

The Future of Privacy Law?

A Comment on Solove/Hartzog's 'Kafka in the Age of AI and the Futility of Privacy as Control'

by Mark Fenwick, Paulius Jurcys, Timo Minssen & Marcelo Corrales Compagnucci *

Abstract: This Comment engages with Daniel Solove and Woodrow Hartzog's thought-provoking claim that 'privacy as control' is increasingly 'futile' in an era of ubiquitous AI. We share their concern about the profound power imbalances and structural opacity that characterize today's data-driven systems, and we recognize the urgency of rethinking traditional privacy frameworks considering these challenges. At the same time, we respectfully suggest that their critique may understate the potential of a reimagined, more robust vision of individual control within privacy law.

We raise three considerations in support of this view. First, the prevailing model of individual control they critique does not fully reflect the richer, human-centric approach to personal data that law can (and should) aspire to. Second, casting individual and structural approaches as oppositional creates a false dichotomy. In

practice, these dimensions are interdependent and mutually reinforcing. Effective privacy governance requires both empowered individuals and robust structural safeguards to establish a data ecosystem that genuinely serves the interests of individuals and the public. And third, rather than endorsing resignation or fatalism, Kafka's work can also be read as a call to reclaim dignity and agency in the face of bureaucratic and technological opacity.

We therefore propose a revitalized framework for privacy law that affirms both meaningfully supported and feasible personal agency and strong institutional safeguards. By integrating these complementary dimensions, we aim to contribute to a constructive and forward-looking dialogue on how privacy can endure as a viable and principled right in an increasingly complex and algorithmically mediated world.

Keywords: Privacy as Control; Individual Control Model; Societal Structure Model; Private-by-Default; Human-Centric Data Ecosystem; Data Ownership; Data Portability

© 2026 Mark Fenwick, Paulius Jurcys, Timo Minssen & Marcelo Corrales Compagnucci

Everybody may disseminate this article by electronic means and make it available for download under the terms and conditions of the Digital Peer Publishing Licence (DPPL). A copy of the license text may be obtained at <http://nbn-resolving.de/urn:nbn:de:0009-dppl-v3-en8>.

Recommended citation: Mark Fenwick, Paulius Jurcys, Timo Minssen & Marcelo Corrales Compagnucci, The Future of Privacy Law? A Comment on Solove/Hartzog's 'Kafka in the Age of AI and the Futility of Privacy as Control', 17 (2026) JIPITEC 216 para 1.

A. Introduction

- 1 Daniel Solove and Woodrow Hartzog's recent essay, *Kafka in the Age of AI and the Futility of Privacy as Control*,¹ presents a bracing challenge to one of the

foundational tenets of modern privacy law. In an era defined by ubiquitous data collection and the accelerating deployment of artificial intelligence (AI), they argue, the prevailing 'privacy as control' paradigm, centered on individual consent and self-management of personal data, has become

* Mark Fenwick: Professor of International Business Law, Graduate School of Law, Kyushu University, Fukuoka, Japan.

Paulius Jurcys: Senior Lecturer, Faculty of Law, Vilnius University, Lithuania; School of Law, University of California, Berkeley, USA; Co-Founder, Prifina.

Timo Minssen: Professor of Law and Director of the Centre for Advanced Studies in Bioscience Innovation Law (CeBIL), Faculty of Law, University of Copenhagen, Denmark.

Marcelo Corrales Compagnucci: Associate Professor and Associate Director, Centre for Advanced Studies in Bioscience Innovation Law (CeBIL), Faculty of Law, University of Copenhagen, Denmark; Inter-CeBIL Research Affiliate, Petrie-Flom Center, Harvard Law School.

1 Daniel J Solove and Woodrow Hartzog, 'Kafka in the Age of AI and the Futility of Privacy as Control' (2024) 104 BU L Rev 1021.

inadequate against pervasive surveillance and fundamentally ‘futile.’

- 2 Drawing on Franz Kafka’s bleak allegories of oppressive systems of bureaucratic power that obscure, overwhelm, and ultimately disempower the individual, Solove and Hartzog contend that digital life today mirrors this condition of alienation.² Individuals routinely interact with opaque algorithmic systems, unknowingly cede data under the pretense of informed choice, and internalize the blame when harms inevitably follow.³ They refer to Kafka’s *The Trial* as offering a powerful metaphor for contemporary privacy’s impotence: affording individuals formal rights is futile when the broader system remains structurally impenetrable. Thus, they call for a paradigm shift: from a rights-based ‘Individual Control Model’ to a ‘Societal Structure Model’ that imposes direct legal constraints on institutional data practices, sidestepping the illusory autonomy of the individual.⁴
- 3 This comment engages Solove and Hartzog’s thesis with deep appreciation for its urgency and analytical acuity. Their diagnosis of the Kafkaesque dimensions of today’s data ecosystem is, in many respects, unassailable. Yet, this comment raises three rejoinders to their conclusion that individual empowerment is a dead end. First, the model of individual control that Solove and Hartzog critique is not the ideal to which privacy law should aspire, but a diluted and instrumentalized version of it. When properly articulated and institutionally supported, individual data rights can be far more meaningful than the feeble versions we see today. Second, the authors’ framing sets up a false dichotomy between individual and structural approaches. In practice, these are mutually reinforcing. Structural reform is indispensable; however, it need not come at the expense of individual dignity and agency. Third, Kafka’s critique of oppressive social systems should not be read as a call to abandon personal empowerment; rather, it should be understood as a challenge to redesign institutions in ways that restore intelligibility, voice, and moral accountability to those enmeshed within them.
- 4 In sum, we argue that the reports of the death of privacy-as-control are exaggerated. We believe that a more nuanced framework combining genuine individual agency, dominion, and control with vigorous structural reform offers a better path to the goal of privacy in the age of AI. Privacy law, at its best, can and must pursue both empowered individuals and restrained institutions, and especially private companies and Big Tech. The following sections

