September 02, 2026

*Submitted electronically: drcf@ofcom.org.uk*

**RE:** ***Comments of the Association for Competitive Technology in Response to the DRCF Call for Input - Consumer interest and AI (Phase 2: Consumers, regulators, policymakers & industry tools)***

In response to the Digital Regulation Cooperation Forum's (DRCF) call for input titled *'Consumer interest and AI',* published on 3 June 2026,[1] the Association for Competitive Technology (ACT) respectfully submits these comments. These comments respond to the second phase of the DRCF's inquiry, concerning the tools and frameworks available to consumers, industry, regulators, and policymakers to manage risks associated with generative and agentic artificial intelligence (AI) while avoiding undue restrictions on AI opportunities and growth.

ACT welcomes the opportunity to comment and shares the DRCF's objective of ensuring that consumers can benefit from AI while material risks are appropriately mitigated. We offer these comments from the perspective of the small business and startup technology developer community, whose members are developing and deploying AI across consumer and enterprise contexts and are often among those most exposed to unintended consequences when regulatory obligations are unduly broad or difficult to operationalise. ACT believes that the DRCF's approach should adhere to the following principles:

- Remain risk-based, evidence-led, and focused on identifiable harms;
- Build on existing law and recognised international standards before creating new AI-specific obligations;
- Allocate responsibility to the actors best positioned to understand and mitigate a particular risk; and
- Ensure that compliance tools are practical for smaller developers who do not have dedicated legal or compliance teams.

These principles can support meaningful consumer protection while preserving the innovations that allow consumers to realise AI's benefits.

## I. Statement of Interest

ACT is a not-for-profit global policy trade association for the small business technology developer community. Our members are entrepreneurs, innovators, and independent developers within the global app ecosystem who build products and services across every industry, including in the United Kingdom. We work with our members and policymakers to promote a policy environment that rewards and inspires innovation while providing resources that help developers raise capital,

[1] Digital Regulation Cooperation Forum (DRCF), Call for Input - Consumer interest and AI (3 June 2026), https://www.drcf.org.uk/news-and-events/news/call-for-input-consumer-interest-and-ai.

create jobs, and continue to build technology. The value of the global ecosystem ACT represents is approximately £74 billion and serves as a key driver of the internet of things (IoT) revolution.[2]

AI is increasingly embedded throughout the digital economy and in the applications and services consumers already rely on every day. From shopping and banking, to healthcare, accessibility, cybersecurity, productivity, and communications, AI can enable faster and better-informed decision making, streamline routine processes, improve the detection of financial and identity theft, and make digital services more accessible and responsive to consumers' needs. Many of these benefits will be delivered not through standalone AI products, but through AI capabilities integrated into existing apps and connected devices. This innovation builds on an already significant economic ecosystem; for example, in 2021 the UK app economy generated nearly £74 billion in revenue throughout all sectors and supported more than 400,000 jobs across the UK in 2021.[3]

The small business technology developers ACT represents build the applications and services through which consumers and businesses across the United Kingdom increasingly benefit from AI. For these developers, consumer trust is essential to the adoption of AI-enabled products and services, while clear and predictable rules are essential to continued innovation and investment.

## II. The DRCF Should Preserve a Risk-Based, Context-Specific Approach Grounded in Existing Law (Questions 19, 24, and 25)

AI is not a standalone sector but a general-purpose technology that cuts across regulated domains. Regulatory approaches should therefore prioritise coordination and clarity over new, centralised regimes. Early choices in AI policy will shape how the technology is deployed across the economy. Rules that treat AI as a single category or assume a static market tend to slow deployment and raise compliance costs where much of the societally beneficial innovation is happening, especially for small businesses. Good AI policy uses existing legal frameworks where they fit and addresses genuine gaps with targeted rules.

