

Comment & Analysis

Nicolas Alder*, Kai Ebert*, Ralf Herbrich, Philipp Hacker

AI, Climate, and Transparency: Operationalising and Improving the AI Act

This paper critically examines the AI Act's provisions on climate-related transparency, highlighting significant gaps and challenges in its implementation. We identify key shortcomings, including the exclusion of energy consumption during AI inference, the lack of coverage for indirect greenhouse gas emissions from AI applications, and the lack of standard reporting methodology. The paper proposes a novel interpretation to bring inference-related energy use back within the Act's scope and advocates for public access to climate-related disclosures to foster market accountability and public scrutiny. Cumulative server level energy reporting is recommended as the most suitable method. We also suggest broader policy changes, including sustainability risk assessments and renewable energy targets, to better address AI's environmental impact through transparency measures that facilitate competition among model providers and allow conscious consumer decisions. Finally, we provide first insights on the General-Purpose AI Model Code of Practice as it relates to climate reporting.

I. Introduction

The climate implications of artificial intelligence (AI), including energy and water consumption, are increasingly subjected to public scrutiny and academic research¹. In one respect, the application of AI in environmental issues provides numerous promising opportunities for more effectively tackling climate and sustainability challenges². These include advancements in carbon measurement for cloud computing³, utilising earth observation to aid electricity grids in achieving carbon neutrality⁴, enhanced decision-making for carbon capture technologies⁵, and accelerating scientific processes across various fields⁶.

In another respect, energy efficiency targets for data centres are under discussion,⁷ and there is concern that their energy consumption could surpass the available supply of renewable energy. Major companies like Google have reported that increased energy demand related to AI endangers their carbon zero strategies⁸.

As in many collective action problems, regulation may play a major part in mitigating the negative impact of AI on climate while fostering socially beneficial use cases. Additionally, recent studies suggest that thoughtful regulation even fosters necessary green innovation in the industry⁹. Globally, several initiatives are underway to establish legal frameworks for AI. The most prominent example is probably the EU AI Act¹⁰, which has just become applicable at the beginning of August 2024. The Act also includes significant sections concerning climate impacts, primarily reporting obligations. Thus, one might hope, the climate effects of AI could become a relevant market parameter; have reputational repercussions; and enable public scrutiny, by analysts and NGOs.

Against this background, this article offers a novel legal-technical analysis of the transparency provisions of the AI

Act, making three core contributions. First, we demonstrate that the Act falls short in several critical areas. Second, we show that even within its current scope, operationalising its mandates poses significant challenges. Third, we propose targeted reforms for legislators: (i) reinterpret Art. 11 and 53 to encompass inference energy; (ii) operationalize energy disclosure via cumulative server-level measurement and provide PUE numbers; (iii) extend reporting to indirect greenhouse gas emissions and data-center water use; and (iv) mandate public release of all climate-related data. These recommendations would provide the Act with the rigour and accountability essential for sustainable AI governance and inform broader policy initiatives, including the United Nations' Global Digital Compact currently under discussion at the United Nations.

* Equal Contribution by Nicolas Alder, Hasso Plattner Institute in Potsdam, Germany, email: nicolas.alder@hpi.de; Kai Ebert, European University Viadrina in Frankfurt (Oder), Germany, email: kebert@europauni.de.