elaborate on this position, aiming to contribute constructively to the vital task of redesigning privacy law for the algorithmic age, where both agency and structure must be preserved if privacy is to endure as a meaningful, enforceable right.

B. Summary of Solove & Hartzog’s Argument

- 5 The central claim of Solove and Hartzog is that the traditional approach of privacy-as-individual-control has failed in today’s environment of ubiquitous data surveillance and AI-driven decision-making. They begin by describing two models for privacy protection.
- 6 The first is the ‘Individual Control Model,’ rooted in the work of Alan Westin and the Fair Information Practice Principles, which seeks to empower individuals with rights to decide how their personal information is collected, used, and shared.⁵ This model underpins most modern privacy laws, from data protection statutes in the 1970s to recent regimes like the European General Data Protection Regulation (GDPR), which they refer to as the ‘crown jewel of data privacy laws,’ and California’s CCPA, which give people notices, consent choices, and data subject rights to determine how enterprises use their personal data.⁶
- 7 The second approach is the ‘Societal Structure Model,’ which focuses on constraining organizations’ behavior through external rules and oversight, rather than relying on individual decisions and private enforcement.⁷ Although many scholars have advocated structural safeguards, policymakers have largely embraced the Individual Control paradigm in crafting data privacy regulations, which have been adopted during the past several decades.⁸
- 8 According to Solove and Hartzog, this emphasis on individual choice and consent is no match for the realities of the digital age, particularly the advent of AI. They invoke Kafka’s writing as a vivid metaphor for the powerlessness of individuals confronted by giant, inscrutable bureaucracies. In Kafka’s *The Trial*, for example, an ordinary man is prosecuted by a secret court without ever learning the charges he faces. This is a stark representation of how people today face sprawling data systems that make important decisions without transparency or accountability.⁹

5 Solove and Hartzog (n 1) 1025.

6 Solove and Hartzog (n 1) 1023 (‘Privacy laws are popping up at the state level in the United States like popcorn kernels in a sizzling frying pan.’) and 1026.

7 Solove and Hartzog (n 1) 1027.

8 Solove and Hartzog (n 1) 1029–31.

9 Solove and Hartzog (n 1) 1031.

2 Solove and Hartzog (n 1) 1023.

3 Solove and Hartzog (n 1) 1025–26.

4 Solove and Hartzog (n 1) 1024.

Modern privacy laws nominally give people rights (e.g., the right to access, correct, or use the data, rights to request not to be tracked or even not to sell one's data to third parties, and a requirement of consent in the collection and processing of data).¹⁰ Yet, in practice, these rights often provide little more than an 'illusion' of control.¹¹

- 9 Faced with opaque privacy policies and an endless stream of 'I Agree' prompts, individuals are becoming increasingly overwhelmed and cynical about the effectiveness of data protection laws. Most consumers lack the knowledge or endurance to truly understand or influence how their data is used, especially as AI systems silently analyze and repurpose that data in unexpected and unforeseeable ways.¹² Solove and Hartzog describe this state of affairs as *Kafkaesque*: people have formal privacy rights yet remain effectively powerless, much like the characters in a Kafka novel who cannot navigate or resist the faceless bureaucracy around them.¹³
- 10 The article by Solove and Hartzog also points to behavioral evidence that even under favorable conditions, individual self-determination falters. Even if manipulative 'dark patterns' are removed, people often still act against their own privacy interests.¹⁴ Despite warnings, individuals tend to 'embrace the technologies that hurt them' and, when 'given control over their data,' will frequently 'give it right back,' acting against their own interests.¹⁵ When personal data is misused in troubling ways, people even tend to blame themselves rather than the opaque systems that exploit their information.¹⁶
- 11 This dynamic is foreshadowed by Kafka's bleak view of human nature and opaque systems. In the age of AI, it means that placing the burden of privacy protection on each individual is essentially setting people and the whole data ecosystem up to fail. Companies and government agencies have little incentive to change intrusive practices when they can point to users' supposed 'consent' as justification for their bad behavior. In Solove and Hartzog's view, simply 'trying to give people control' does not truly empower them and can make matters worse.¹⁷
- 12 Given this critique, Solove and Hartzog contend that privacy law must adopt a different paradigm. Instead of pursuing what they see as a futile quest to empower each person in isolation, the law should

impose robust constraints on data practices at a society-wide level. The authors conclude that the primary goal of privacy regulation should no longer be maximizing individual control, but rather directly controlling the institutions and technologies that have rendered individuals so powerless.¹⁸

- 13 In short, they argue that individual control in the age of AI is either absent or illusory, and that only structural reform (rules directly limiting surveillance practices and ensuring fairness and accountability in data use) can meaningfully protect people's privacy moving forward.¹⁹
- 14 While this diagnosis is compelling in many respects, we propose three counterarguments to challenge and expand upon their conclusions.

