

cerre



AGENTIC AI AND CONSUMER PROTECTION

ISSUE PAPER

July 2026

Christoph Busch

As provided for in CERRE's bylaws and procedural rules from its "Transparency & Independence Policy", all CERRE research projects and reports are completed in accordance with the strictest academic independence.

The CERRE Regulatory Lab series, of which this note is part, received the support and/or input from several CERRE members including Apple, Booking.com, and Microsoft. Each draft note is discussed in Chatham House meetings between the CERRE corporate members, national regulatory authorities, EU institutions and civil society organisations. However, they bear no responsibility for the contents of this report. The views expressed in this CERRE report are attributable only to the authors in a personal capacity and not to any institution with which they are associated. In addition, they do not necessarily correspond either to those of CERRE, or of any sponsor or of members of CERRE.

© Copyright 2026, Centre on Regulation in Europe (CERRE)

info@cerre.eu – www.cerre.eu



Table of Contents

<u>TABLE OF CONTENTS</u>	<u>1</u>
<u>1. INTRODUCTION</u>	<u>2</u>
<u>2. IS EU CONSUMER LAW AGENT-READY?</u>	<u>3</u>
<u>3. ADAPTING CONSUMER LAW FOR THE AGENTIC ECONOMY</u>	<u>9</u>
<u>4. HOW TO TRANSITION TOWARDS THE NEW AGENTIC CONSUMER LAW?</u>	<u>11</u>
<u>ABOUT CERRE.....</u>	<u>15</u>
<u>ABOUT THE AUTHOR</u>	<u>16</u>



1. Introduction

Online shopping is going through a profound transformation.¹ According to a recent survey, an increasing number of consumers no longer start their customer journey with a search engine like Google or an online marketplace like Amazon, but rather with ChatGPT, Gemini, or Claude: 46% of consumers in the European Union already use artificial intelligence (AI) tools to discover new products; 63% use AI tools to compare offers and learn more about product features.²

But AI assistants that support shopping with recommendations are just the beginning of a more fundamental structural change affecting online commerce. The next stage of evolution is agentic AI. Unlike chatbots, AI agents can not only talk but also take action. They do more than just assist with product searches by providing purchase recommendations and product comparisons. AI agents can independently perform complex tasks, such as purchasing products from online stores, booking airline tickets, or switching electricity providers. In the near future, consumers could therefore be able to delegate more of their transactional decisions to AI agents. Even though this may still sound like science fiction, forecasts suggest that agentic shoppers will soon play an increasingly important role in digital markets. According to some estimates, between 10% and 20% of e-commerce transactions could be handled by an AI agent in 2030.³ Other forecasts suggest that agentic commerce could orchestrate USD 3 trillion to USD 5 trillion worth of transactions globally by 2030.⁴

Against this background, in May 2026, the Centre on Regulation in Europe (CERRE) organised a Regulatory Lab on Agentic AI and Consumer Protection. The event brought together stakeholders and policy makers to discuss whether EU consumer law is ready for the rise of AI agents and how consumer law should develop in order to enable innovative business models built on agentic AI while, at the same time, ensuring a high level of consumer protection.

Building on these conversations and earlier research⁵, this note outlines some of the key issues that were discussed at the Regulatory Lab and develops further ideas on how consumer policy should respond to the challenges created by the ‘agentic turn’ in digital markets.

The note is structured into three parts. Following this introduction, the first part explores whether EU consumer law is agent-ready and explains how AI agents challenge the existing consumer protection framework. The second part looks ahead and offers some initial recommendations for building a consumer law rules that work for humans and machines. Finally, the third part outlines how to transition to the next generation of consumer law that is ready for the agentic economy.

¹ See, generally, Mark Bartholomew and Shmuel I. Becher, *The End of Shopping*, 68 *William & Mary Law Review* (forthcoming).

² McKinsey, *Europe’s agentic commerce moment, Decision influence is here; execution is coming*, 2 March 2026, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/europes-agentic-commerce-moment-decision-influence-is-here-execution-is-coming>

³ Morgan Stanley, *Here Come the Shopping Bots*, 8 December 2025, <https://www.morganstanley.com/insights/articles/agentic-commerce-market-impact-outlook>.