- 1 Emma Strubell/Ananya Ganesh/Andrew McCallum, 'Energy and Policy Considerations for Modern Deep Learning Research' (2020) 34 EAAI-20 <https://doi.org/10.1609/aaai.v34i09.7123> (accessed 12.7.2025); Jesse Dodge et al., 'Measuring the Carbon Intensity of AI in Cloud Instances' (2022) FAccT'22 <https://dl.acm.org/doi/10.1145/3531146.3533234> (accessed 12.7.2025); Lynn H. Kaack et al., 'Aligning Artificial Intelligence with Climate Change Mitigation' (2022) 12 Nature Climate Change 518; Alexandra Sasha Luccioni/Sylvain Viguier/Anne-Laure Ligozat, 'Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model' (2023) 24 JMLR 1; Pengfei Li et al., 'Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models' (2023) arXiv:2304.03271 <https://arxiv.org/abs/2304.03271> <https://doi.org/10.48550/arXiv.2304.03271> (accessed 12.7.2025); Philipp Hacker, 'Sustainable AI regulation' (2024) 61 CMLR 345.
- 2 David Rolnick et al., 'Tackling Climate Change with Machine Learning' (2022) 55 CSUR 1; Carole-Jean Wu et al., 'Sustainable AI: Environmental Implications, Challenges and Opportunities' (2022) 4 Proceedings of Machine Learning and Systems 795; Josh Cowsils et al., 'The AI Gambit: Leveraging Artificial Intelligence to Combat Climate Change – Opportunities, Challenges, and Recommendations' (2023) 38 AI & Society 283.
- 3 Dodge et al. (n 1).
- 4 Claudio Persello et al., 'Deep Learning and Earth Observation to Support the Sustainable Development Goals: Current Approaches, Open Challenges, and Future Opportunities' (2022) 10 IEEE Geoscience and Remote Sensing Magazine <https://doi.org/10.1109/MGRS.2021.3136100> (accessed 12.7.2025).
- 5 Anirban Chandra et al., 'Fourier Neural Operator Based Surrogates for CO₂ Storage in Realistic Geologies' (2025) arXiv:2503.11031 <https://doi.org/10.48550/arXiv.2503.11031> (accessed 12.7.2025).
- 6 Hanchen Wang et al., 'Scientific Discovery in the Age of Artificial Intelligence' (2023) 620 Nature 47.
- 7 Further EU legislation can be expected following the Commission's report by 15.5.2025 under Art. 12(5) Energy Efficiency Directive (EU) 2023/1791.
- 8 Google, 'Google 2024: Environmental Report', 2024, <https://www.gstatic.com/gumdrop/sustainability/google-2024-environmental-report.pdf> (accessed 6.6.2025).
- 9 Wanli Zhang et al., 'Revisiting the Porter Hypothesis: A Multi-Country Meta-Analysis of the Relationship between Environmental Regulation and Green Innovation' (2024) 11 Humanit Soc Sci Commun 232.
- 10 Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 [2024] OJ L 2024/1689.

II. Climate Transparency and the AI Act: Gaps and Interpretation Challenges

Any change for the better starts with information about what is wrong. However, at the moment, it is often unclear what the exact impact of the development and usage of an AI model is concerning energy and water consumption. The AI Act seeks to provide a remedy by forcing certain AI providers to make climate-related disclosures. However, the patchwork of provisions includes seven significant ambiguities and loopholes.

First, for high-risk AI systems, providers are required under Art. 11(1) to document the computational resources used in development, training, testing, and validation, as per Annex IV(2). However, there is no explicit requirement to disclose energy consumption, limiting the comparability of, and transparency on, the environmental impact of these high-risk systems to estimates based on the documented computational resources.

Second, the AI Act imposes transparency obligations on providers of general-purpose AI (GPAI) models, particularly concerning energy consumption. Under Art. 53(1)(a), providers must maintain up-to-date technical documentation that includes information specified in Annex XI, which requires known or estimated energy consumption of the model, with estimates potentially based on computational resources. However, this requirement focuses on the model's development phase, excluding the inference phase, as it forms part of the requirements in Annex XI Section 1 para. 2, which refers to "information of the process for the development [of the model]." If indeed the energy consumption is restricted to training and excludes inference (see below, V.), this constitutes a significant oversight given the potentially much greater cumulative energy consumption during inference¹¹. To address this gap, a novel interpretation can be considered. Art. 53(1)(a) and (b), in conjunction with Annex XI and Annex XII, require providers to include in the documentation for downstream AI system providers and authorities information on the technical means needed to integrate the GPAI model into AI systems. Although energy consumption is not explicitly mentioned, these provisions should, arguably, be interpreted to include information on hardware requirements, allowing downstream providers to estimate the energy consumption for inference. This novel interpretation would indirectly ensure transparency regarding inference energy use.

A third issue arises with open-source (OS) GPAI models, which are generally exempt from transparency obligations unless they pose a systemic risk (Art. 53(2)). Recital 102 emphasises transparency for OS models but does not include energy consumption in the information that must be disclosed. Rather, the focus is on parameters, model architecture, and usage information, leaving a gap in transparency regarding the energy impact of these models.

