Thobani S, 'The "Consent Illusion": Where is the Fallacy?' (2025) Special Issue 1 2025 EJPLT.
DOI: https://doi.org/10.57230/EJPLTSI251ST

EJPLT
ISSN 2704-8012



The 'Consent Illusion': Where is the Fallacy?

L' "illusione del consenso": dov'è l'equivoco?

SHAIRA THOBANI
Associate Professor of Private Law
University of Turin

Abstract

The article explores the so-called "privacy fallacy" examined in Ignacio Cofone's book, which concerns the unsuitability of contract rules to protect privacy values and to prevent privacy harms, and analyses the potential implications of the solutions proposed by the author within the framework of the European Union. After a brief outline of the general regulatory paradigms of data governance and the role of consent, the article examines the extent to which consent can be considered an illusion in light of divergent regulatory approaches, and, finally, addresses the challenges associated with entirely dispensing with consent.

L'articolo analizza la cosiddetta "privacy fallacy" esaminata nel libro di Ignacio Cofone, riguardante l'inadeguatezza del consenso nel proteggere la privacy e prevenire i danni, e analizza le potenziali implicazioni delle soluzioni proposte dall'autore nel contesto dell'Unione Europea. Dopo una breve panoramica sui diversi modelli regolatori della governance dei dati e il ruolo del consenso, l'articolo esamina in quale misura il consenso possa essere considerato un'illusione alla luce di diversi approcci normativi, e, infine, affronta le sfide associate alla completa eliminazione del consenso come base legittimante il trattamento.

© The author(s) 2025, published by Suor Orsola Benincasa Università Editrice.
This contribution is licensed under a Creative Common Attribution 4.0 International Licence CC-BY-NC-ND, all the details on the license are available at:
https://creativecommons.org/licenses/by-nc-nd/4.0/

Keywords: data protection; consent; privacy

Summary: 1. Introduction. – 2. Regulatory paradigms of data governance. – 3. The Consent Illusion. – 4. The Illusion of Doing Without Consent. – 5. Conclusions.

1. Introduction.

The "privacy fallacy" to which the book refers to consists of two false assertions: the first one is that contract rules are suitable to protect privacy values and to avoid privacy harms; the second one is that, in order to protect the value of privacy, only tangible consequences of privacy violations need be avoided. We will concentrate here on the first fallacy, the so-called "consent illusion".

Starting from the premise that data subjects are not able to take meaningful decisions with regard to data, the main aim of the book is to design a system that does not allow those who exploit data to be able to do so on the basis of the consent of data subjects. The focus is not to prevent data exploitation per se, but only to avoid harmful data practices. To this end, corporations shall be prohibited from relying on consent when such consent serves as the basis for exploitative practices that result in harm to the data subjects. Thus, the main goal is to prevent corporations from causing privacy harm for which they are not held accountable because they can legitimately rely on the justification that data subjects have somehow waived their right to complain by giving consent – thereby shifting the cost of data practices onto the individuals themselves. The most problematic aspect of data processing is therefore identified as the privacy harms it may produce, and the book proposes corporate accountability as the appropriate remedy.

As the book is written from a US perspective (albeit with insightful references to other legal systems and rules, such as the EU General Data Protection Regulation, "GDPR"), it may be worthwhile to examine the potential implications of such an approach within the framework of the European Union.

After a brief outline on the general regulatory paradigms of data governance and the role of consent, we will focus on the extent to which consent can be considered an illusion in light of the differing regulatory backgrounds, and finally on the challenges associated with entirely dispensing with consent.

2. Regulatory Paradigms of Data Governance.

Before addressing the issue of the "consent illusion", and in order to better understand the rules on consent within a broader context, it is necessary to briefly outline the main regulatory paradigms of data governance. Data governance broadly refers to the rules concerning data access, control, sharing use and re-use, with the aim of determining who is entitled to extract value out of data, what safeguards should be adopted to prevent harm, what interests

deserve protection, and how such protection should be implemented. To this end, the legislator is required to make normative choices.