C. The Unrealized Potential of an Individual Control Model (Response #1)

- 15 Solove and Hartzog build a compelling case against what is, in effect, a straw man: the feeble version of individual control that exists in practice today. Privacy self-management as currently implemented often means consumers clicking 'OK' on endless pop-ups, reading opaque policies (or more likely not reading them), and navigating cumbersome opt-out processes.
- 16 It is no wonder this regime has failed, and Solove and Hartzog are surely right to be dismissive. However, the flaw lies not in the concept of individual control itself, but in how superficially it has been deployed by legal and corporate actors in existing articulations, including the 'crown jewel,' GDPR. In critiquing the status quo, Solove and Hartzog aim at an easy target: a model of 'individual control' that has been watered down and co-opted by the data industry.
- 17 A more pertinent question is what a truly robust model of individual control would look like. What if an individual control model was designed to genuinely empower individuals in the data market,

10 California Consumer Privacy Act of 2018 (CCPA), as amended by the California Privacy Rights Act of 2020 (CPRA).

11 Solove and Hartzog (n 1) 1036.

12 Solove and Hartzog (n 1) 1035.

13 Solove and Hartzog (n 1) 1035–36.

14 Solove and Hartzog (n 1) 1035.

15 Solove and Hartzog (n 1) 1021, 1035.

16 Solove and Hartzog (n 1) 1021, 1035.

17 Solove and Hartzog (n 1) 1024.

18 Solove and Hartzog (n 1) 1040ff.

19 Solove and Hartzog (n 1) 1041 ('The law won't succeed in giving individuals control; instead, the law must try to control the larger forces that exploit people and to protect individuals, communities, and society-at-large from harm.').

not to placate users with hollow checkboxes and little substantive control?

I. From Enterprise-Centric to Genuine, 'Human-Centric' Empowerment

- 18 Solove and Hartzog's examples highlight people's failure to exercise privacy rights effectively. This is largely a predictable outcome of the current enterprise-centric notice-and-consent paradigm. Under the GDPR and CCPA, for example, individuals *nominally* have rights to access or delete data and to opt out of certain processing. In reality, these rights are often buried behind complex procedures that are created to retain the siloed data frameworks of the few dominant data giants.
- 19 Companies frequently make it difficult to opt out of tracking or data sales unless a user navigates a maze of settings or submits formal requests. The result is that only a tiny minority of users ever invokes these rights. Most individual consumers remain on the default setting, which can be described as 'pervasive data collection by default.' This is not because people have consciously chosen to surrender control. It is because control, as furnished today, is an illusion, or at best a privilege for the technologically savvy, the legally sophisticated, and those pig-headed or stubborn enough to devote their lives to protecting their privacy.
- 20 If individual empowerment is to be a meaningful principle of digital life, then modest reforms to consent architecture or regulatory frameworks will not suffice. Solove and Hartzog are surely right about that.
- 21 What is required, instead, is a more radical reorientation: a shift from an enterprise-centric data infrastructure to a human-centric one, in which personal data is *private by default*.²⁰ Under the current regime, corporations possess de facto control over individual data, and users are left to navigate a byzantine world of 'opt-outs,' disclaimers, and fine print to retain control.²¹ The predictable result is widespread passivity, fatigue, and a routinized surrender of privacy.
- 22 By contrast, a private-by-default framework would reverse the presumption. Rather than requiring individuals to assert their rights repeatedly, it would require service providers to request access and to justify that request in ways that are clear, specific, and grounded in actual user interest. Such a model does not reject consent; it honors it by restoring its substantive meaning, namely, personal control over some essential aspect of a person's identity.
- 23 More broadly, this shift would bring privacy law into closer alignment with moral intuitions and widely shared social norms and attitudes. People do not think of their personal data as a free-floating abstraction; they think of it as 'theirs,' something akin to their voice or their home or other piece of meaningful property that is intimately connected to their unique personal identity.²²
- 24 A private-by-default architecture gives legal form to this deeply held intuition about the meaning and value of data. It promotes trust by ensuring that data is technically secure and subject to user governance and control. It encourages companies to innovate along ethical dimensions (technological design, transparency, and accountability) rather than relying on data extraction by default. Most importantly, it invites a reframing of the individual's role in the digital ecosystem: from data subject to data sovereign. What is at stake beyond the issue of technical design and technological infrastructures is the broader question of what it means to live with dignity in the age of AI.
- 25 Building on a critique of the Solove and Hartzog approach, consider an analogy to consumer protection in other domains. If a product is sold with fine-print warnings that consumers rarely read, we do not conclude that safety does not matter to people; we conclude that the disclosure-based approach failed to actually protect them. Likewise, the failure of today's consent forms and privacy dashboards to safeguard privacy does not prove that individuals cannot or should not be empowered. It suggests that the method of empowerment was poorly designed and executed.
- 26 Research supports the view that individuals' apparent complacency is a byproduct of flawed systems, and not a lack of concern. For example, one recent survey found that while the median consumer would only pay about \$5 per month to preserve their privacy, they would demand roughly \$80 to allow

20 Paulius Jurcys, Mark Fenwick, Michael Fertik and Souichirou Kozuka, "'Private-by-Default': A Principle and Framework for Designing a New World of Personal AI" <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4839183> accessed 7 March 2026.