That approach is also consistent with the United Kingdom's longstanding pro-innovation framework. The Government has emphasised that regulation should focus on the use of AI rather than the technology itself, and that context matters because the same technology can produce different risks in different applications. Existing regulators are best placed to conduct detailed risk analysis and enforcement within their areas of expertise. The DRCF is therefore well positioned to improve coordination among the Competition and Markets Authority (CMA), Information Commissioner's Office (ICO), Office of Communications (Ofcom), and Financial Conduct Authority (FCA) without creating a parallel cross-sector AI compliance regime.[4] A small developer adding a generative feature to an existing offering may face online safety duties overseen by Ofcom, data

[2] Association for Competitive Technology (ACT), *The State of the App Economy*, https://actonline.org/wp-content/uploads/2026/09/US-App-Economy-Fast-Facts.pdf.

[3] Deloitte, The App Economy in Europe (Aug. 2022) (commissioned by the Association for Competitive Technology (ACT)), https://actonline.org/wp-content/uploads/220912_ACT-App-EU-Report.pdf; techUK, UK Tech SMEs: A Global Force to Be Reckoned With (2023), https://www.techuk.org/resource/uk-tech-smes-a-global-force-to-be-reckoned-with-techukdigitaltrade.html.

[4] Department for Science, Innovation & Technology, A pro-innovation approach to AI regulation, paras. 45-50 (29 Mar. 2023), https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper.

protection obligations overseen by the ICO, and consumer protection obligations overseen by the CMA, all arising from a single product decision. Straightforward guidance from UK government showing how these obligations fit together would be particularly beneficial to the ACT community.

The CMA's recent work on agentic AI illustrates why such a coordinated approach is needed. UK consumer protection law applies whether decisions are made by people or AI, and existing law focuses on the effects on consumer decision-making rather than the novelty of the underlying technology. Businesses may not mislead or exert undue pressure on consumers merely because an algorithm or agent mediates the interaction. The appropriate policy question is therefore not whether AI is involved, but whether a particular design or use creates a meaningful risk of harm that existing law does not adequately address.[5]

ACT accordingly supports outcomes-focused approaches where they clarify how existing statutory obligations apply and give regulated firms flexibility to achieve appropriate consumer outcomes. However, a generalised, cross-sector AI analogue to the FCA Consumer Duty should not become an open-ended new duty layered on top of existing regulatory remits. Regulatory scrutiny and obligations should scale with demonstrated risk and harm, not apply uniformly across AI use cases. Where the DRCF identifies a genuine gap, any new measure should be targeted, proportionate, clearly within the relevant regulator's remit, and accompanied by practical guidance that smaller firms can implement.

## III. Industry Tools, Standards, Guidance, and Sandboxes Should Translate Risk Principles into Practical Compliance (Questions 17, 19-21, and 24)

ACT recommends that the DRCF focus on practical tools that can be scaled to the relevant use case and the actor's role in the AI value chain. These tools should include, as appropriate:

- clearly defining intended and reasonably expected uses, limitations, and relevant risks;
- risk assessments and testing that are proportionate to the likely harms and consequences of the use case;
- documentation sufficient to support safe downstream deployment without compelling disclosure of proprietary information or trade secrets;
- monitoring of real-world performance, including errors, bias, complaints, security incidents, and unintended outcomes;
- feedback, escalation, and human-oversight mechanisms appropriate to the consequence of the decision or action; and
- for agentic systems, clear limits on permissions, credentials, tools, integrations, and the actions an agent may execute without additional user confirmation.

These recommendations draw directly from ACT's AI Roles and Interdependencies Framework. Policy frameworks should encourage design of AI systems that are informed by real-world workflows, human-centred design and usability principles, and end-user needs. Those developing or offering AI systems intended for consumer use should provide truthful and easy-to-understand representations regarding intended use and risks that would be reasonably understood by intended and expected users. AI solution developers should aim to provide interfaces that allow end users to

[5] Competition and Markets Authority (CMA), Agentic AI and consumers (9 Mar. 2026), https://www.gov.uk/government/publications/agentic-ai-and-consumers/agentic-ai-and-consumers.

safely and effectively act upon system outputs, including through explanations, feedback mechanisms, and human-oversight options where appropriate.[6]

For agentic AI, an underlying AI model can present materially different risks depending on how it is deployed. The key governance question is how much real-world authority an agent is granted, including what actions it can execute, whether it operates in a bounded or unbounded environment, and what credentials, tools, and system integrations it can access. Security expectations and oversight should scale in proportion to that authority, with the strongest controls reserved for agents deployed with elevated privileges, consequential capabilities, or operation in high-impact environments.