With regard to the circulation of data, one end of the regulatory spectrum entails full recognition of data subjects' private autonomy, meaning that data can circulate only with the data subject's consent. In theory, this model places utmost importance on the role and interests of data subjects. In practice, however, the actual effectiveness of such recognition depends on the conditions under which consent is given. It is well established that granting someone the power to choose can serve as a tool of empowerment only if that person is able to make fully informed and voluntary decisions; otherwise, it risks becoming a source of oppression, particularly to the detriment of the weaker party - who may give consent unknowingly, without understanding its implications, or simply because no alternative is available.[1]

At the opposite end of the spectrum lies the complete absence of recognition of data subjects' autonomy in deciding what happens with their data. The problem then becomes: who decides? Several answers are possible. A first approach is full liberalization, which implies that everyone can decide – meaning that any actor is permitted to collect and process data without prior restrictions or specific requirements. In this scenario, the only means of protecting the various interests at stake is to provide for ex post remedies, such as civil liability.

An alternative approach is that of fully controlled circulation, whereby data may be collected and processed only in the cases expressly provided for by the legislator. In this model, without necessarily excluding ex post remedies, the protection of the interests involved in data processing in ensured ex ante, through restrictions to data collection and use – thereby shaping processing practices from the outset. Clearly, numerous intermediate solutions exist. For instance, GDPR recognizes the relevance of data subject consent, but also provides for other legitimate bases for processing, each of them broader or narrower according to the relevance assigned to competing interests.

After defining the role of private autonomy and identifying which interests deserve protection, the legislator can choose and combine a wide array of instruments to pursue its objectives.[2]

A first option is to implement command-and-control mechanisms that impose prohibitions and obligations aimed at achieving specific behaviors or outcomes. For instance, certain entities may be prohibited from processing data, or there could be a general ban on particular types of processing or on the collection of specific categories of data. Naturally, employing command-and-control tools in the regulation of data governance entails the adoption of mandatory (as opposed to default) rules – i.e. rules that cannot be altered by agreement between the parties – and therefore establishes protection that cannot be waived by data subjects.[3] Moreover, as the book points out, relying exclusively on this system to prevent harm risks proving highly ineffective

[1] In general terms, see M Fabre-Magnan, *L'institution de la liberté* (PUF, 2018).

[2] For an overview on ex post and ex ante remedies, see S Romanosky, A Acquisti, 'Privacy Costs and Personal Data Protection: Economic and Legal Perspectives' (2009) 24(3) Berkeley Tech. L.J. 1061-1102.

[3] See D Hirsch, 'From individual Control to Social Protection: New Paradigms for Privacy Law in the Age of Predictive Analytics' (2020) 79(2) Md. L. Rev., 439-505

when the rules, albeit mandatory, are merely procedural, as they could easily be complied with by simply ticking boxes. Procedural rules do not impose substantive prohibitions or obligations, but rather set formal requirements: for instance, a substantive rule would be one that prohibits high-risk processing activities, whereas a procedural rule would be the one that requires to keep a data processing register.

A second option is to adopt market-based instruments which do not directly prohibit specific activities or results, but instead seek to render them economically unfeasible or unprofitable. Several types of market-based instruments are aimed at harm prevention, and some include civil liability among them, as it compels those who cause harm to internalize the negative externalities of their actions by compensating the victims. If the costs of damages exceed the potential benefits, harmful activities will naturally be excluded from the market. Once again, the effectiveness of this approach in the context of data processing depends on the extent to which data subjects are allowed to exercise private autonomy by waiving in advance their right to compensation. The more it is deemed legitimate for individuals to give valid prior consent to the losses and harms they may suffer as a result of the collection and using of their data, the less effective this system becomes – especially in situations in which data subjects are not in a position to provide meaningful consent.

Command-and-control instruments and civil liability can, of course, be used in combination through a mixed system of public and private enforcement, as the book suggests, by providing both substantive prohibitions aimed at reducing risk and obligations to compensate when risk materializes.