21 Mark Fenwick, Michael Fertik, Paulius Jurcys and Timo Minssen, 'Data Portability Revisited: Toward the Human-Centric, AI-Driven Data Ecosystems of Tomorrow' (2025) 27 Vand J Ent & Tech L 373.

22 Mark Fenwick, Paulius Jurcys, Aidas Liaudanskas, Brian D Earp, Sebastian Porsdam Mann and Souichirou Kozuka, 'Voice Cloning in an Age of Generative AI: Mapping the Limits of the Law and Principles for a New Social Contract with Technology' <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4850866> accessed 7 March 2026.

access to their personal data: a sixteen-fold gap.²³ This striking ‘super-endowment effect’ suggests that individuals attach far greater value to keeping their data private than their day-to-day behavior might indicate. Sunstein and Winegar concluded that standard metrics of privacy value are distorted by information deficits and cognitive biases and warned policymakers against dismissing people’s privacy interests based on those distorted measures.²⁴ Put simply, just because users click ‘Accept’ by default does not mean they do not care about privacy; it may mean the control mechanisms are too burdensome or skewed to reflect their true preferences.

- 27 Empirical evidence also suggests that when meaningful, easy-to-use privacy controls are available, individuals actively engage with them. For instance, when Apple’s iOS began explicitly asking users whether to allow app tracking, the vast majority (on the order of 75% of users) declined.²⁵ Until now, almost all users of mobile applications and IoT devices have been tracked by default without realizing it.
- 28 This stark change shows that when privacy controls are straightforward and user-friendly, people will often make use of them. Under today’s typical ‘opt-out’ regimes, the deck is stacked against users: the user’s consent is assumed until affirmatively revoked, and revocation is made laborious. However, *opt-in* and private-by-default data frameworks can flip that script. If data were private by default (more specifically, if there were a data infrastructure where an individual’s data is private-by-default, and individuals were truly in control of who can access their data, if individuals were asked for an *ex ante* permission), many would likely guard their data more closely. Solove and Hartzog’s argument underestimates how different the outcomes could be

if individual control were implemented in a human-centric rather than enterprise-centric fashion.

II. Re-Designing Individual Control for Real Impact

- 29 A more radical vision of individual control would address the problems Solove and Hartzog identify, rather than accepting them as inevitable and moving on to other possible solutions. For instance, one flaw they note is that people blame themselves when their data is misused. This is a form of victim-blaming encouraged by a system that says, ‘Well, you consented. What did you expect?’ Yet, a robust private-by-default individual control model could invert that dynamic by strengthening individuals’ negotiating power and shifting responsibilities onto data collectors. As an analogy, consider how consumer finance laws require clear, simple disclosures and sometimes mandate default rules (like automatic fraud protection on credit cards) instead of expecting consumers to catch every risk.²⁶ Privacy law could similarly require that individuals be given effective tools *and* strong default protections, so that the onus is not entirely on each person to police every data transaction.
- 30 To realize meaningful individual empowerment, we must transition from today’s enterprise-centric architecture (where personal data is collected, stored, and processed by centralized platforms) toward a human-centric, private-by-default infrastructure.²⁷
- 31 We need this shift on a technological level. In this new framework, personal data would reside in secure personal data instances or ‘vaults’ (e.g., on a user’s device backed up in personal data clouds), available exclusively to the individuals themselves (unless they grant permission to other service providers to access the data).²⁸ Rather than sending user data off to corporate service providers, data should stay on the individual user side, and AI-powered applications would then operate on the user’s side. In practical terms, this means that companies seeking to develop applications and services would have to request access, rather than assume it.
- 32 Under this human-centric architecture, data portability looks different, for example. Instead of transferring bits and pieces of information from one corporation to another, and having digital profiles and accounts scattered across services and devices, the individual user’s data remains with each user; third-party app developers can only access it if the

23 Angela G Winegar and Cass R Sunstein, ‘How Much Is Data Privacy Worth? A Preliminary Investigation’ (2019) 42 J Consum Policy 425, 426.

24 Winegar and Sunstein (n 23) 426.

25 See, for example, John Koetsier, ‘ATT Opt-In Rates 2024: Down, Down, Down (But Here’s How to Improve)’ (Singular, 21 August 2024) <<https://www.singular.net/blog/att-opt-in-rates-2024/>> accessed 7 March 2026; ‘App Tracking Transparency: Four Years of Impact on User Privacy’ (AppleMagazine) <<https://applemagazine.com/app-tracking-transparency-cuts-u-s-tracking-rate/>> accessed 7 March 2026.

26 12 CFR pt 1026 (Regulation Z).

27 On this transition, see Jurcys and others (n 20).

28 Jurcys and others (n 20).

individual gives prior permission.²⁹ This framework strengthens privacy and security and also lowers the costs and risks of moving between different services. Gone are the so-called ‘walled gardens’ of the current data ecosystem. Instead, individual users hold a comprehensive dataset, and outside services must come to them following the terms set by individuals.