Fourth, regarding fine-tuning, Recital 97 seems to imply that an entity engaging in any, even minuscule, fine-tuning of a GPAI model automatically becomes the provider of a new model, with all corresponding duties. For minor changes, this seems excessive, even though Recital 109 suggests that reporting obligations are limited to that fine-tuning. However, Art. 25(1)(b) holds that, for high-risk AI systems (e.g., in recruitment), only a substantial modification bestows provider status upon the modifying entity. This rule could be analogised for fine-tuning, such that only substantial model

modifications via fine-tuning lead to provider status, protecting smaller entities. Commission Guidelines are expected over the summer on this issue, and they will likely contain a Floating Point Operations (FLOP) threshold: if the fine-tuning uses less than 1/3 of the FLOPs of the original model, a presumption is triggered that the modification is not substantial.¹² However, as scholars have suggested, providers still need to conduct a consequence scanning exercise to determine if their modification significantly changes the risk profile of the model, leading to a "compute and consequence" test for new provider status.¹³

Fifth, the AI Act overlooks the greenhouse gas (GHG) effects of AI applications, such as those used in oil and gas exploration¹⁴. This omission leaves a significant gap, as these applications can substantially contribute to climate change, yet their environmental impact remains unreported.

Sixth, while the Act requires energy consumption to be documented, this information is only available to authorities, not downstream providers (unless our suggested interpretation is adopted), and certainly not to the general public. Without broader access to this data, transparency and accountability are significantly curtailed, hindering market effects based on climate reporting, independent research and verification, and public scrutiny by analysts and NGOs.

Finally, the Act also fails to address the use of toxic materials and water consumption, a critical factor in data centre operations. While most data centres in the EU must report their water usage under the Energy Efficiency Directive¹⁵, the AI Act lacks a specific attribution to AI, as stipulated for energy consumption, and computing outside the EU is not covered. Given the significant water usage for cooling in data centres, this omission leaves a major aspect of AI's environmental impact unreported.

III. Operationalising the Requirements: Implementation Challenges

As the previous section showed, under the current version of the AI Act, GPAI providers must log the energy consumption used for training GPAI models. To operationalise this provision, it is crucial to clarify how energy consumption should be measured or estimated. We discuss three methods: measurement at the data centre level; at the cumulative server level; and at the individual graphic-processing unit (GPU) level.

Energy efficiency in data centres is measured by the Power Usage Effectiveness (PUE) metric. It denotes the ratio of total energy used by the data centre to the energy consumed by its computational hardware. A lower PUE indicates higher energy efficiency, with a global average PUE of 1.58 recorded in 2023¹⁶. When measuring energy consumption at the data

11 Sasha Luccioni/Yacine Jernite/Emma Strubell, 'Power Hungry Processing: Watts Driving the Cost of AI Deployment?' (2024) Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency 85; Wu et al. (n 2) 795.

12 See European Commission, 'Targeted Consultation in Preparation of the Commission Guidelines to Clarify the Scope of the Obligations of Providers of General-Purpose AI Models in the AI Act', 2025, 9, <https://digital-strategy.ec.europa.eu/en/news/commission-seeks-input-clarify-rules-general-purpose-ai-models> (accessed 15.7.2025).

13 Philipp Hacker/Matthias Holweg, 'The Regulation of Fine-Tuning: Federated Compliance for Modified General-Purpose AI Models' Working Paper (2025), <https://ssrn.com/abstract=5289125> (accessed 15.7.2025).

14 Kaack et al. (n 1) 518.

15 Directive (EU) 2023/1791 on energy efficiency and amending Regulation (EU) 2023/955 [2023] OJ L 231.

16 Google (n 8).

centre level, the advantage lies in capturing the total power usage, including both direct computing energy and overhead like cooling. This provides a comprehensive overview and encourages efficient data center selection. However, it can obscure the energy impacts of specific model architecture or software inefficiencies, as these are influenced by the data center's overall efficiency. Estimating with the PUE ratio is practical but may lack precision for specific model-level insights.