The choices regarding data governance and the tools adopted to implement such choices interact in various ways, resulting in different regulatory scenarios. Systems that rely on data subjects' private autonomy for data governance may differ significantly depending on the existence of substantive prohibitions concerning the conditions under which consent is considered valid, and on what cannot be done with data regardless of consent. Indeed, if data may be processed only with the data subject's consent, but obtaining valid consent proves very difficult due to stringent requirements – and, in any case, even if consent is obtained, the controller is still prohibited from using the data in certain ways or for certain purposes – then the principle of autonomy clearly loses much of its practical relevance and should not be overemphasized. Another key differentiating factor among systems based on private autonomy is the type of remedy available when the data are processed without valid consent: one thing is an ex ante remedy, such as a blanket prohibition on processing data without consent; another is to provide ex post remedies, such as compensation for damages, but only when and if harm occurs.

Similarly, systems that provide for full liberalization of data circulation (and therefore do not require consent or any other legitimate basis to process data) may also vary substantially, depending on whether they impose substantive ex ante prohibitions on how data may be used, or, conversely, allow corporations to use data freely as long as they provide compensation in the event of harm. In the same way, the significance and role of civil liability in protecting against harm depend on the rules governing compensation and procedure, as the book

underlines.

In other words, the effectiveness of each model lies in the details.

3. The Consent Illusion.

Against this background, it is time to focus on consent.

The main assumption of the "consent illusion" fallacy is that the incapacity of data subjects to make informed and deliberate decisions regarding the management of their personal data (a phenomenon that is well documented)[4] does not stem from a lack of interest, but rather from the inability of to make a genuine choice due to the uncertainty surrounding the risks of data processing. In other words, individuals do care about their data and they wish to avoid privacy harms, but they are hindered in making coherent decisions by the opacity of the consequences of data, which they cannot fully grasp. This vulnerability is then exploited by data collectors, who seek to obtain consent through various manipulative strategies. Under such conditions, reliance on individual control as a means of protecting privacy appears to be an ill-suited solution.

The book identifies several reasons why data subjects are unable to make genuine choices concerning their personal data.

A first set of reasons relates to the limitations of data subjects' rationality in understanding the implications of data processing. As the book underlines, the inability to fully comprehend what will happen with their data is both subjective (stemming from cognitive constraints on the part of data subjects) and objective (arising from the fact that how data will be used is not only unknown to the data subjects but in many cases unknowable to anyone at all).

As regards to the subjective impossibility to be informed in a meaningful and effective way, not only are data subjects' cognitive abilities constrained by bounded rationality, but also, in order to fully understand privacy policies and assess the value of their data, the average data subject would need to possess adequate digital literacy covering all problematic aspects of data protection – and be willing to apply this knowledge repeatedly throughout the day whenever their consent is requested. This is, quite clearly, not the case. As a result, data subjects are particularly susceptible to nudging and manipulative design choices implemented by those seeking to exploit data.

It is important to emphasize that the limitations on the conscious agency of data subjects apply regardless of the decision's outcome: in theory, individuals may be nudged either to give or to deny consent. Consider a system in which the legislator mandates that the only way to request consent is by offering data subjects a choice between a large, brightly colored button stating "No, I refuse"

[4] A Acquisti, J Grossklags, 'Privacy and Rationality in individual Decision Making' (2005) 3(1) IEEE Secur. Priv., 26-33; DJ Solove, 'Privacy Self-Management and the Consent Dilemma' (2013) 126 Harv. L. Rev. 1880-1903; S Kokolakis, 'Privacy Attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon' (2017) 64 Comput. Secur., 122-134; L Gatt, R Montanari, IA Caggiano, 'Consenso al trattamento dei dati personali e analisi giuridico-comportamentale. Spunti di riflessione sull'effettività della tutela dei dati personali' (2017) 2(2) Politica del diritto, 337-353; N Gerber, P Gerber, M Volkamer, 'Explaining the privacy paradox: A systematic review of literature investigating privacy attitude and behavior' (2018) 77 Comput. Secur., 226-261.

