition to socially desirable reforms if their implementation is postponed, because such postponement creates uncertainty about the reform's gainers and losers, and because some of the future losers are not party to the present decision process (Porat and Yadlin 2006). In keeping with the same logic, it may be desirable to expedite negotiations about the design of general norms so as to conclude them before more information becomes available to the negotiators. A case in point is negotiations of international instruments for addressing global warming and other climate changes, before it is known how each country would be affected by those changes.

Evidence Law: Admissibility Rules, Privileges, and Presumptions

Similar to the conventions of scientific research, the law of evidence sets considerable limits on the determination of facts. Among other things, the rules of evidence (and sometimes substantive legal rules) dictate that certain pieces of information are not made available to judicial fact finders. Some inadmissibility rules are based on the premise that the prejudicial effects of certain types of evidence outweigh their probative value. For example, information about a defendant's past convictions may be relevant to a determination of liability in a

given case, but may also skew decisions toward a finding of guilt. Other types of evidence, such as hearsay testimony, may be excluded due to their allegedly limited probative value (or the difficulty to assess their probative value). Finally, some inadmissibility rules stem from policy considerations that are unrelated to the probative weight of the evidence. For example, evidence obtained through illegal police practices might be deemed inadmissible in order to incentivize the police to behave appropriately in future cases and to protect the fairness of the judicial process (for a general discussion on these rules in U.S. law, see Broun et al. 2013, vol. 1:897–991, 1013–1187 and vol. 2:175–257).

Relevant information may also be kept out of the reach of judicial fact finders when the people who possess the information, or the people to whom it refers, enjoy the legal privilege to refuse to disclose this information, as in the cases of attorney–client and physician–patient relationships (Broun et al. 2013, vol. 1:527–642). Clearly, these privileges promote values that compete with the primary goals of evidence law; namely, the accuracy of judicial fact-finding and the optimal allocation of the risks of error between litigants.

Rules of burden of proof (including rebuttable legal presumptions that shift the burden from one party to the other) do not ordinarily exclude any information from the reach of the court. However, the stronger the presumption and the stricter the limitations on contradicting it, the more its effect resembles that of an exclusionary rule. For example, the *marital paternity presumption* (the presumption that the mother’s husband is the father of a child) used to be conclusive, and even today is difficult to contradict in some jurisdictions (Glennon 2000). Thus, under California law,⁵ a motion for a paternity test can only be filed within two years from the child’s birth.

Inasmuch as these doctrines effectively minimize the total sum of adjudicatory errors (or weighted errors, if some types of errors are considered more harmful than others), they are perfectly rational from a cost-benefit standpoint. Even if the exclusion of a piece of evidence or the adoption of a conclusive presumption dramatically increases the risk of judicial error *in a particular case*, it may still be true that adhering to such rules would minimize *the total sum* of errors (Berman 2004). This argument does not necessarily apply to rules of evidence that serve other purposes, each of which requires a delicate balancing of the pertinent considerations, which we cannot offer here.

Finally, it should be noted that the efficacy of inadmissibility is challenged when judicial fact finders are exposed to inadmissible evidence. This can happen during adjudication when witnesses and attorneys intentionally or inadvertently reveal the inadmissible evidence. Information from external sources, such as the media, may also be inadmissible. Subject to certain nuances, the picture emerging from numerous empirical studies is that judicial decision makers are unable to completely disregard inadmissible evidence, so it affects their decisions (Stebly et al. 2006; Wistrich et al. 2005).

⁵ See California Family Code §§ 7540, 7541.

Evidence Law: Enhancing Evidence Credibility through Blinding

Evidence law can enhance accurate fact-finding not only by depriving judicial fact finders of certain information, but also by using blinding techniques in the process of gathering and preparing the evidence (for discussion on comparable uses of blinding in other contexts, see MacCoun, this volume). Two primary examples are double-blind lineups and depriving experts of information that may bias their expert opinion.