The CMA's illustrative personal shopping and finance agent usefully demonstrates this principle in a consumer context: autonomy can be bounded by user-defined constraints; users can be alerted before high-impact or irreversible decisions; high-risk actions can require confirmation; and providers can monitor performance, feedback, and complaints with human oversight, clear escalation, audit logs, and risk assessments. These are more useful safeguards than generic disclosure or consent prompts because they intervene at the point at which real-world consequences can occur.[7]

Recognised international standards can help translate these principles into repeatable practices. The DRCF should encourage use of the National Institute of Standards and Technology (NIST) Artificial Intelligence Risk Management Framework (AI RMF), ISO/IEC 42001 on AI management systems, and ISO/IEC 23894 on AI risk management, among other recognised standards and references. These frameworks are valuable because they are adaptable across organisations and use cases and can reduce the need for businesses to build separate compliance programmes for each market or regulator.[8]

This interoperability is especially important for small and medium-sized enterprises (SMEs). Unlike larger firms, small developers cannot easily navigate separate compliance programmes for every market they seek to enter or operate in. Where UK guidance maps cleanly to recognised international references, small businesses can build once and serve customers worldwide. Government should therefore continue to support private-sector and multistakeholder standards development and avoid converting voluntary standards wholesale into rigid regulatory checklists that may become outdated as technology evolves.

DRCF also inquires about what modes of guidance and other regulatory tools are most helpful for achieving compliance. In ACT's experience, the UK government can and should prioritise direct, low-cost access to a regulator who will answer a specific question about a specific product. A

[6] Association for Competitive Technology (ACT), AI Roles and Interdependencies Framework (May 2024), https://actonline.org/wp-content/uploads/2026/09/ACT-AI-Roles-Interdependencies-Framework-final-text-May-2024-UK-English.pdf.

[7] CMA, Agentic AI and consumers, Illustrative case study: the personal shopping and finance agent (9 Mar. 2026), https://www.gov.uk/government/publications/agentic-ai-and-consumers/agentic-ai-and-consumers.

[8] National Institute of Standards and Technology (NIST), Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (26 Jan. 2023), https://doi.org/10.6028/NIST.AI.100-1; ISO/IEC 42001:2023, Information technology - Artificial intelligence - Management system, https://www.iso.org/standard/42001; ISO/IEC 23894:2023, Information technology - Artificial intelligence - Guidance on risk management, https://www.iso.org/standard/77304.html.

leading example of such an effort is the DRCF's own AI and Digital Hub, which offered free and informal cross-regulatory advice from each of its member regulators, though that service ran as a one-year pilot and is now closed.[9] ACT encourages the DRCF to reinstate it on a permanent basis and to continue publishing anonymised case studies so that firms that never submit a query still benefit from the answers given to those that did.

ACT also notes that sector sandboxes serve a similar function and that DRCF member regulators already operate several, including the Financial Conduct Authority's AI Lab with its AI Live Testing cohorts and Supercharged Sandbox, and the Information Commissioner's Office Regulatory Sandbox.[10] ACT encourages the DRCF to consider eligibility criteria that is inclusive of early stage and SME developers.

## IV. Consumer Choice, Opt-Outs, Consent, and Other Levers Should Be Tied to Material Risk and Genuine Control (Questions 18, 22, and 23)

ACT does not recommend a general right to opt out of the use of AI merely because AI is involved in delivering a product or service. Such a rule would be infeasible in operation in light of the fact that AI is routinely integrated into ordinary business and software processes (in other words, it is not bolted on as a separable feature that can be switched off on request) and would be inconsistent with a scaled approach to risk management. AI is already used for low-risk or beneficial functions such as fraud detection, cybersecurity, accessibility, search, recommendations, customer-service triage, and back-office efficiency. Requiring a separate AI opt-out for every such use could reduce product quality and security while creating little meaningful consumer protection.

Instead, policymakers should identify the consumer interest at stake and tailor the lever accordingly. They should preserve space for low-risk, informational uses that expand access and avoid collapsing those uses into the same category as high-risk decision-making.