and a smaller, less visible one stating "Yes, I accept", while any attempt to circumvent the choice (e.g., by scrolling down in the hope that the banner disappears) is interpreted as a refusal – and, crucially, no adverse consequences follow from a denial of consent. This too constitutes nudging. However, this form of nudging (i.e., steering individuals toward refusing consent) is rarely taken into account in legal discourse on the validity of consent and is not generally regarded as problematic (for example, it would be considered entirely valid under the GDPR). This is because concern about the quality of consent tends to arise only when it enables data processing, not when it prevents it – even though in both scenarios the decision may equally lack genuineness. Indeed, if individuals do not understand what will be done with their data and what the consequences may be, this lack of awareness affects their decision-making irrespective of whether they say yes or no. Hence, the decisive factor appears not to be the (lack of) informed awareness of the data subjects, but rather whether their decision results in greater or lesser volumes of data being processed: in the former case, uninformed consent is perceived as problematic, in the latter it is not. This reveals a crucial aspect of data protection governance: the assessment on whether consent is valid, when based on the data subjects' subjective impossibility to be informed, may vary depending on the broader normative stance toward data processing – that is, whether the system favors more or less collection and use of data. Where there is a general suspicion toward data processing, uninformed decisions will render consent invalid only if the data subject is nudged to consent. Conversely, the law will fictionally deem the decision genuine if the data subject is nudged to refuse.[5]

The problem however goes beyond nudging: if individuals are incapable of making informed decisions regarding their personal data, this incapacity persists even the absence of external pressures or manipulation. Thus, even in neutral conditions, when no nudging is applied and the individual is called upon to decide autonomously, the resulting choice will likely be highly random. This further reinforces criticisms against a system that bases data processing on the principle of private autonomy, highlighting the impossibility of addressing the issue without first taking a stance on the broader regulatory approach to data processing – namely, whether it ought to be promoted or restrained.

As regards the objective impossibility for data subjects to be informed, the book underscores that the inability to make genuine decisions also derives from the fact that not even data controllers can foresee all the consequences – and therefore the risks – of data processing. This implies that, even if data subjects were both willing and capable of fully reading and understanding privacy policies, it would be objectively impossible to provide them with complete and accurate information about the use of their data and the corresponding risks, many of which remain unknowable. This is another argument advanced in the book to challenge the legitimacy of relying on private autonomy in the context of data processing: data subjects cannot objectively know what they are consenting to.

However, completely excluding private autonomy is not the only regulatory

[5] From a technical standpoint, this derives from the fact that the GDPR requires the existence of a valid legal basis (in this case, consent) for carrying out data processing, but not for refraining from processing.

option. The collection of data for unknowable and unpredictable purposes is a human activity that the legislator can regulate – or even prohibit. In other words, the fact that data may be processed in unknowable ways is not an inherent necessity, but the outcome of a regulatory framework that can be adjusted. For instance, if processing for purposes unforeseeable at the time of collection were prohibited – which aligns with a restrictive interpretation of the GDPR's purpose limitation principle, whereby personal data may only be processed for purposes that the controller can specify in clear and detailed terms at the moment of collection – the problem of unknowable risk would be mitigated. These risks are not natural givens, but rather the result of broader decisions concerning data governance, which extend beyond the narrow issue of consent. Hence, the presence of unknowable risks – the so-called "objective impossibility" of obtaining genuine consent – derives from regulatory choices, not from natural constraints.

The point becomes especially evident in the context of artificial intelligence systems and the so-called black box problem. It is often impossible, even for developers of self-learning algorithms, to understand why certain outputs are generated based on certain inputs. If such opacity makes it impossible to foresee consequences, a regulatory choice must be made: either to allow AI systems that produce non-explainable outputs or to prohibit them. This is a political decision involving competing interests, including the public interest in fostering technological innovation. Regulatory responses do not need to be binary, and intermediate solutions are possible – for instance, prohibiting such systems only when the outputs are used to make decisions that significantly affect individuals, whether they are the subjects themselves or third parties. Once again, the role and extent of private autonomy must be assessed within the broader regulatory context and cannot be considered in isolation.