⁴ McKinsey, *The agentic commerce opportunity: How AI agents are ushering in a new era for consumers and merchants*, 17 October 2025, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-agentic-commerce-opportunity-how-ai-agents-are-usher-in-a-new-era-for-consumers-and-merchants>

⁵ Christoph Busch, *Consumer Law for AI Agents*, 26 *German Law Journal* 1367 (2025).



2. Is EU consumer law agent-ready?

This part briefly outlines how AI agents challenge the existing consumer protection framework and discusses how AI agents could mitigate some consumer vulnerabilities while possibly exacerbating others.

Defining AI agents

Before discussing whether EU consumer law is agent-ready, it is necessary to clarify what we talk about when referring to AI agents. There is not yet a generally agreed definition of AI agents.⁶ According to a leading textbook on AI, an agent is, in a very general sense, "anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators".⁷ However, for the purposes of this note, which focuses on issues of consumer protection, this definition (which would include a thermostat) is too broad. A more suitable working definition can be borrowed from a recent report on AI regulation by the UK Department for Science, Innovation and Technology: AI agents are "autonomous AI systems that perform multiple sequential steps – sometimes including actions like browsing the internet, sending emails, or sending instructions to physical equipment—to try and complete a high-level task or goal."⁸

Beyond human-centric consumer law

In order to understand why the rise of agentic AI challenges the very foundations of EU consumer law, it is necessary to recall a feature of EU consumer law that seems so self-evident that it is often not explicitly stated: EU consumer law in its current form is *human-centric*. It is designed with the needs and weaknesses of human consumers in mind. Put differently, the design of consumer law rules is based on the premise that purchasing decisions are made by humans. Therefore, EU consumer law aims to protect human consumers and ensure their autonomy.

However, this premise is challenged by AI in general and AI agents in particular. The agentic economy will be more complex and diverse than the old human-centric economy, for which the existing consumer laws were built. In the brave new world of AI-driven retail, human consumers will coexist with "augmented consumers" (who are assisted by AI tools such as ChatGPT or Claude) and AI agents acting autonomously as "algorithmic consumers" or "Custobots" on behalf of humans.⁹

⁶ For an overview see OECD, The agentic AI landscape and its conceptual foundations, February 2026, https://www.oecd.org/content/dam/oecd/en/publications/reports/2026/02/the-agentic-ai-landscape-and-its-conceptual-foundations_a9d4b451/396cf758-en.pdf; see also Melissa Heikkilä, What Are AI Agents?, MIT Tech. Rev., 5 July 2024, <https://www.technologyreview.com/2024/07/05/1094711/what-are-ai-agents/>.

⁷ Stuart Russell and Peter Norvig, Artificial Intelligence, 4th ed. 2021, p. 54.

⁸ UK Department for Science, Innovation and Technology, A pro-innovation approach to AI regulation: government response (6 February 2024), <https://www.gov.uk/government/consultations/ai-regulation-a-pro-innovation-approach-policy-proposals/outcome/a-pro-innovation-approach-to-ai-regulation-government-response>; see also Friso Bostoen and Jan Krämer, AI Agents and ecosystem contestability (November 2024), p. 5.

⁹ See e.g., Michal S. Gal and Niva Elkin Koren, Algorithmic Consumers 30 Harvard Journal of Law & Technology 309 (2017); see also Don Scheibenreif and Mark Raskino, When Machines Become Customers, 3rd ed. 2025.



Against this background, the shift towards agentic commerce challenges several core concepts and regulatory tools of consumer law – from mandatory disclosures and design-based consumer law standards to rules against unfair commercial practices.

Information disclosures and consent requirements

For starters, the rise of agentic commerce raises the question of how pre-contractual information duties and consent requirements should be redesigned once an AI agent, rather than a human, is the addressee of the disclosure. Unlike humans, AI agents are not constrained by the cognitive limitations that make information overload a problem for human consumers. They can process far larger volumes of information and, more importantly, filter out what is actually relevant to the individual consumer's preferences.

In this sense, AI agents have the potential to address one of the most persistent criticisms of the "information paradigm" in consumer law, namely that the long lists of standardised pre-contractual disclosures mandated by EU directives tend to overwhelm human consumers rather than enabling informed decisions.¹⁰ Agents could also help reduce information asymmetries more broadly by drawing on a wider range of data sources than any human consumer could realistically consult, including customer reviews, user forums, product tests, and online media, thereby improving comparison shopping across brands and retailers.