- 33 In practical and legal terms, such a reimagined version of individual data control means engaging with and developing a new notion of data ownership or dominion over personal data. Existing rules often emphasize consent; yet, they do not treat data as something that truly belongs to individuals. On the contrary, companies have been treated as the de facto owners of data collected from the users of their products and services.
- 34 A new human-centric data infrastructure would form a technical foundation for legal recognition of personal data as a new type of property.³⁰ This would open up markets in which people could license their information on conditions they find acceptable, or choose not to disclose at all, as a default preference. With property rights firmly in place and privacy as the baseline, users would no longer be forced to navigate a fragile system of notice and consent in which all the power lies with the owner of the service acting as a gatekeeper. Rather, individuals would have substantive power and entitlements. In the end, this shift would offer a renewed social contract in the age of AI that respects autonomy, promotes trust, and ensures that individuals benefit from the opportunities afforded by AI.
- 35 To be sure, none of these changes is easy; all would require moving beyond the familiar notice-and-consent template. A revitalized Individual Control Model would not leave people to fend for themselves; it would use regulation and technology to amplify individual agency and correct the power imbalances of the digital marketplace.
- 36 Solove and Hartzog are correct that the current diluted form of individual control has not protected people from Kafkaesque abuse. Yet, their conclusion that personal empowerment is futile assumes that the flawed status quo is the best we can do. A better response is that we have never genuinely tried to empower individuals in the robust ways that truly matter.
- 37 Instead of abandoning the ideal of personal control, the solution should aim to radicalize it: to make individual privacy rights truly effective and human-centric. By spotlighting the failures of the weakest version of privacy self-management, Solove and Hartzog in fact highlight why we must strengthen, not abandon, individual agency over their data in the age of AI. Where individuals are given real alternatives, easy exits, and a stake in their data, ‘privacy as control’ could be far more potent than Kafka’s pessimistic vision imagines. In the end, protecting privacy will require both meaningful individual control *and* strong structural safeguards. The shortcomings of self-management should spur us to reinvent individual control, not renounce it.
- 38 The concept of data ownership is often misunderstood as an all-or-nothing proposition. In reality, ownership (whether of land, intellectual property, or digital information) is better understood as a ‘bundle of rights,’ comprising various entitlements such as access, control, exclusion, and transfer.³¹ Recognizing this layered structure of ownership opens up the possibility of a ‘staggered’ approach to data governance. More sensitive data (e.g., health information, biometric identifiers, or payment history) could be afforded a stronger and more complete bundle of rights, including restrictions on reuse, sale, or algorithmic processing. Conversely, less sensitive or more transformed data (e.g., anonymized datasets or information significantly modified by a company’s proprietary technologies) might be governed under a looser set of rights.
- 39 This flexible, graduated model avoids the rigidity of absolute control while preserving a human-centric vision of privacy. It enables the design of systems that uphold the dignity and autonomy of individuals without chilling innovation, especially when firms are rewarded for meaningful investment in data stewardship or novel data transformation. Structuring dominion and ownership over data in this layered way can better match legal entitlements with the practical realities and complex configurations of interests in data processing in a digital economy. This should be supported by effective, easy-to-use remedies that make it possible to sanction data breaches and misuse.

²⁹ Jurcys and others (n 20).

³⁰ Paulius Jurcys, Christopher Donewald, Mark Fenwick, Markus Lampinen, Vytautas Nekrošius and Andrius Smaliukas, ‘Ownership of User-Held Data: Why Property Law Is the Right Approach’ (2021) *Harvard Journal of Law & Technology Digest* <<https://jolt.law.harvard.edu/digest/ownership-of-user-held-data-why-property-law-is-the-right-approach>> accessed 7 March 2026.

³¹ Jeremy Waldron, ‘What Is Private Property?’ (1985) 5 *OJLS* 313; Jeremy Waldron, ‘The Rule of Law and the Measure of Property. Lecture 2: Unravelling the Form and Substance of Property’ <https://www.law.nyu.edu/sites/default/files/ECM_PRO_068939.pdf> accessed 7 March 2026.