Thus far, the focus has been on issues of rationality. There is, however, a further reason why reliance on consent may not always be appropriate. As is widely recognized, data processing – especially when conducted on a large scale – may have consequences extending far beyond the individuals whose data are processed and who have given consent.[6] For example, discriminatory outputs resulting from data processing can affect individuals who never consented to such processing. Moreover, broader collective interests may be at stake: targeted political advertising can undermine democratic processing, and filter bubble may contribute to the radicalization of social groups. In these cases, the problem is not one of rationality but of entitlement. Even if individuals could fully comprehend the consequences of data processing, their understanding would not be sufficient to legitimize their decision, because the impact extends beyond their own personal interests. Accordingly, data subjects cannot be entitled to make individual decisions that have significant external effects on others or on society as a whole. Naturally, no issue of entitlement arises where the primary consequences of processing involve only the

[6] A Mantelero, 'Personal data for decisional purposes in the age of analytics: From and individual perspective to a collective dimension of data protection' (2016) 32(2) CLSR 238-255; B Van Der Sloot, 'The Quality of Life: Protecting Non-Personal Interests and Non-Personal Data in the Age of Big Data' (2021) 29(5) Eur. Rev. Priv. L.,757-784; ML Montagnani, M Verstraete, 'What Makes Data Personal?' (2023) 56(3) UC Davis Law Review, 1165-1232.

individual providing consent. For instance, if a hospital asks for access to a patient's data in order to review the results of prior medical tests, the patient's interests are the only ones affected – provided that the data are used exclusively for the purposes of delivering medical care. The same applies when a person authorizes the publication of their photograph, again on the condition that the picture is not repurposed for unrelated uses.

4. The Illusion of doing without consent.

The aforementioned limitations of consent lead the book to argue that it should have no relevance, as the control and empowerment it provides do data subjects are merely illusory. However, as we shall see, it is not easy to dispense with consent altogether: indeed, even in a system as the one proposed, consent re-emerges through the back door.

As previously mentioned, the book advances a system of corporate accountability whereby corporations are held liable for harm (understood in a broad sense that goes beyond material consequences) resulting from the data processing activities they undertake. In this view, civil liability and damage compensation are deemed to be the most effective tools (alongside public enforcement) to address the problematic consequences of data processing. What role remains for consent and individual choices in such a scenario?

In the first place, it follows from its very definition that the effectiveness of any instrument of private enforcement, such as civil liability, entirely relies on individuals' choice to seek compensation. However, individuals will only pursue compensation if they are not only aware that they have suffered harm, but also if they are aware that this harm may be related to the processing of their personal data. Considering the very low level of awareness and data literacy among data subjects (which is at the root of the critique to private autonomy), it will not be that easy, in practice, for both of these conditions to be satisfied. Individuals will typically seek compensation only in cases where data processing infringes upon very personal interests that they are capable of perceiving – such as the disclosure of intimate or false information (as illustrated in many real-life examples cited in the book, which would in any case already be addressed by ordinary liability rules in many jurisdictions) or the persistent annoyance of marketing calls. Many other privacy violations – even if, in theory, they could give rise to compensation – will probably remain outside the awareness of data subjects. And, in a liability-based system, anything that is not perceived by the victims is effectively irrelevant, since, as noted, the system's effectiveness depends on individuals recognizing that they have suffered harm. Private autonomy and individual choices thus resurface in full force – not to waive rights in advance (as with consent to data processing), but to activate remedies ex post. Procedural rules can certainly help (as the book underlines), but they do not overcome the fact that even joining a class action requires a decision that may be impaired by the same lack of awareness affecting the validity of consent.

In other words, liability only compensates harm that data subjects recognize. According to the book, this is a better solution than relying on ex ante consent,

since, once the harm has occurred, it is arguably easier for individuals to recognize it than it is to foresee it at the time of giving consent – given the impossibility of grasping the full consequences of data processing beforehand. However, it follows from the above that this largely depends on the requirements set by the legislator for valid consent and on the extent to which such requirements can overcome the "subjective" and "objective impossibility" arguments. As previously mentioned, the "objective impossibility" argument does not hold if the legislator prohibits data processing in cases where it is not possible to determine its consequences. The "subjective impossibility" argument can also be substantially mitigated by establishing strict consent requirements that prevent manipulative tactics, vagueness, and power imbalances – at least when the processing involves information whose misuse affects deeply personal interests of data subjects. For example, if, before sending a picture to someone, I am clearly informed that the recipient will share my picture online, I will likely pay close attention; similarly, if I am explicitly asked whether my browsing behavior may be used to assess my creditworthiness – potentially leading to loan denial – the average data subject is likely to refuse consent (the issue being that such information is rarely provided so clearly).