Eyewitness identification often plays a key role in criminal (and sometimes civil) trials. However, reliable identification hinges on accurate encoding, retention, and retrieval of information, all of which are imperfect and prone to biases. Dozens of studies have demonstrated (a) that people are not very good in encoding strangers' faces and are particularly bad at identifying members of other races, (b) that memories tend to fade over time and may be contaminated (e.g., by exposure to media reports), and (c) that during the retrieval phase—often involving the use of lineups—witnesses are over-inclined to choose someone in the lineup and are influenced by (conscious or unconscious) clues given by the lineup administrator (Simon 2012:50–80; Zamir and Teichman 2018:568–572). Thus, there is a growing consensus in the forensic literature that double-blind lineups, in which the administrator does not know the identity of the suspect, can significantly increase the reliability of identification (Wells et al. 1998). However, since there is generally a trade-off between type-I and type-II errors, and since the adoption of this and comparable recommendations requires the allocation of more resources, this recommendation is far from being universally accepted (Stebly and Loftus 2013).

Comparable concerns are expressed about expert opinions that are regularly used by litigants. When experts are hired by one of the litigants or are otherwise motivated to arrive at a particular conclusion, the reliability of their opinion is compromised due (at least) to confirmation bias and motivated reasoning. Several strategies, which we cannot explore here, have been proposed to mitigate these concerns by blinding experts to the identity of the party who hires them and by depriving them, at least initially, of information that might bias their investigation (for a collection of contributions on this issue, see Robertson and Kesselheim 2016:129–220).

Anonymity and Omitted Details

The notion that some types of information may adversely affect the impartiality of decision makers extends beyond the courtroom. For example, employers who recruit new employees, scholars who review their peers' manuscripts for publication, and professors who grade their students' papers are all vulnerable to all sorts of biases and prejudices that may lead their decisions astray (see also MacCoun, this volume). As previously noted, even well-intentioned people may be unable to overcome automatic and possibly unconscious cognitive

biases. Unidirectional—and sometimes even bidirectional— anonymity and exclusion of certain bits of information (e.g., an applicant’s religion or sexual orientation) may therefore facilitate unbiased decision making. Thus, for example, the British governmental online guidelines for interviewing new employees provides a list of *protected characteristics* (e.g., age, gender reassignment, and marital status) that an employer must not ask candidates about during job interviews (GOV.UK 2018). In the same spirit, Airbnb, the global company that offers an online marketplace for lodging, changed its policy (October 22, 2018) such that potential hosts can only view the photos of potential guests *after* accepting a booking request, thereby reducing race-based and other forms of discrimination. While it may be difficult to hide some of these characteristics from human decision makers, the advent of computerized decision making provides a unique opportunity to implement such measures (see also Teichman et al., this volume).

At times, even information that conveys positive probative value may be excluded from the decision process, as illustrated by the various prohibitions on discrimination in the insurance industry. In the United States, for instance, health insurers are prohibited from considering genetic information when setting premium rates or rules of eligibility.⁶ Turner (2009:316) notes that this law may be understood as a method of *discouraging* the deliberate ignorance of the individual who would otherwise have an incentive to remain ignorant of his genetic information. Similarly, the European Court of Justice has invalidated an exemption in the equal treatment in the goods and services directive which allowed for statistically based gender discrimination in insurance premiums, ruling that it is incompatible with the objectives of the directive and with articles 21 and 23 of the European Charter of Fundamental Rights.⁷ While such prohibitions may curtail the efficiency of the market, they are arguably justified on the basis of considerations such as the protection of susceptible populations or protections of individual privacy (Avraham et al. 2014:201–221). The exclusion of information from decision makers should, however, be considered with caution, as it can have unintended and even counterproductive consequences. For example, many jurisdictions across the United States have recently adopted “ban the box” policies, prohibiting employers from enquiring about applicants’ criminal records, at least until late stages of the hiring process. These policies are meant to help people with criminal records find employment. However, employers who believe that ex-offenders are less qualified may turn to indirect ways of statistical discrimination to avoid hiring them. Indeed, both Agan and Starr (2018) and Doleac and Hansen (2016) find that these policies have harmed the job prospects of young low-skilled black and

⁶ 42 U.S. Code § 300gg–53(a), Prohibition of health discrimination on the basis of genetic information.