D. Individual and Structural Solutions Are Complementary, Not Competing (Response #2)

- 40 A second issue with Solove and Hartzog is that they frame privacy reform as a choice between empowering individuals or imposing structural safeguards. This appears like a false binary. In reality, effective privacy protection will require a combination of individual rights and institutional constraints. Far from being at odds, personal agency and structural regulation reinforce each other, each addressing weaknesses the other cannot.
- 41 Modern privacy regimes already blend individual and structural approaches. The GDPR, for example, gives data subjects a suite of personal rights (access, deletion, objection, portability) and simultaneously obligates data controllers to abide by strict principles (data minimization, purpose limitation, security, etc.) regardless of user intervention. These two pillars are meant to work in tandem. The European Data Protection Board has emphasized that controllers must implement appropriate technical and organizational measures to ensure that data protection principles and data subjects' rights are effective 'by design and by default.'³²
- 42 In other words, structural measures such as privacy-by-design requirements are put in place precisely to facilitate and bolster individual control. A law might mandate, for instance, that companies only collect the minimum data necessary (a structural rule), which in turn reduces the burden on individuals to constantly police what is collected about them. Likewise, giving individuals the right to access their data (individual empowerment) complements that structural rule by allowing them to verify and challenge whether an organization truly respected the collection limits.
- 43 Consider how privacy laws operate in practice: consumer data regulations often require businesses to offer *user-friendly mechanisms* for exercising rights, which essentially is a structural mandate aimed at enabling individual agency. Even Solove and Hartzog's preferred 'Societal Structure Model' would inevitably rely on individual participation at some stages (e.g., people still need channels to report abuses or seek redress).
- 44 Many structural measures are explicitly designed to amplify individual choice, not supplant it. As noted
- in Response #1, data portability and interoperability requirements are regulatory impositions on companies; their purpose is to liberate individuals from lock-in and give them greater freedom to choose privacy-respecting services.
- 45 Similarly, restrictions on certain data practices (e.g., bans on using sensitive data for ad targeting without explicit consent) create a clearer space for meaningful consent. Such restrictions ensure that if consent is sought, it is for something genuinely optional and not for abusive baseline activities. In the context of AI, structural rules such as algorithmic transparency and impact assessments can empower individuals by informing them how decisions about them are made, thereby enabling them to challenge or avoid harmful uses. Rather than seeing structural protections as negating the role of the individual, we should view them as establishing a safer environment in which individuals can then exercise their agency.
- 46 Conversely, empowering individuals can also advance structural goals. If the environment is changed through regulation (say, by making privacy the default setting or outlawing the most egregious data practices), individuals' choices will carry more weight. Cass Sunstein's concept of 'libertarian paternalism' advises regulators to establish strong default rules and preserve individual freedom to opt out.³³
- 47 In a privacy context, this means implementing privacy-protective defaults (a structural measure) while still allowing people to voluntarily share data if they truly wish (an exercise of personal choice). Such a *hybrid approach* harnesses structure to guide outcomes in the socially desired direction yet preserves individual autonomy. In practice, it often yields better results than relying on either strategy alone: most people enjoy robust privacy by default, and those who genuinely wish to exchange some privacy for benefits remain free to do so.
- 48 Hence, privacy discourse should move past the false choice of 'individual vs. structural' and recognize that a staggered, layered approach is more effective. Even Solove and Hartzog's own proposals imply a blend: they call for laws to 'bring the system under control,' which would likely include prohibitions and duties for data holders. Yet, any such regime would also need avenues for individual remedies (such as the ability to know when one's rights have been violated and to seek enforcement). Rather than pitting the Individual Control Model against the Societal Model, the better view is to fuse them. Strong structural safeguards (backed by regulators and courts) can ensure that individuals are not overwhelmed or deceived when making privacy

32 EDPB, 'Guidelines 4/2019 on Article 25 Data Protection by Design and by Default' (Version 2.0, 20 October 2020) <https://www.edpb.europa.eu/sites/default/files/files/file1/edpb_guidelines_201904_dataprotection_by_design_and_by_default_v2.0_en.pdf> accessed 7 March 2026.

33 Richard H Thaler and Cass R Sunstein, *Nudge: The Final Edition* (Penguin 2021) 324.

choices. At the same time, preserving avenues for individual agency can make structural rules more adaptive and democratically legitimate (by allowing personal preferences and consent to play a role where appropriate).

- 49 The rise of AI amplifies the stakes dramatically. Unlike traditional data systems, modern AI systems are not just opaque; they are often described as ‘black boxes’ because their functioning and the use of data are incomprehensible even to their creators. Machine learning models generate outputs through processes that cannot be fully audited or predicted in real time, leading to outcomes that are neither explainable nor contestable. These systems are complex and non-deterministic. They generate decisions, outputs, and inferences that are often indecipherable. The logic behind a given result may be emergent, shaped by opaque training data, or probabilistic inference across vast datasets.
- 50 This extreme epistemic opacity renders existing privacy safeguards dangerously inadequate. When the companies deploying these systems cannot themselves anticipate the full range of behaviors their AIs might exhibit, the result goes beyond just disempowerment and leads to a radical systemic volatility. People are subject to decisions about credit, hiring, policing, and even their own health, driven by algorithms whose logic is unavailable and whose errors are untraceable. This goes beyond the Kafkaesque paradigm: a world where individuals are governed by machinic logics they cannot even see or name, let alone challenge. Under such conditions, the need for both robust structural constraints *and* meaningful individual rights becomes even more pressing. Legal and technical systems must place intelligibility and human override at their core and as baseline conditions for deployment.
- 51 In such a context, structural regulation alone becomes insufficient: how do we constrain what we cannot describe, explain, or predict? What safeguards could be placed on AI systems and autonomous AI agents that do not abide by concepts of time, space, or gravity as we humans do? Kafka’s vision of impenetrable systems has become newly literal. Individuals are now targets, inputs, and sometimes even co-creators in systems that elude explanation. Protecting them demands active orientation, intelligibility, and control.
- 52 This is precisely why a revitalized, human-centric vision of privacy anchored in structural guarantees as well as grounded in and guaranteed by meaningful personal agency is essential. AI systems increasingly operate in direct interface with individuals: shaping choices, generating content, mediating relationships, and simulating trust. If we rely solely on institutional guardrails, we risk reproducing the very Kafkaesque condition we seek to overcome: a world where

individuals are trapped within inscrutable logics they cannot influence or resist. However, if individual empowerment is well-designed, it can serve as a complement to structural constraints. Furthermore, individual empowerment could be a counterweight to the unique disorientation that AI, in the hands of bad or careless actors, potentially induces. Rather than retreating from the idea of control, we should aim to reinvent it: real tools for real users, in systems built for human-scale understanding.