These situations largely correspond to those in which individual become aware of harm and recognize its link to the processing of their data – i.e. the cases in which they are likely to seek damages. In other words, depending on how consent is regulated, even consent-based systems may not prevent individuals from consciously granting or refusing consent when the consequences of the processing are sufficiently clear and appreciable – the very same cases in which they would pursue compensation after harm occurs. It remains to be demonstrated that relying on data subjects to sue data controllers is more effective in preventing privacy harm than relying on public authorities to enforce strict consent requirements. On the contrary, if consent is considered valid even when it refers to unknowable consequences, and it is permissible to obtain it through manipulative techniques, then it clearly follows that a consent-based system cannot adequately protect data subjects.

Even when stringent consent requirements are in place, the issue of subjective impossibility persists in cases where data processing does not implicate highly personal interests of data subjects. In such circumstance, individuals are often unable to fully comprehend the consequences of the processing, even if adequately informed. This occurs particularly when the processing has no immediate impact on data subjects, or when its effects pertain to broader collective interests, such as systemic risks. It must be emphasized, however, that this issue is not resolved by replacing a consent-based model with one grounded in liability. Both frameworks are inadequate in addressing systemic risks that affect general interests, due to an underlying entitlement issue that renders it problematic to rely on individual remedies-such as consent or liability – that presuppose individual initiative. In such cases, as the book suggests, recourse to public enforcement mechanisms becomes necessary.

A second way in which consent re-emerges indirectly is when it is employed – as the book proposes – as a criterion for assessing whether a privacy violation

constitutes harm. In certain contexts, it is indeed undeniable that consent can legitimize specific information flows. For instance, if an individual shares intimate photographs, they cannot subsequently object to that person accessing the material (although they may, of course, object if it is disclosed to third parties without their authorization). Similarly, if one subscribes to a newsletter that clearly and transparently describes and promotes itself as delivering targeted content based on data collected from a specific source, one cannot reasonably object to profiling practices underpinning that content. In such cases, from a liability perspective, individual consent shifts from being a condition for the lawfulness of the processing to serving as a criterion for determining the harmfulness of the processing activity.

It would indeed be difficult to deny the relevance of consent when decisions concerning one's intimate sphere are at stake. Only the data subject is entitled to determine which intimate information may be disclosed, and to whom. In this regard, consent operates not as contractual consent, but rather as a form of defence within the framework of civil liability. While the book associates the fallacy of the so-called "consent illusion" with the predominance of contractual theory in the governance of data flows, this is not a universal paradigm. In continental Europe, for example, consent to data processing Is rooted in the theory of fundamental rights. Within this framework, consent is not conceived as a mechanism to bind individuals, but rather as a safeguard to ensure that no one interferes with an individual's private sphere without their authorization.

To assess the true divergence between the role of individual choice in a consent-based model and that in a liability-based one, it is therefore necessary to consider the extent to which the criteria for valid contractual consent differ from those required for consent as a valid defence. In systems where the threshold for contractual consent is significantly lower than that required for consent as a defence, the liability model offers greater protection for data subjects. Conversely, where the criteria for valid consent in both frameworks are essentially equivalent, the gap between the two models is considerably reduced.

5. Conclusions.

It is undeniable that, in the majority of personal data collection practices observed in daily life, consent often represents a mere illusion. However, in order to fully comprehend the role of consent, it is not sufficient to determine whether it can serve as a lawful basis for processing. Rather, it is necessary to consider the broader regulatory framework, including the legal requirements for valid consent and the substantive prohibitions established by the legislator concerning data processing activities.