⁷ *Association belge des Consommateurs Test-Achats (ASBL) v. Conseil des ministres* 2011. ECJ Case C-236/09.

Hispanic men, who are more likely than other demographic groups to have a criminal record. Similar concerns arise even when the decision process is computerized (Kleinberg et al. 2018).

Ignoring and Forgetting

Forgetting is important at both the individual and the social level (see Schooler, this volume). Our ability to forget past events helps us overcome traumatic events, allows us to start over with a clean slate, and promotes our autonomy by liberating us from the shackles of the past (Mayer-Schönberger 2009:16–49). A society that never forgets limits its opportunity to receive a second chance and creates a chilling effect on its members, who know that the impact of every mistake they make will be perpetuated.

One notable way in which the law recognizes the importance of forgetting is the expungement of criminal records. Criminal records impose a significant burden on convicts attempting to reintegrate into society, as they hinder the ability to find employment, receive credit, or rent an apartment. These effects are not limited to those who have been found guilty, but extend to those who have been charged and acquitted, and even those who have only been arrested (Jacobs and Crepet 2008). To mitigate the effect of this mark of Cain, many jurisdictions allow for a process of expungement, which may mean that a criminal record is sealed, vacated, or completely destroyed. The conditions under which expungement is attainable vary, and most jurisdictions allow for a more lenient treatment of juvenile records in light of the greater emphasis on rehabilitation. In some cases, the process is automatically initiated after a certain amount of time has passed, whereas in others it may be granted only upon petition.

The implications of expungement also vary. For example, Section 651:5 of the New Hampshire Criminal Code states that “the person whose record is annulled shall be treated in all respects as if he or she had never been arrested, convicted or sentenced....” Other jurisdictions state that such a person is entitled, even under oath, to deny that the expunged incident ever occurred.⁸

The notion of forgetting and starting over with a clean slate is similarly evident in the regulation of reports about people’s financial history. Credit reports are used by creditors to determine whether, and at what interest rate, they will offer credit to a consumer. Reports may also be used by insurance agencies, property owners who wish to rent out their property, and sometimes even prospective employers. The report includes information about previous loans, debts, defaults, and bankruptcies. While this information is certainly useful, many countries set a limit, after which the information may no longer be included in the report. For example, in the United States, bankruptcies are

⁸ See Connecticut General Statutes § 54-142a(e), Erasure of criminal records; see also section 4 of the U.K. 1974 Rehabilitation of Offenders Act.

omitted from the report after ten years,⁹ and many countries in the European Union impose even stricter standards (Feretti 2008:103–121). These rules balance the utility of the data in predicting default rates against the interest of allowing individuals a clean slate.

The advent of the digital age poses a particularly difficult challenge to legal attempts to promote forgetting. While in the past forgetting was the norm and remembering the exception, digitization has reversed this state of affairs (Mayer-Schönberger 2009). Digitization—and the concomitant ability to index, store, and search huge amounts of information—has made forgetting a lot more difficult than it used to be. Particularly challenging to the goal of permitting a clean slate is the pervasive use of search engines that have perfected the capacities of indexation, search, and retrieval, thus creating an eternal memory.