Another question concerns consent requirements. Article 8(2) of the Consumer Rights Directive (CRD) requires that consumers "explicitly acknowledge" a payment obligation, typically via a button. But it remains unclear whether an AI agent can meaningfully perform this acknowledgement on the consumer's behalf, or whether a separate confirmation step involving the human is still required.

From graphical user interface to the protocol layer

Today, many provisions of EU consumer law presuppose that a human consumer is clicking through a graphical user interface (GUI). Examples of design-based provisions include the above mentioned "buy now" button (Article 8(2) CRD), the recently introduced right for a consumer to withdraw from an online contract within a specified period (Article 11a CRD) or the "compliance by design" rule for online marketplaces (Article 31(1) of the Digital Services Act or DSA).

However, the assumption on which these rules are built no longer holds once agentic transactions move from graphical interfaces to the protocol layer where agent-to-agent (A2A) communication via APIs will replace browser-based interaction. While so-called 'web browsing agents' still mimic human click-through behaviour, other agents interact directly through APIs using standardised communication protocols such as the Universal Commerce Protocol (UCP) or the Agentic Commerce Protocol (ACP). In such a scenario, the agents never encounter buttons or prominently displayed

¹⁰ See, generally Geraint Howells, Christian Twigg-Flesner and Thomas Wilhelmsson, Rethinking EU Consumer Law, 2018, pp. 94-112; see also Anne-Lise Sibony, Can EU Consumer Law Benefit from Behavioural Insights? An Analysis of the Unfair Practices Directive, 6 European Review of Private Law 901 (2014).



disclosures. This will make compliance with GUI-based requirements structurally impossible.

This raises the question of whether functionally equivalent obligations — for instance a ‘withdrawal API’ alongside the mandatory withdrawal button — should be introduced to preserve the substantive protection these rules were meant to deliver, even where the interface itself changes. Taken together, these developments suggest that the existing EU consumer law requirements may need to be translated into agent-ready requirements that operate on the protocol layer rather than on graphical interfaces.

Conflicts of interest

A further risk specific to the agentic economy concerns the alignment of incentives between consumers and the providers of AI agents. Possible conflicts of interest could arise where the provider of the AI agent has commercial interests that diverge from those of the consumer it is supposed to represent. More specifically, an agent operated by a business that also sells or promotes its own products or those of paying partners may be structurally incentivised to steer purchasing decisions accordingly rather than solely pursuing the consumer's actual interests. Such risks of self-preferencing behaviour echo concerns already addressed, in the context of platform regulation, by the prohibition on self-preferencing under Article 6(5) DMA for designated gatekeepers.¹¹

Unlike a search engine ranking, where at least some transparency requirements already exist, an agent's internal decision-making process is typically opaque to the consumer, making it correspondingly harder to detect whether a recommendation solely reflects the consumer's stated preferences or whether and how the commercial interests of the agent's operator are taken into account. Addressing this risk may require dedicated conflict-of-interest or “loyalty by design” obligations for agent providers combined with transparency requirements about any commercial relationships that could influence an agent's recommendations or transactional decisions.¹²

Manipulation by AI agents

The more consumers rely on AI agents that have access to their purchasing history, calendars, playlists, or email accounts, the more powerful — and potentially the more manipulative — these agents can become. Intimate knowledge of a consumer's preferences and vulnerabilities could enable highly targeted recommendations that go well beyond anything achievable in today's digital markets.

From this perspective, some observers have raised the concern that AI agents could become “manipulation engines” rather than tools of consumer empowerment.¹³ This view sits alongside established concerns about algorithmic manipulation and algorithmic harm more broadly.¹⁴ Some protection may already be available through existing algorithmic-transparency rules for rankings and recommender systems, such as Article 6a(1)(a) CRD and Article 7(4a) of the Unfair Commercial Practices Directive (UCPD), both of which are formulated in a technology-neutral manner and are

¹¹ See Friso Bostoen and Jan Krämer, *Is the DMA ready for agentic AI?*, CERRE Report, July 2025.

¹² See Article 12 ELI Guiding Principles and Model Rules on Digital Assistants for Consumer Contracts (requiring providers of AI agents to disclose potential conflicts of interest).

¹³ Kate Crawford, *AI Agents Will Be Manipulation Engines*, 23 December 2024, [Wired.com](https://www.wired.com/story/ai-agents-will-be-manipulation-engines/).