- 53 On this account, individual empowerment and structural regulation are not mutually exclusive, not an either/or. Rather, they are mutually reinforcing. The most promising path forward is to embrace both. As many observers, including Prof. D. Solove himself, note, protecting privacy in the age of AI requires both ethics in technology frameworks and enforceable legal rights.³⁴ A comprehensive framework can establish broad ‘no-go’ zones for data misuse and give individuals a meaningful voice in how their data is handled. Rather than abandoning one approach in favor of the other, regulators should focus on how structural and individual measures can be intelligently designed and combined to produce a privacy regime greater than the sum of its parts.

E. Kafka’s Lesson Is to Dismantle Oppressive Systems, Not Embrace Fatalism (Response #3)

- 54 There is a tension at the heart of Solove and Hartzog’s use of Kafka. Their reading is powerful as a critique of the Individual Control Model *as implemented*. However, their invocation of Kafka as evidence that individual rights are inherently futile stretches beyond what Kafka’s work supports. Kafka indeed portrays opaque, dehumanizing systems that crush individuals. These portrayals are more accurately understood as a *call for awareness, outrage, and structural reform*, not as a rejection of all individual agency or rights. The tragedy in *The Trial* or *Before the Law* is not that resistance is foolish; it is that resistance has been structurally undermined. This is precisely what makes these stories so haunting and morally urgent. They invite readers to recognize the cruelty of such systems and demand better (not to accept them as inevitable or unassailable).
- 55 Solove and Hartzog draw on Kafka to argue that attempts to empower individuals (particularly through legal rights and procedural safeguards) are misguided and even counterproductive.³⁵ They observe that individuals often surrender control

34 See, for example, Daniel J Solove, ‘Artificial Intelligence and Privacy’ (2025) 77 Fla L Rev 1.

35 Solove and Hartzog (n 1) 1031, 1035.

willingly, blame themselves for harms they suffer, and struggle to make rational privacy decisions even when given the tools to do so.³⁶ These are important insights, and the authors are right to challenge the overly optimistic assumptions embedded in many privacy frameworks. They are also persuasive in showing how the *illusion* of control, as exemplified by endless consent forms, unread privacy policies, and performative checkbox rights, has lulled policymakers into a false sense of accomplishment.³⁷

- 56 But where Solove and Hartzog go astray is in using Kafka to suggest that individual rights are doomed in principle, not just in practice. They frame Kafka as a prophet of resignation: a writer who reveals the futility of personal agency in the face of systemic power.³⁸ Yet Kafka's work, in its very bleakness, performs a moral function that is fundamentally oppositional. His narratives of disempowerment and absurdity are not endorsements of fatalism; they are aesthetic and ethical provocations. By forcing readers to witness the horror of faceless authority and procedural indifference, Kafka's stories generate a kind of moral clarity. If individuals are unable to resist in Kafka's world, it is not because resistance is meaningless, but because *every channel of resistance has been corrupted or foreclosed*. That is precisely why readers experience Josef K.'s fate as a tragedy, not a farce.
- 57 To cite Kafka as proof that individual rights should be abandoned is to misread the moral arc of his work. His fiction does not instruct us to give up on the individual. Surely, it demands that we construct a world in which individuals are no longer crushed by illegible and unaccountable systems. In fact, Kafka's legacy has more often inspired the *expansion* of individual rights than their retreat. Legal scholars, judges, and policymakers have repeatedly drawn on Kafka's imagery to support reforms that enhance due process, transparency, and accountability. One need only look to *Goldberg v Kelly* (1970),³⁹ in which the US Supreme Court ruled that welfare recipients must be granted an evidentiary hearing before losing benefits. While Kafka is not cited directly, the structure of the Court's reasoning reads like a direct rebuttal to the world of *The Trial* (in *Goldberg*, the court rejected summary, opaque decision-making in favor of procedural fairness). The decision affirms the importance of formal rights and procedures, including the individual right to an oral hearing before an impartial decision-maker, the right to

confront and cross-examine witnesses, and the right to a written statement setting out the evidence relied upon and the basis of the decision, that afford meaningful opportunities to understand, contest, and influence the forces that govern our lives.