The repercussions of eternal memory have led to an intense debate about the potential legal *right to be forgotten* (Leta Jones 2016). The intent of such a right is to grant people a certain amount of control over the online circulation of details about them. Depending on its scope and exact definition, such a right could possibly allow one to de-index certain results from Google, erase posts from Facebook and other social media outlets, and request the revision or removal of other online references. Underpinned by a desire to protect reputation and privacy and to allow one to shape one's own identity free from the burdens of the past, such legally induced amnesia also has serious implications for free speech and the free flow of information online. For this reason, detractors of the right to be forgotten are concerned about the creation of “black holes” of information and attempts to rewrite history (Rosen 2012). Thus, while the European Union has recently introduced its new General Data Protection Regulation, which explicitly recognizes a right to be forgotten, in the United States the First Amendment would most likely prevent similar initiatives (Larson 2013). Without delving into the normative debate, we note that while it may restrict certain forms of speech, the right to be forgotten may also facilitate expression and prevent the chilling effect of knowing that every Facebook post or tweet we make may come back to haunt us in the future.

Attorney–Client Relationships and Perjury

Sometimes the law does not necessarily encourage deliberate ignorance but nonetheless tolerates it by *not* equating deliberate ignorance with actual knowledge. For example, under the model rules for professional conduct provided by the American Bar Association (ABA), attorneys have a duty of candor toward the court that prohibits them from *knowingly* deceiving the court (ABA Rule 3.3). In this context, “knowingly” is construed as pertaining to *actual* knowledge only (ABA Rule 1.0(f); Roiphe 2011:190). This regime has cultivated a practice of deliberate ignorance amongst defense lawyers. While they

⁹ 15 U.S. Code § 1681c(a)(1), Requirements relating to information contained in consumer reports.

are not permitted to *knowingly* allow a witness to commit perjury, they may deliberately avoid information, thereby enabling their clients to present false testimony (Roiphe 2011:197). In an effort to defend this regime, it has been argued that requiring lawyers to investigate their clients' statements would undermine the attorney–client relationship and induce clients to hide information from their lawyers (Luban 1999:976–980). Arguably, however, it follows that lawyers should simply be allowed to introduce false testimony under certain conditions, as the deliberate ignorance route impairs communication, which is at the core of the attorney–client relationship (Roiphe 2011).

At any rate, outside the context of perjury, the willingness to accept deliberate ignorance on the part of lawyers is limited. Notably, the role of deliberately ignorant lawyers in the Enron scandal has led to stricter regulation of corporate lawyers. Such lawyers now have a duty to investigate suspicions of client misconduct, report it within the corporation, and in certain cases even withdraw representation and inform the Securities and Exchange Committee (Cramton et al. 2004).

The Downside of Informed Consent

The *informed consent* doctrine serves several interrelated purposes in tort law, chief of which are (a) facilitating patients' compensation for injuries they have suffered as a result of medical treatment, in incidences where, but for the lack of informed consent, the treatment was adequate and involved no negligence (Peck 1984; Raab 2004); and (b) protecting patients' autonomy (Jones 1990). While contemporary discussions highlight patient autonomy, the former, traditional purpose still plays a major role in practice, as courts strive to provide relief to unfortunate patients who have incurred bodily injury, irrespective of physician negligence (Brennan et al. 1996). To fulfill the former purpose, but not the latter, patients must establish that had they received the relevant information, they would have refused to undergo the treatment, thereby avoiding the resulting injury (Maclean 2009:183–188). Patients' autonomy is compromised whether or not they would have consented to the treatment, so it does not hinge on such causality (although the *magnitude of the harm* to autonomy may hinge on it).

Neither of the two purposes *necessarily* assumes that patients actually *wish* to be fully informed about their condition and the potential prospects and risks of the proposed treatment. Arguably, respect for patients' autonomy requires that they receive all the relevant information, irrespective of whether they wish to receive it (Ost 1984; Turner 2009:347), and once patients are fully informed, they might refuse to undergo a given treatment even if they initially preferred not to receive the relevant information (for further discussion of the right not to know, see Berkman, this volume).