¹⁴ See Oren Bar-Gill and Cass R. Sunstein, *Algorithmic Harm: Protecting People in the Age of Artificial Intelligence*, Oxford 2025; Cass R. Sunstein, *Manipulation*, Cambridge 2025.



therefore likely to extend to agent-based recommendations. In the case of other, platform-specific transparency obligations, such as Articles 27 and 38 DSA, it is less clear whether they apply to AI agents. The answer to this question depends on whether a given AI agent can be considered as an online platform under Article 3(i) DSA.

A further layer of protection could come from the general prohibition on AI-enabled manipulation in Article 5(1)(a) AI Act, although this provision is formulated in comparatively vague terms and its practical application to agentic commerce remains untested. Ensuring that agents act in consumers' best interests will therefore likely require a combination of sharpened transparency obligations, substantive limits on manipulative agent design, and attention to the underlying incentive structure of the entity that provides the agent.

Manipulation of agents

There is reason to expect that AI agents will be less susceptible than human consumers to at least some of the manipulative techniques traditionally captured under the heading of harmful online choice architectures or "dark patterns".¹⁵ Agents are presumably immune to "confirm shaming" and emotional advertising messages and may also be more resistant to "click fatigue," "nagging," or artificial urgency created through countdown timers and scarcity notifications. If this holds true, the widespread adoption of AI agents could meaningfully reduce consumers' exposure to the choice-architecture techniques that current instruments such as Article 25 DSA were designed to combat.

However, this resilience should not be overstated. Being less vulnerable to dark patterns aimed at humans does not mean AI agents are invulnerable to manipulation altogether. More empirical research on how agents actually respond to persuasive and manipulative interface design is needed before firm conclusions can be drawn, particularly given that findings from the emerging fields of "machine psychology" and "AI psychometrics" remain inconclusive as to whether large language models replicate human-like decision biases.¹⁶ Policy makers should therefore treat the mitigation of dark-pattern risks as a plausible but unproven benefit of agentic commerce.

In addition, the threat may shift toward a new type of 'agentic dark patterns'. More specifically, third parties — rogue traders, hackers, or other bad-faith actors — could manipulate the AI agents themselves in order to trick them into unwanted transactions or the disclosure of sensitive data. Because AI systems perceive their environment differently from humans, they can be manipulated in ways that would have no effect on a human consumer. One notable example are so-called "adversarial image perturbations", i.e. subtle, human-imperceptible modifications to images designed to alter how an agent categorises what it "sees".¹⁷ Further risks may arise from prompt injection attacks that embed malicious commands capable of hijacking an agent's behaviour, for example to authorise

¹⁵ See Christoph Busch and Amelia Fletcher, Harmful Online Choice Architecture, CERRE Report, May 2024.

¹⁶ See, e.g., Thilo Hagendorff et al. Machine Psychology (2024), <https://arxiv.org/pdf/2303.13988>; Max Pellert et al., AI Psychometrics: Assessing the Psychological Profiles of Large Language Models Through Psychometric Inventories, 19 Persps. on Psych. Sci. 808, 808-26 (2024).

¹⁷ See, e.g., Kevin Eykholt et al., Robust Physical-World Attacks on Deep Learning Models (2017), <https://arxiv.org/pdf/1707.08945>.



unauthorised transfers of funds or leak confidential user data.¹⁸

This raises the question of whether the rules on misleading and aggressive commercial practices in the UCPD as well as the prohibition of dark patterns in Article 25 DSA, that were designed to protect human consumers from the exploitation of behavioural biases, will still serve their purpose if transactions are delegated to AI agents.

Sector specific consumer protection laws

Beyond the horizontal instruments discussed above, such as the CRD, the UCPD and the DSA, the rise of agentic commerce also raises questions for sector-specific consumer protection regimes such as the Package Travel Directive (PTD) or the Insurance Distribution Directive (IDD), which likewise rest on the assumption that a human consumer receives standardised information and makes an informed choice at a clearly identifiable moment.

These regimes often contain even more detailed, sector-tailored disclosure and form requirements than the horizontal instruments, which could either compound the ‘information paradigm’ problem discussed above — since agents may be well suited to process detailed sectoral disclosures that would overwhelm a human traveller or policyholder — or create additional friction where sector-specific consent or documentation formats presuppose direct human interaction. In addition, the question arises how the terms and concepts defined by sector specific legislation map onto agentic transactions, (e.g., "package" (Art. 3(2) PTD), "linked travel arrangement" (Art. 3(5) PTD), "organiser" (Art. 3(8) PTD), "insurance intermediary" (Art. 2(3) IDD) or "ancillary insurance intermediary" (Art. 2(4) IDD)). This is an area that merits closer scrutiny beyond the forthcoming Digital Fairness Act which focuses on horizontal instruments of consumer protection.