At the cumulative server level, i.e., for all utilized servers within one data center, energy measurement with power distribution units is highly accurate, closely reflecting model size, data volume, and software efficiency. This method is recognized in the industry and can provide detailed insights into energy consumption. However, not all data centers currently track power demand at this level, and implementing such systems can be time-consuming¹⁷. While cloud providers like AWS and Azure may have these capabilities, widespread reporting standards are lacking, potentially disadvantaging smaller companies.

Finally, measuring energy usage at the GPU level within a server is straightforward with on-chip sensors for components like NVIDIA GPUs, which offer user-friendly monitoring. However, this approach significantly underestimates total energy consumption as it only accounts for a single component, missing the broader picture of server-wide energy use. Therefore, it is not recommended for comprehensive energy tracking.

IV. Discussion and Policy Proposals

The AI Act is a first step toward mandatory AI-related climate reporting, but is riddled with loopholes and vague formulations. To remedy this, we make six key policy proposals. Such mechanisms should not only be included in the evaluation report due in August 2028 (Art. 111(6)), but in any interpretive guidelines by the AI Office and other agencies, reviews and potential textual revisions beforehand.

The primary weakness of the AI Act is the exclusion of inferences from explicit and mandatory energy consumption reporting. While we offer a solution for interpretation, it is unclear whether courts, agencies and companies will follow this route. This significantly hampers the assessment of future AI energy usage, related carbon emissions, and effects on (renewable) energy infrastructure. Hence, future guidance from the AI Office, and delegated acts by the Commission (Art. 53(5) and (6)), should explicitly include inference as a reporting category, both in Annex XI (for the AI office) and XII (for downstream actors).

Another major challenge is the failure to include indirect emissions by AI applications (e.g., for oil and gas exploration) and water consumption within the reporting obligations. This should be remedied at the provider (water) and the deployer level (applications).

Third, the consequences of minor fine-tuning operations on GPAI remain unclear. It would be beneficial to tie the energy reporting requirement to the mechanism of training (fitting model weights) and incorporate a minimum computational cost threshold, as this would encompass energy-intensive training and fine-tuning for reasonably sized workloads.

The open source exemption, fourth, should be revoked. There is no convincing reason to abstain from climate reporting only because other parts of the model are made public and transparent.

Fifth, energy consumption measurements ought to be conducted at the cumulative server level and reported accordingly. This reflects the total computation-related power usage. Furthermore, the PUE factor of each data centre, as measured and reported under the Energy Efficiency Directive (EU) 2023/1791 and Delegated Regulation (EU) 2024/1364¹⁸, provides information for a relevant estimate of the overall energy consumption. By reporting these numbers separately, we can differentiate between the power usage specific to the model (server-level computation) and the efficiency of the data centre, thus reflecting the realistic overall energy investment. Estimations for server-level power consumption should utilise peak utilisation values from the hardware manufacturer (e.g., NVIDIA). When actual measurements are available, they must be prioritised over estimations. The ultimate aim is to secure as precise power consumption data as possible, allowing for flexibility for model providers with limited access to data infrastructure (such as for finetuning with substantial modification), while also ensuring that estimations are not misused to avoid accurate measurement reporting. These considerations should inform both the technical standards drafted under Art. 40 AI Act and the possible implementation of the Global Digital Compact at the international level.

Finally, sixth, all climate-related disclosures must urgently be made available to the general public, not only to authorities and, potentially, downstream actors. Trade secrets and intellectual property do not stand in the way if only aggregate numbers at the cumulative server level are reported. Only in this way, market pressure can build up, reputational effects set in, and public scrutiny via analysts, academics, and NGOs unfold its incentivising force.

Shortcomings in the AI Act Concerning Climate Reporting, and Policy Proposals

Shortcomings	Policy Proposals
1. Inference Energy Consumption Exclusion	Explicitly include inference in energy reporting obligations in Annexes XI and XII.
2. Indirect Emissions and Water Consumption	Extend reporting obligations to include water consumption and indirect GHG emissions from AI applications.
3. Fine-Tuning Uncertainty	Clarify uncertainty of reporting obligations by tying them to computational cost and training mechanisms.
4. Open-Source Models	Revoke exemption to ensure comprehensive climate reporting.
5. Lack of Standard Reporting Methodology	Measure energy consumption at the cumulative server level, with separate PUE.
6. Lack of Public Access to Energy Data	Make all climate-related disclosures publicly available to foster transparency and market accountability.