Nonetheless, the doctrine of informed consent *is* associated with the assumption that people wish to be informed about their medical condition, the

available treatments, and the attendant prospects and risks involved. If patients typically prefer that their physician make the decision for them, and agree to follow the physician's advice whatever information they receive, then the causal link between not getting the relevant information and consenting to undergo the relevant treatment is severed. Arguably, this would be the case even if a minority of patients do not share this typical preference, as long as the plaintiff fails to prove that s/he belongs to that minority. As for patient autonomy, it plausibly requires that the patient's (explicit and perhaps even implicit) choice not to receive information is respected as well (Andorno 2004; but see Harris and Keywood 2001).

Contrary to much of the legal discourse, the available empirical data shows that patients do typically prefer to follow their physician's advice, rather than make the necessary decisions by themselves (Schneider 1998:35–46). And while patients' desire for information is much stronger than their desire to make decisions (e.g., Ende et al. 1989), many also prefer not to know at least some aspects of their medical condition, prognosis, and the risks involved in a given treatment (e.g., on the preferences of elderly cancer patients, Elkin et al. 2007; on patients' desire to get information about various aspects of a surgery, Asehnoune et al. 2000; on the aversion of some cancer patients to receive technical information before major cancer surgeries, McNair et al 2016:261 and Schneider 1998:110–111). Insofar as this is true, there is a tension between the legal norms that strongly incentivize physicians to provide as much information as the patient can reasonably comprehend and patients' common preference (which is not necessarily irrational) to remain uninformed.

This tension might be eliminated by entitling patients to damages for injuries arising from medical treatment on a no-fault basis. Under such a regime, the law would not have to frustrate patients' options to remain ignorant and still compensate them for their loss and suffering in appropriate cases. Physicians would still be liable for violating the duty to provide patients with relevant information, but only in instances where patients do wish to be informed (or at least do not wish to remain uninformed). Of course, moving from fault-based to nonfault liability for medical injuries, and delineating the scope of such liability, would require careful consideration of a multitude of factors that lie beyond the scope of the present discussion (Studdert and Brennan 2001; Weiler 1993).

Fostering Settlements and Plea Bargains

People's decisions are influenced by the anticipated feeling of regret; that is, by the expectation that if it transpires that they have made the wrong choice, they would experience regret. The anticipation of regret depends on what one expects to know *ex post*. The decision maker may expect to know the outcomes of both the chosen option and the forgone one(s) (*full knowledge*), or only those of the chosen option (*partial knowledge*) (Ritov and Baron 1995).

Sometimes, one option entails full knowledge while the other involves only partial knowledge. In such cases, the latter option is typically more attractive to regret-averse people, because not knowing the outcome of the forgone option largely shields one from the anticipated regret. This is a common explanation for the pervasiveness of settlements in civil proceedings and of plea bargaining in criminal ones (Zamir and Teichman 2018:505–507, 522). Inasmuch as the legal system strives to encourage settlements and plea bargains, courts should make sure not to reveal how they would have decided the case in the absence of a settlement or a plea bargain.

Conclusion

In this chapter, we have surveyed ways in which the law overcomes some instances of deliberate ignorance and fosters others. Along with major issues, such as the doctrine of willful blindness and institutional veils of ignorance, we touched upon more specific and even peculiar examples. In addition to describing existing legal norms, we highlighted cases in which the law should arguably combat deliberate ignorance more effectively or, conversely, facilitate more of it. Given the limited scope of the paper, it should primarily be taken as an invitation to further discuss these theoretically fascinating and practically important issues.

Acknowledgments

We are grateful to Christoph Engel, David Hagmann, Ori Herstein, Daphna Lewinsohn-Zamir, Ofer Malcai, Hanan Raviv, Nora Szech, Sonja Utz, and Limor Zer-Gutman for helpful comments, and to Inbal Elbaz for excellent research assistance. This research was supported by the I-CORE Program of the Planning and Budgeting Committee and the Israel Science Foundation (Grant No. 1821/12).