Liability for unexpected actions of AI agents

A further open question concerns liability for transactions or other actions that an AI agent takes in ways the consumer did not anticipate or authorise — for instance, purchasing the wrong product, exceeding a spending limit due to technical malfunctions or hallucinations, or falling victim to a prompt injection attack that triggers an unauthorised transfer of funds.

While a growing body of scholarship has analysed these questions through the lens of agency law,¹⁹ exploring how doctrines of authority and attribution developed for human agents might apply to AI systems, existing EU consumer law offers little specific guidance on how the resulting risk should be allocated between consumer, agent provider, and trader.

Given the technical vulnerabilities discussed above, questions of liability are also likely to intersect with cybersecurity and product-safety frameworks such as the AI Act and the Product Liability Directive, rather than being resolved by contract or consumer law alone. As agentic commerce becomes more widespread, clarifying this liability architecture — and in particular whether existing

¹⁸ See, e.g., Peter Yong Zhong et al., RTBAS: Defending LLM Agents Against Prompt Injection and Privacy Leakage (2025), <https://arxiv.org/pdf/2502.08966>.

¹⁹ See, e.g., Noam Kolt, Governing AI Agents, 101 Notre Dame L. Rev. (forthcoming); see also Anat Lior, AI Entities as AI Agents: Artificial Intelligence Liability and the AI Respondeat Superior Analogy, 46 Mitchell Hamline Law Review 1043 (2020).



rules on unauthorised payment transactions or defective digital content provide adequate protection
— will be essential to maintaining consumer trust in delegated decision-making.



3. Adapting consumer law for the agentic economy

Building on the analysis of the insufficient agent-readiness of EU consumer law, this part explores how to build a consumer law that works for humans and machines and how to transition towards the new agentic consumer law.

The core challenge for the next generation of EU consumer law is not to invent entirely new legal concepts, but to ensure that the existing framework of consumer protection continues to function once the addressee of these rules is no longer necessarily human, or at least where there is a new intermediary layer comprised by one or more AI agents. This may require a number of adjustments that will be sketched out in the following sections.

Ensuring disclosure, consent, and unfair commercial practices rules remain effective

A key task will be to ensure that existing instruments of consumer protection — disclosure, consent, and protection against unfair commercial practices — will remain effective even if the settings and actors involved in consumer transactions change. As discussed above, this will likely require differentiated approaches for each pillar. Disclosure rules may need to shift from limiting information (to avoid overload) toward providing richer, machine-readable information that agents can filter on the consumer's behalf. Consent mechanisms may need to be reconceived to clarify that an 'explicit acknowledgement' of a payment obligation can be performed by an agent rather than only by a human clicking a button. In addition, unfair-practices rules will need to capture both the traditional risk of manipulating humans and the emerging risk of manipulating agents.

What unites these adaptations is a shift from prescribing specific human-facing mechanisms toward defining the underlying regulatory objective — enabling informed decisions, ensuring genuine consent, protecting freedom from manipulation — in a way that remains technology-neutral enough so that they can be applied to agentic transactions. This is less a matter of drafting entirely new doctrines than of testing each existing concept against the question: Does this rule still protect the consumer's autonomy when an AI agent is stepping in between them and the trader?

Compliance by protocol: integrating EU consumer law into the protocol layer

As agent-to-agent communication increasingly replaces browser-based transactions on graphical user interfaces, an important question is whether consumer protection can still be achieved purely through substantive legal rules addressed to traders and agent providers, or whether it must additionally be embedded directly into the technical protocols and communication standards that structure agentic commerce — such as the Universal Commerce Protocol (UCP) or the Agent Payment Protocol (AP2).

If information duties, consent mechanisms, and dispute-resolution pathways can be encoded as mandatory elements within these protocols themselves, compliance could become largely self-enforcing, in the same way that today's 'withdrawal button' requirement is enforced through interface design. This would mark a significant evolution of the 'compliance by design' approach, moving from



design requirements addressed to individual traders' interfaces toward design requirements embedded in the shared technical infrastructure and protocols that multiple traders and agents rely upon ('compliance by protocol').