17 The Green Grid, 'IT & Power Efficiency Survey Results (for EUC)', 26.10.2023, <https://www.thegreengrid.org/en/resources/library-and-tools/572-IT-%26-Power-Efficiency-Survey-Results-%28for-EUC%29> (accessed 6.6.2025).

18 Commission Delegated Regulation (EU) 2024/1364 on the first phase of the establishment of a common Union rating scheme for data centres [2024] OJ L 2024/1364.

V. The Code of Practice

Importantly, however, the AI Act rules on climate effects have just received a major update as they constitute one part of the recently published General-Purpose AI Code of Practice (CoP).¹⁹ It complements the AI Act by offering a voluntary, yet concrete framework for providers of general-purpose AI models to demonstrate compliance with the regulatory obligations in the GPAI chapter of the AI Act. Published on July 10, 2025, the Code emerged from a transparent, multi-stakeholder drafting effort led by the European AI Office and independent researchers with participation from nearly 1,000 experts across industry, civil society, academia, rights-holders, and Member State representatives. It addresses three core chapters—Transparency, Copyright, and Safety & Security—each mapping to corresponding obligations under Articles 53 and 55 of the AI Act.

The Transparency chapter provides a user-friendly Model Documentation Form to aid providers in meeting the AI Act's requirements for technical documentation, copyright information, and training data summaries under Article 53. The Copyright chapter guides providers through implementing policies necessary to comply with EU copyright law obligations under the same article. The Safety & Security chapter applies only to providers of the most advanced, high-risk models—those subject to Article 55—and offers concrete measures for systemic risk taxonomy, risk assessment, incident reporting, and cybersecurity.

The Code is now undergoing adequacy assessment by the AI Office and Member States. Once endorsed (Art. 56(6) AI Act), it will serve as a practical early implementation tool to ease compliance burden. Signatories will benefit from a presumption of conformity (Art. 53(4) and 55(2) AI Act). Mistral and OpenAI were the first companies who committed to indeed sign and abide by the CoP.

The CoP, in its Transparency chapter, introduces reporting requirements on energy consumption—notably during both the training and inference phases of GPAI models.²⁰ Providers must disclose the amount of energy used for training, expressed in megawatthours (MWh) and recorded with at least two significant figures (e.g. 1.0×10^2 MWh).

In the absence of a delegated act pursuant to Article 53(5) of the AI Act specifying standardised measurement methodologies, providers must describe the method used to measure or estimate training energy consumption. This includes explaining whether the estimation is based on direct measurements or indirect calculations derived from known computational resources. If critical data are unavailable, providers must identify the specific information they lack.

The real surprise comes in the section on inference. Published after the initial publication of this paper as a Working Paper, the CoP takes up our general suggestion to interpret the AI Act in a way compatible with inference reporting, but based on a different legal reasoning. The CoP requires disclosure of the benchmarked amount of computation, reported in floating point operations and recorded with at least two significant figures (e.g. 5.1×10^{17} FLOPs). Providers must describe the task used to benchmark inference computation—such as generating 100,000 tokens—and detail the hardware setup, for example using 64 Nvidia A100 GPUs. As a basis in the AI Act, the CoP cites Annex XI Section 1 para. (2)(e), which requires the disclosure of the “known or estimated energy consumption of the model.”²¹ Indeed, this requirement, read on its own, does not limit the consumption to training. How-

ever, as mentioned (Part II.), the heading of Section 1 para. 2 suggests that the item indeed may refer to training only (“development”). If this reading was adopted, then our interpretation provides a further justification for keeping inference in the CoP, and requiring reporting more generally from GPAI providers: it forms part of the “technical means needed to integrate the GPAI model into AI systems” (Annex XI Section 1 para. (2)(a), Annex XII para. (2)(a)). This reading has the additional benefit that the benchmarked inference reporting must be disclosed not only to the AI Office and National Authorities (Annex XI), but also to downstream deployers (Annex XII). The lack of access for the general public, however, remains problematic.