The DRCF also asks what the legal basis is for the use of consumers' data in the event there is no general right to opt out. The Data (Use and Access) Act 2025, in effect since 5 February 2026, removed the general prohibition on solely automated decision-making (other than where special category data are involved) so that a controller may rely on legitimate interests rather than being confined to consent, contractual necessity, or legal authorisation.[11] This flexibility is important because a developer running, for example, transaction monitoring to detect payment fraud or automated screening to detect account takeover, cannot realistically obtain meaningful consent from the person whose conduct is being screened (and they may indeed be enabled to withhold it). Under current law, a controller must tell the individual that a decision has been taken by

[9] DRCF, AI and Digital Hub, https://www.drcf.org.uk/ai-and-digital-hub (describing the Hub as a one-year pilot, now closed for queries, that offered free and informal advice from the CMA, ICO, Ofcom, and FCA to innovators whose questions spanned at least two of those remits).
[10] Financial Conduct Authority, AI Lab, https://www.fca.org.uk/firms/innovation/ai-lab (comprising AI Live Testing, the Supercharged Sandbox, AI Spotlight, and AI Sprint); Information Commissioner's Office, Regulatory Sandbox, https://ico.org.uk/for-organisations/advice-and-services/regulatory-sandbox/.
[11] Data (Use and Access) Act 2025, pt. 5, ch. 1, https://www.legislation.gov.uk/ukpga/2025/18/part/5/chapter/1/crossheading/automated-decisionmaking/enacted.

automated means, allow the individual to make representations, provide access to meaningful human intervention, and allow the decision to be contested,[12] giving a consumer recourse at the point where an automated decision actually affects them. ACT supports modernised frameworks that assure that an individual's data is properly protected while allowing the flow of information and responsible evolution of AI, and a balanced framework should avoid undue barriers to data processing and collection while imposing reasonable data minimisation, consent, and consumer rights frameworks. A blanket opt-out taken at the outset does not provide that control to the individual.

The DRCF should therefore avoid 'cookies-like' AI prompts that generate repeated interruptions without changing the consumer's substantive rights or the risk of the underlying activity.[13] Further, cooling-off periods are worth distinguishing from AI-specific mechanisms because consumers already hold cancellation rights and those rights attach to the transaction rather than to the technology that produced it. A separate AI-specific cooling-off regime layered on top of existing cancellation rights would create two overlapping clocks for a single transaction and would impose a significant implementation burden on SMEs without improving the consumer's position.

## V. Accountability Should Follow Knowledge and the Ability to Mitigate Risk (Question 26)

ACT urges the DRCF to align with the concept of shared responsibility across the AI value chain. The AI ecosystem includes infrastructure providers, model developers, application developers, deployers, and end users. Not all actors have visibility into or control over how systems are used in practice. Assigning liability broadly at one point in the chain, without regard to control or context, creates incentives to limit functionality rather than improve safety.

Obligations should be allocated across the AI value chain based on who can mitigate risk. Specifically, those in the value chain with the ability to minimise risks based on their knowledge and ability to mitigate should have appropriate incentives to do so. Foundation-model and upstream developers can provide information necessary for downstream actors to understand intended uses, relevant limitations, and reasonably foreseeable risks within their control. Application developers and deploying organisations can assess whether the system is fit for the intended context, configure permissions appropriately, monitor performance, and maintain suitable human oversight. End users can follow reasonable instructions and use available controls where their own conduct affects the risk.

This shared-responsibility model is compatible with ordinary consumer-law accountability. A business that uses an AI agent to engage with its customers remains responsible for complying with consumer law in that interaction. But that does not justify making the consumer-facing business automatically responsible for defects or risks that arise from upstream conduct it could not reasonably observe or mitigate. The DRCF should preserve the distinction between

[12] *Id*.
[13] Information Commissioner's Office (ICO), When is consent appropriate? (retrieved Aug. 2026), https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/consent/when-is-consent-appropriate/; ICO, A guide to lawful basis (updated 2 Apr. 2026), https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/a-guide-to-lawful-basis/.

responsibility for one's own conduct and strict, undifferentiated liability for the entire AI supply chain.