Differentiating between types of agentic systems

A further design question is whether EU consumer law should apply uniformly to all agentic transactions, or whether it should differentiate based on the degree of human involvement at the moment of the transaction — distinguishing, for instance, between 'human-in-the-loop' scenarios where the consumer reviews and confirms each transaction proposed by the agent, and 'human-out-of-the-loop' scenarios where the agent is authorised to conclude transactions autonomously within pre-set parameters.

Where a human remains in the loop, the case for maintaining relatively light-touch, disclosure-based protections at the point of the transaction itself may be stronger, since the consumer retains a final opportunity to catch problematic terms or bad agent decisions. Where the human is out of the loop, however, correspondingly stronger ex-ante design requirements would seem necessary to compensate for the absence of a final human check. In addition, it might be necessary to introduce mandatory circuit-breakers into the agentic workflow, for example, requiring renewed human confirmation above a certain transaction value or for certain product categories.

Further differentiations may be necessary between different types of AI agents, e.g., agents embedded directly within a marketplace's own infrastructure, independent third-party shopping agents that browse and transact across multiple external sites or agents deeply integrated into a device's operating system (OS).

These different deployment models raise distinct regulatory concerns: a marketplace-embedded agent may raise self-preferencing risks similar to those already addressed under Article 6(5) DMA for designated gatekeepers, a third-party shopping agent raises questions about liability allocation across multiple traders and platforms it interacts with, and an OS-embedded agent raises particularly acute concerns about market gatekeeping, since it could become the primary or even sole access point through which a consumer's purchasing decisions are mediated.

A uniform regulatory approach that does not distinguish between these models risks either being too permissive toward the most concentrated deployment (OS-level integration) or imposing unnecessary burdens on more open, competitive models (independent third-party agents). Rather, it might be advisable to develop a differentiated typology of AI agents that can serve as a basis for risk-based regulation for a heterogeneous agent landscape.



4. How to transition towards the new agentic consumer law?

The final section of this note addresses the question of how to manage the transition towards an agent-ready consumer law: is incremental change or a more revolutionary approach more desirable? What is the right timing for regulatory intervention? What are the right regulatory tools? What risks could flow from regulatory inaction?

Incremental versus revolutionary reform

At the time of writing, the European Commission is preparing its proposal for the Digital Fairness Act (DFA) which is expected in the third quarter of 2026.²⁰ While its name seems to suggest a major reform of EU consumer law, it currently seems more likely that the DFA will pursue targeted amendments of existing EU consumer protection instruments rather than a wholesale rewrite of EU consumer law.

As part of the update to EU consumer law, the DFA should also address the challenges described above that the rise of agentic commerce poses for consumer law. Since the transition towards agentic commerce is still in its early stages and the future development of business models is not yet clearly foreseeable, an incremental reform strategy appears the more sensible and politically realistic path.

As a first step, it should be clarified how existing regulations apply to agentic transactions. This could be done, for example, through an updated version of the guidance on the UCPD²¹ and the CRD²². Such an approach would build on the technology-neutral wording already present in many provisions of the UCPD and CRD, extending established regulatory approaches to cover agentic transactions through interpretation and targeted clarification rather than fundamentally new categories of rules. At the same time, the most glaring gaps should be addressed to ensure that consumers retain control in agentic transactions. The DFA offers a good opportunity to do so, which should not be missed.²³

In a second step, the existing regulatory framework should be further developed to ensure that consumer protection is also guaranteed at the protocol layer. To achieve this, it will be necessary to translate the existing design-based provisions into protocol-based rules. However, this can only be done effectively once there is more clarity regarding which communication standards will shape the future of agentic commerce.

In parallel, under the Apply AI Strategy,²⁴ the European Commission has set up an expert group on contracting to assist the European Commission in developing horizontal model contract terms for AI

²⁰ See Christoph Busch, Amelia Fletcher and Michèle Ledger, Towards an EU Consumer Law Fit for the Digital Age, CERRE Report, February 2026.

²¹ Guidance on the interpretation and application of Directive 2005/29/EC of the European Parliament and of the Council concerning unfair business-to-consumer commercial practices in the internal market, OJ C 526, 29.12.2021, pp. 1–129.

²² Guidance on the interpretation and application of Directive 2011/83/EU of the European Parliament and of the Council on consumer rights, OJ C 525, 29.12.2021, pp. 1–85.