Finally, we do note, however, that some information on energy consumption may ultimately be accessed by researchers. While the public, summarised versions of the Framework and Model Reports (Measure 10.2, Safety and Security Chapter) will likely not contain any information on energy, independent evaluators receive free access to the most capable model versions to facilitate post-market monitoring (Measure 3.4, Safety and Security Chapter). This not only mirrors, to a certain extent, Art. 40 DSA, but also provides researchers with some access to models—and potentially their inference calculations as it can be argued that environmental effects do form part of systemic risks.²² However, the measure only applies, like the entire Safety and Security Chapter, only to GPAI models with systemic risk, which limits its effectiveness in tracking model inference energy consumption.

Overall, the CoP does address the problem of inferences in a laudable way; and it provides some limited external oversight via researchers. But generally, the policy suggestions we have sketched remain to be enacted, through guidelines, implementing acts, a revision of the Code of Practice, or even the AI Act itself.

VI. Conclusion

This paper tackles some of the complexities at the intersection of AI, climate and regulation. The AI Act does contain significant climate reporting obligations. By drawing on technical and legal research, we show that they contain too many loopholes, and are difficult to operationalize. Perhaps most importantly, even though recent research has shown inference to be a major driver of AI-related GHG emissions, this key area is omitted from the AI Act. A novel interpretation of the Act's reporting obligations might bring inference back within its scope. Furthermore, none of the climate disclosures are initially open to the public. We suggest changing this urgently to kickstart market pressure, induce reputational effects among consumers, and enable crucial public scrutiny, e.g. by academics and NGOs, beyond Measure 3.4 of the CoP's Safety and Security Chapter. For model providers, our transparency proposals enable essential global competition for green innovation in AI models, which currently exists through very indirect and intransparent economic measures, such as electricity prices or (local) water regulations. For

19 European Commission, ‘The General-Purpose AI Code of Practice’, 10.7.2025, <https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai> (accessed 15.7.2025).

20 Nuria Oliver/Rishi Bommasani, ‘Transparency Chapter’ in the Code of Practice for General-Purpose AI Models, 9, 10.7.2025, <https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai> (accessed 16.7.2025).

21 Oliver/Bommasani, (n 20), 9, footnote n 1.

22 See Atoosa Kasirzadeh et al., ‘Too Big to Regulate? Systemic Risk in Market, AI, and Platform Regulation’, Working Paper, forthcoming (on file with the authors).

consumers, both organisations and individual end-users, this allows for more informed decisions from both environmental and reputational standpoints. Additionally, downstream consumers can select the most cost-effective model for their specific use cases, as inference energy is directly linked to power consumption and hardware expenses.

However, climate reporting can only be a first step in addressing the massive and fast-rising environmental impact of AI

models and systems. It must be complemented by substantive obligations, including sustainability risk assessment and management, renewable energy targets for data centres, and potentially even (tradable) caps on the energy and water consumption of data centres and similar major consumption drivers in the AI value chain²³. ■

23 Hacker (n 1) 345.

Michael Rose, Jörg Hoffmann, Dietmar Harhoff*

Digital Consumer Law, Competition and the (Un-)Informed Consumer: Evidence from a Survey among German Consumers

I. Introduction

In recent years, policymakers at the European level have enacted legislation aimed at harnessing digital possibilities to improve consumer welfare, enhance competition, and reduce costs. Examples include the Second Payment Services Directive (PSD2) or the Digital Content Directive.¹ Policymakers were motivated by a desire to spur competition, stimulate the EU digital economy and to achieve a high level of consumer protection.

However, action on the part of consumers requires awareness of their rights. As *Ben-Shahar and Schneider* put it, '[d]isclosures cannot use information until they get it. But they often do not get it because they do not know that it exists, where it is, or how to find it, and they often doubt that they need it or that it would justify its acquisition costs'.²

To probe the awareness of the new consumers rights, we conducted a representative survey of German consumers in October 2024. Participants are faced with two simple everyday scenarios. All scenarios share two key characteristics: first, they reflect changes in the legal framework that originated at the EU level, and second, the aim to foster competition. These scenarios include whether banks are required to share data during a bank account transfer and whether sellers of electronic hardware must provide security updates.