## VI. 'Tolerable Risk' Should Be Contextual, Evidence-Based, and Capable of Changing as Systems and Safeguards Evolve (Questions 27 and 28)

Tolerable risks can exist in consumer-facing AI, just as they exist in other technologies and markets. The relevant task is not to eliminate all risk, which is neither realistic nor necessarily welfare-enhancing, but to determine whether residual risk is reasonable in light of the likely severity and likelihood of harm, the benefits of the use case, available mitigations, and the consumer's ability to understand, contest, or reverse consequential outcomes. Risk tolerance should therefore be contextual rather than expressed through a single numerical threshold applicable to all AI.[14]

The NIST AI Artificial Intelligence Risk Management Framework RMF is helpful on this point. It does not prescribe a universal risk tolerance and recognises that acceptable levels of risk are highly contextual and application- and use-case-specific, and may change over time as systems, policies, and norms evolve. The OECD's 2026 Due Diligence Guidance for Responsible AI likewise evaluates severity through factors such as scale, scope, and irremediable character, considered alongside likelihood. These approaches are more suitable than a static threshold attached to a technology.

Low-risk uses may require only baseline governance and ordinary legal compliance. Uses with elevated consequences may justify stronger testing, monitoring, documentation, and human oversight. Systems authorised to execute high-impact or irreversible actions may justify mandatory confirmation, more robust evaluation, and rapid escalation and remediation. Any classification should be transparent, evidence-based, limited to intended and reasonably expected uses, and revisited at regular intervals as capabilities and safeguards mature.

Regarding the feasibility of metrics, thresholds, or proxy indicators, ACT discourages using a single arbitrary threshold and instead recommends a composite approach drawing on both design-side and operational-side measures. Design-side measures capture how much reach an agent is given and how readily its actions can be constrained or undone. Operational-side measures capture what deployment actually shows about how often the system errs and how quickly those errors are caught and corrected. Taken together, they should give an accurate picture of risk than any one figure could.

With respect to DRCF's flagging agentic collusion as a risk,[15] we urge DRCF to recognise that algorithmic tools can be procompetitive by helping firms, including small businesses, respond

[14] NIST, Artificial Intelligence Risk Management Framework (AI RMF 1.0), sec. 1.2.2 Risk Tolerance, NIST AI 100-1 (26 Jan. 2023), https://doi.org/10.6028/NIST.AI.100-1; OECD, OECD Due Diligence Guidance for Responsible AI (19 Feb. 2026), https://doi.org/10.1787/41671712-en.

[15] DRCF, Call for Input - Consumer interest and AI, question 4 (3 June 2026) (listing agentic collusion, with the possibility of price spikes, among the risks put to consumers in the first phase of this call for input).

more quickly to supply, demand, inventory, and consumer behaviour, reduce overhead, and improve efficiency. The mere use of an algorithm should not be treated as evidence of anticompetitive conduct. The CMA has itself recognised that the risks it describes do not imply that the use of algorithms or AI in pricing is inherently problematic, and that the assessment turns on context, design, and governance.[16] Concerns may arise where shared tools or pooled competitively sensitive data facilitate an agreement or conduct that harms competition, but those concerns should be addressed through fact-specific enforcement rather than broad presumptions against algorithmic tools.

## VII. Conclusion

The United Kingdom has an opportunity to provide a practical path forward that protects consumers while enabling responsible AI deployment and growth. Small businesses need clear rules they can understand and apply without bespoke legal advice at every step. A framework with too little clarity chills investment, and a framework built around broad process mandates favours the largest firms. The DRCF can support both trust and innovation by keeping consumer protection focused on outcomes, meaningful risks, practical safeguards, and the actors best positioned to address them.

ACT appreciates the DRCF's consideration of the above views and stands ready to engage further with the DRCF and its member regulators as this work develops.

Sincerely,

Mike Sax
Founder and Chairperson
Association for Competitive Technology (ACT)

Stephen Tulip
UK Country Manager
Association for Competitive Technology (ACT)

[16] Karen Croxson and Juliette Enser, AI and collusion: frontiers, opportunities and challenges, CMA Competition and Markets blog (4 Mar. 2026), https://competitionandmarkets.blog.gov.uk/2026/03/04/ai-and-collusion-frontiers-opportunities-and-challenges.