The book appears to assume a particularly expansive interpretation of consent and to adopt a notion of personal data as a form of property that can be disposed of once and for all,[7] an approach which may not be universally accepted or applicable across all jurisdictions. Accordingly, it implies (linking

[7] On the issue, see PM Schwarz, 'Property, Privacy and Personal Data' (2004) 117(7) Harv. L. Rev., 2055-2128; SH Hazel, 'Personal Data as Property' (2020) 70(4) Syracuse L. Rev., 1055-1113.

this to the concept of private autonomy) that once consent is given, e.g. for the publication of a picture in a specific context and for a specific purpose, it is thereby lawful to reuse that information for any other purpose. Such an interpretation, however, may contradict the very notion of individual control and, in certain systems, may fall outside the scope of contractual autonomy. It certainly exceeds the boundaries of what constitutes valid consent in the EU under GDPR.

As previously observed, when the legal requirements for valid consent are sufficiently stringent to render the acquisition of such consent difficult, the overall viability of a consent-based system must be re-evaluated, as the role of consent is likely to diminish significantly. This is precisely the case within the EU legal framework. Initially, consent was regarded by many data controllers as the most secure basis for data processing, particularly given the uncertainties surrounding alternative lawful grounds such as legitimate interest. This led to an overreliance on consent requests, even in situations where such requests were unnecessary. Subsequently, as it became evident that the criteria for valid consent – previously not given sufficient consideration – were being interpreted with increasing strictness by data protection authorities, controllers began to exercise greater caution in seeking consent and turned instead to alternative legal bases.[8]

Furthermore, it is important to note that the GDPR is fundamentally constructed to prevent harm through ex ante regulatory mechanisms. These include a regime of controlled data circulation – where processing is lawful only if one of the legal bases under Article 6 is satisfied – and the principle of accountability principle, according to which the controller is responsible for, and must be able to demonstrate compliance with, the rules governing data processing (Article 5(2)). Moreover, GDPR and other European legal instruments contain both general and specific prohibitions against processing operations that pose serious risks to the fundamental rights and freedoms of natural persons, irrespective of the presence or absence of consent.[9] While the system is far from flawless, eliminating consent altogether would entail removing one of the key legal bases for data processing, thereby rendering the structure incomplete. This entails moving from a semi-controlled circulation regime to a semi-liberalized model for data circulation predominantly based on ex post liability, which would constitute a substantial regression in the level of data protection currently ensured by the EU framework – although it could be regarded as an advancement in protection within other legal systems where data protection is conceived differently.

Additionally, in a system where obtaining valid consent is particularly

[8] Emblematic is the case decided by the EUCJ, 4th July 2023, C-252/21, in which Meta sought to rely on any alternative legal basis for processing in order to avoid grounding the activity on consent.

[9] For a general prohibition of high-risk processing activities, see Article 36 of the GDPR. With regard to automated decision-making, Article 22 GDPR provides for specific safeguards and remedies, even where the data subject has given consent. For more specific prohibitions, see, for example, the provisions prohibiting the processing of minors' data and special categories of data for profiling and marketing purposes (Articles 26(3) and 28(2) of Regulation (EU) 2022/2065). See also the rules addressing the downstream risks resulting from broader data flows—namely, those that do not regulate data processing per se, but rather directly prohibit or restrict certain activities that may derive from such processing (e.g., Regulation (EU) 2024/1689 on Artificial Intelligence and Regulation (EU) 2024/900 on the transparency and targeting of political advertising).

challenging, there already exists considerable space for liability: provided that harm has been suffered and can be demonstrated, compensation may be sought in all cases where processing has occurred without a legal valid basis. Moreover, due to the principle of accountability a significant portion of the burden of proof is shifted to the controller. The core difficulty in liability claims for unlawful data processing in the EU does not typically lie in the assertion that data subjects provided consent (as is often invalid), but rather in the judicial finding that no harm – or significant harm – has occurred. This issue, however, pertains to the second "privacy fallacy" discussed in the book, which will not be addressed in this context.

Is consent an illusion? In many instances, it undoubtedly is. Nevertheless, an outright dismissal of consent appears not only to constitute an illusion in itself, but also to represent a counterproductive response within a regulatory framework that, arguably, does not too seriously rely on it. It is therefore essential to carefully assess where exactly the fallacy lies in the so-called "consent illusion".