- 58 Solove and Hartzog acknowledge these kinds of legal reforms, but they suggest that they are ultimately insufficient. Perhaps. But insufficient is not the same as futile. The failures of current privacy laws do not mean that any or all efforts to empower individuals are doomed. They mean that our current implementations are poorly designed, under-enforced, or too easily captured by the very institutions they are meant to regulate. That is a call for improvement. Indeed, many of the authors' own structural proposals would still depend, at some level, on individual agency: on the individuals' ability to report harms, to seek redress, to demand justification, to challenge abusive practices. Even the best structural constraints still require individuals to recognize when something has gone wrong. And for that, meaningful rights and access to information remain essential.
- 59 To be sure, Solove and Hartzog are correct in asserting that individual rights, standing alone, are insufficient to shoulder the full weight of modern privacy protection. They are right to emphasize the importance of structural safeguards such as data minimization rules, algorithmic accountability, duties of loyalty, and substantive restrictions on harmful uses irrespective of consent. However, the force of their argument diminishes when it slides toward a form of determinism that seems to suggest that individuals will invariably act against their own interests, and that initiatives aimed at empowering them are ineffective and potentially counterproductive. This framing not only risks a misreading of Kafka's critique; it may inadvertently entrench the very disempowerment it seeks to expose. If individuals are led to believe that their actions are meaningless, their autonomy illusory, and their resistance futile, they may well succumb to apathy. In that moment, the Kafkaesque allegory becomes a self-fulfilling prophecy.
- 60 A more generous reading of Kafka may offer a useful complement to Solove and Hartzog's critique, without capitulating to its more pessimistic conclusions. Rather than construing Kafka's stories as repudiations of individual agency, we might instead interpret them as stark indictments of legal and bureaucratic systems that promise rights while systematically undermining their realization. The core tragedy in Kafka's work relates to the institutional betrayal of pursuing agency. The tragedy becomes obvious when formal entitlements exist in name only, and the very structures designed to empower are configured to frustrate and

36 Solove and Hartzog (n 1) 1035.

37 Solove and Hartzog (n 1) 1030, 1036.

38 For a version of this interpretation of Kafka, see, for example, Gilles Deleuze and Félix Guattari, *Kafka: Toward a Minor Literature* (Dana Polan tr, University of Minnesota Press 1986).

39 *Goldberg v Kelly* 397 US 254 (1970).

disempower. The lesson, then, is to render individual empowerment real.

- 61 Viewed through this lens, the imperative is to reconsider the project of individual rights so that they meet the realities of the age of AI. Today is the time when we need to craft individual rights that move beyond formal aspirations towards functionally meaningful measures that individuals can rely on in practice. This requires designing mechanisms and frameworks of control that are accessible, enforceable, and embedded within institutional frameworks that alleviate, rather than exacerbate, the burdens on individuals. It also demands a recognition that individual and structural safeguards are interdependent pillars of a coherent privacy architecture.
- 62 Kafka's work, far from undermining the case for individual agency, amplifies the urgency of fortifying it and supporting its realization. His narratives expose the human cost of legal systems that profess to confer rights while denying the means to exercise them. The appropriate lesson is to search for a renewed and reimagined framework. Such a new framework must empower individuals with meaningful tools for resistance through law, technological infrastructure, and collective will. Only then can we begin to reimagine a digital order in which Kafka's dystopias are no longer prophetic, but cautionary reminders that help us to face the future with open eyes.

F. Concluding Remarks

- 63 Franz Kafka's specter continues to loom over contemporary debates on privacy. His work reminds us of what is at stake. The above discussion has challenged the notion that individual empowerment is a futile endeavor in the age of AI. On the contrary, when individual control is reimagined in a robust way and buttressed by structural safeguards, it can be a powerful force for transparency and accountability in the data ecosystem, while still being supportive of technological advances to the benefit of mankind.
- 64 The purported dichotomy between individual and structural approaches dissolves under critical scrutiny. Durable privacy frameworks are not built on choosing between the two, but on interweaving them: structure must enable agency, just as agency demands structure to be meaningful. While some authors would like to posit Kafka's enduring legacy as an indictment of rights-based frameworks, in this paper, we invite the readers to redeem such-rights-based approaches through systems designed to uphold human dignity and enable its defense rather than erode it. That requires a dual-front strategy: on one hand, acknowledging the limits of enterprise-centric, extractive data practices that default to

surveillance and opacity; on the other, advancing human-centric frameworks that equip individuals with meaningful tools to assert dominion, ownership, and control over their personal data, while imposing enforceable obligations and oversight mechanisms on data controllers and institutions.

- 65 Looking ahead, the challenge is one of operationalization. Scholars, technologists, and regulators must work collaboratively to translate this integrated vision into practical tools and enforceable norms. Research should advance technical architectures (such as private-by-default data models, personal data vaults, controlled personal AI agents, and calibrated nudging systems) that empower users with varying capacities and backgrounds by design. In parallel, legal reforms must move beyond procedural formalism to establish substantive duties of fairness, transparency, and redress. In the context of AI, this means more than prohibiting harm. It requires ensuring that individuals are informed, empowered, and meaningfully involved in decisions that shape their lives. It means rejecting deterministic narratives of technological inevitability in favor of governance models grounded in democratic values and normative commitments.
- 66 Ultimately, we invite privacy law experts and technologists to move beyond the feeling of resignation and despair toward a practical, realistic, and enforceable frameworks fit for the agentic age. Such frameworks could not only serve the interests of individuals but also restore public trust in technologies, platforms, and public institutions, strengthen civic infrastructure, and lay the foundation for sustainable and ethical innovation. In doing so, we honor the very aspirations that Kafka's fiction mourned. Perhaps we might even contribute to offering a path toward a more just and human-centered future.