²³ Christoph Busch, Enabling Innovation and Protecting Consumers in the Agentic Economy: Why the Digital Fairness Act Should Regulate Agentic AI, in: Lohsse, Schulze and Staudenmayer (eds.) AI Contracting: Münster Colloquia on EU Law and the Digital Economy IX, Baden-Baden 2026, pp. 151-168.

²⁴ European Commission, Apply AI Strategy, 8.10.2025, COM(2025) 723 final.



contracting and guidance for the choice of AI contracting systems from the users' perspective. The forthcoming model contract terms, that will draw inspiration from the UNCITRAL Model Law on Automated Contracting, could be an important element of a future legal infrastructure for agentic commerce.

Addressing the most glaring gaps

Among the gaps identified in this policy brief, several stand out as particularly urgent because they concern the basic legal infrastructure without which agentic commerce cannot function with legal certainty at all: the absence of clear rules on the validity and enforceability of contracts concluded via AI agents, the absence of any provision protecting AI agents themselves from manipulation (in contrast to the well-developed dark-pattern rules protecting humans), and the unresolved question of how 'explicit acknowledgement' of payment obligations can be satisfied in agent-mediated transactions. Further issues that are particularly relevant at the payment layer concern authorisation and customer authentication.

These gaps are foundational in the sense that market participants cannot confidently build agent-based business models, and consumers cannot be confident that agents are adequately protected, until they are resolved — unlike more incremental questions, such as exact disclosure formatting, which can more easily be addressed through subsequent guidance or case law.

Prioritising these structural questions would allow the DFA to establish the basic legal infrastructure for agentic commerce, leaving more detailed implementation questions to delegated acts, guidance, or future amendments. Left unaddressed, these gaps risk leaving both consumers and businesses in uncertainty. This could undermine trust in agentic transactions and create a major obstacle to the widespread adoption of agentic commerce.

Timing of regulatory intervention

Closely related to the question of regulatory priorities is the question of the right timing. This presents a genuine dilemma: intervening too early risks regulating a technology whose commercial and technical trajectory is still highly uncertain, while intervening too late risks allowing potentially harmful practices and market structures to become entrenched before rules can catch up.

The considerations assembled in this brief suggest that the 'too late' risk currently outweighs the 'too early' risk. Agentic shopping features are quickly moving from experimental to commercially deployed at scale. The technical infrastructure for agentic commerce — ranging from standardised communication protocols to agentic payment solutions — is rapidly being developed, while EU consumer policy has, to date, barely addressed agentic commerce as a distinct issue.

This asymmetry mirrors the pattern already seen with the platform economy, where regulatory intervention through the DSA and DMA arrived only at a point when regulatory intervention became rather difficult and costly. Waiting for an 'iPhone moment' of agentic commerce would repeat this pattern. Policy makers should instead treat the current period of technological flux as the more opportune moment to shape the rules of the agentic economy, precisely because market structures are not yet locked in.



Risks of regulatory inaction

If the DFA proceeds without addressing agentic AI, several types of risk could materialise. One possible scenario is that the lack of ‘agent readiness’ in European consumer law and the lack of legal certainty for agentic commerce could prove to be barriers to innovation. Due to the unclear legal situation, providers might refrain from offering agentic commerce solutions in the EU. This would mean that European consumers would lose out on the benefits of agentic commerce, and no agentic commerce ecosystem would emerge in the EU. It is also conceivable that the unclear regulatory framework causes consumers to hesitate in using agentic commerce solutions because they lack trust in the technology.

Another possible scenario is that some providers ignore the lack of ‘agent readiness’ in consumer law and decide to ‘move fast and break things,’ offering AI agents on the European market. This would also not be a desirable outcome for consumers. In this scenario, regulatory inaction leaves open exactly the gaps discussed in this brief: unclear consent mechanisms, absent protections against agent manipulation, and design-based rules that do not easily translate to the protocol layer. As a consequence, consumers who delegate purchasing decisions to agents would enjoy a lower level of protection than those who transact directly, despite EU consumer law's stated aim of ensuring a high level of protection regardless of the channel used.