In all scenarios fewer than half of respondents know the correct answer: 47 % of German consumers know that banks are required to share data with a new bank upon request in the case of a bank account transfer. 26.7 % of German consumers are aware that laptop sellers are responsible for providing security updates—almost two years after the policy-change. In both cases, about 25 % of Germans don't know.

Since the new consumer rights aim at increasing competition, the unawareness of the new rights likely has repercussions on competition. For instance, PSD2 posits that Banks have to share account data to facilitate switching financial service providers. Behavioral factors, especially the lack of willingness to switch providers, likely result in low customer switching rates. If customers did not learn that switching costs reduced thanks to PSD2, they might perceive the switching costs as too high, and might thus refrain from changing their bank.³ As a result, incumbents are at an advantage, making the active market intervention by the state in this form seem more justified.

Our figures also cast doubt on the reality of the "informed customer" who weighs their options and chooses the utility-enhancing one. The expectation, that consumers deal in a

reasonably "reasonably well-informed and reasonably observant and circumspect" (Recital 18 of the Unfair Commercial Practices Directive (UCP)) underlies much of the European Union's commercial law.⁴

Previous studies have challenged the notion of the "informed consumer" in other domains. In the context of health literacy in Germany for instance, *Kilian, Liebig, and Manthey* suggest that this assumption may be flawed as more educated individuals systematically possess greater knowledge about health risks.⁵ *Lee* demonstrates that the 'tendency of over-insurance' of US consumers is partly the result 'from the lack of information about the probability of loss'.⁶

Uninformed consumers may also at the same time be the reason and the cause of sub-optimal EU competition policy: they may participate less in the policy-making because they do not understand it, and they may not understand it, because only experts engage in competition policy-making.⁷

In contrast to the 'informed consumer', the concept of the 'vulnerable consumer' shifts attention away from the desired

* Michael Rose, Max Planck Institute for Innovation and Competition, Munich, email: michael.rose@ip.mpg.de; Jörg Hoffmann, Max Planck Institute for Innovation and Competition, Munich, email: joerg.hoffmann@ip.mpg.de; Dietmar Harhoff, Max Planck Institute for Innovation and Competition, Munich, email: dietmar.harhoff@ip.mpg.de.

1 Directive (EU) 2015/2366 of the European Parliament and of the Council of 25 November 2015 on payment services in the internal market, amending Directives 2002/65/EC, 2009/110/EC and 2013/36/EU and Regulation (EU) No 1093/2010, and repealing Directive 2007/64/EC, OJ L 337, 23.12.2015; Directive (EU) 2019/770 of the European Parliament and of the Council of 20 May 2019 on certain aspects concerning contracts for the supply of digital content and digital services, OJ L 136, 22.5.2019.

2 Omri Ben-Shahar, Carl E. Schneider, 'The Failure of Mandated Disclosure' (2011) University of Pennsylvania Law Review 159 (3), 647, 709.

3 Paul Klemperer, 'The Competitiveness of Markets with Switching Costs' (1987) The RAND Journal of Economics 18 (1), 138, Steven A Sharpe, 'The Effect of Consumer Switching Costs on Prices: A Theory and Its Application to the Bank Deposit Market' (1997) Review of Industrial Organization 12 (1), 79–94.

4 Natalie Helberger, Marjin Sax, Joanna Strycharz, and Hans-Werner Micklitz, 'Choice Architectures in the Digital Economy: Towards a New Understanding of Digital Vulnerability' (2022) Journal of Consumer Policy 45 (2), 175–200.

5 Caroling Kilian, Moritz Liebig, Jakob Manthey, 'Education-based Differences in Alcohol Health Literacy in Germany' (2025) Drug and Alcohol Review 44 (2), 555–62.

6 Kangoh, Lee, 'Consumer Perception, Information Provision, and Regulation of Insurance Markets' (2017) Journal of Regulatory Economics 51 (1), 1–17.

7 Firat Cengiz, 'Legitimacy and Multi-Level Governance in European Union Competition Law: A Deliberative Discursive Approach' (2016) JCMS: Journal of Common Market Studies 54 (4), 826–45.