An opportunity to rethink the current approaches of consumer law

Finally, the rise of agentic commerce could also be seen as an opportunity to revisit foundational assumptions of consumer law. As already mentioned, the role of mandatory disclosures could be reconsidered in the agentic economy. For example, AI agents could function as enablers of personalised disclosures.²⁵ More specifically, AI agents may serve as techno-legal intermediaries capable of tailoring legally relevant information to the needs and preferences of individual users. This could give a new practical twist to the theoretical debate on Personalized Law.²⁶

Similarly, the shift of regulatory focus away from the ‘magical moment’ of contract conclusion toward the earlier moment of agent configuration and deployment invites a broader rethinking of when, exactly, consumer protection should intervene in a transaction. In this perspective, the rise of agentic commerce could be used as a catalyst for modernising the conceptual architecture of EU consumer law as a whole. In this regard, the recently published ELI Guiding Principles and Model Rules on Digital Assistants for Consumer Contracts (ELI DACC Model Rules), which combines ex-ante design requirements for AI agents with ex-post contractual rights, could serve as a source of inspiration.²⁷

More generally, the rise of AI agents could also provide an opportunity to reconsider which concepts

²⁵ See Christoph Busch, AI Agents as Enablers of Personalized Law: Towards Agentic Disclosures? Oxford Business Law Blog, 16 March 2026, <https://blogs.law.ox.ac.uk/oblb/blog-post/2026/03/ai-agents-enablers-personalized-law-towards-agentic-disclosures>; see, generally, Ariel Porat and Lior Strahilevitz, Personalizing Default Rules and Disclosure with Big Data, 112 Mich. L. Rev. 1417, 1417–78 (2014).

²⁶ See, generally, Omri Ben Shohar and Ariel Porat, Personalized Law: Different Rules for Different People, Oxford 2021; see also Christoph Busch and Alberto De Franceschi (eds.) Algorithmic Regulation and Personalized Law, Munich 2021.

²⁷ European Law Institute, ELI Guiding Principles and Model Rules for on Digital Assistants for Consumer Contracts, May 2025; see also Christoph Busch, Marie Jull Sørensen, Teresa Rodríguez de las Heras Ballell and



will become obsolete and which will endure the ‘agentic turn’. Some of the highly prescriptive, human-facing design rules — precise requirements about button text, font size, or visual prominence — are of limited value once transactions are increasingly concluded via API rather than through a graphical interface. By contrast, the underlying substantive standards these design rules were meant to implement appear likely to endure and indeed to gain in importance: the requirement that consumers give genuine, informed consent to a payment obligation or the prohibition on materially distorting a consumer's (or an agent's) economic behaviour remain just as relevant when a machine is executing the transaction as when a human is.

Christian Twigg-Flesner, Towards the Next Generation of Consumer Law: the ELI Guiding Principles and Model Rules on Digital Assistants for Consumer Contracts, 14 Journal of European Consumer and Market Law 89 (2025).



About CERRE

Providing high quality studies and dissemination activities, the Centre on Regulation in Europe (CERRE) is a not-for-profit think tank. It promotes robust and consistent regulation in Europe's network, digital industry, and service sectors. CERRE's members are regulatory authorities and companies operating in these sectors, as well as universities.

CERRE's added value is based on:

- its original, multidisciplinary and cross-sector approach covering a variety of markets, e.g., energy, mobility, sustainability, tech, media, telecom, etc.;
- the widely acknowledged academic credentials and policy experience of its research team and associated staff members;
- its scientific independence and impartiality; and
- the direct relevance and timeliness of its contributions to the policy and regulatory development process impacting network industry players and the markets for their goods and services.

CERRE's activities include contributions to the development of norms, standards, and policy recommendations related to the regulation of service providers, to the specification of market rules and to improvements in the management of infrastructure in a changing political, economic, technological, and social environment. CERRE's work also aims to clarify the respective roles of market operators, governments, and regulatory authorities, as well as contribute to the enhancement of those organisations' expertise in addressing regulatory issues of relevance to their activities.



About the Author



Christoph Busch is Professor of Law and Director of the European Legal Studies Institute at the University of Osnabrück, Germany. He is a Fellow and former Council Member of the European Law Institute (ELI) and an Affiliated Fellow at the Information Society Project at Yale University. In addition, he is a member of the European Commission's Consumer Policy Advisory Group and the Expert Group on AI Contracting. His research focuses on consumer law, platform governance, and AI regulation.

cerre



Avenue Louise 475 (box 10)
1050 Brussels, Belgium
+32 2 230 83 60
info@cerre.eu
www.cerre.eu

 Centre on Regulation in Europe (CERRE)

 CERRE Think Tank

 CERRE Think Tank