Response ID ANON-7M4M-3R5P-F

Submitted to Agentic AI Call For Evidence
Submitted on 2026-11-16 13:23:41

Section 1: About you and your organisation

1 Are you responding to this consultation on behalf of an organisation?

No

2 What is the name of your organsiation?

What is the name of your organisation? :

3 If you are an organsiation or responding on behalf of one, is your organisation:

Not Answered

4 Which of these best describes the main industry your organisation operates in?

Other (please specify)::

Section 2: Data Security

5 In an agentic AI system, what session isolation controls ensuring data within a session is only accessible to specific parties (eg the user who started the session) are used?

Please answer below:

6 How do organisations evaluate whether to grant an agent access only to internal knowledge systems (eg document repositories, policy libraries, intranet sites) or allow it access to third-party knowledge systems or databases?

Please answer below:

7 Do guardrails and permission controls change dynamically in an Agentic system, if so how?

Please answer below:

8 What are examples of structured logging measures that record AI agent identity alongside agent actions (eg data retrieval, database entry modification), outcome tracking and tamper-evident records?

Please answer below:

9 What system setups or functionalities built in by an agentic AI system provider are effective at supporting data security, and where do deployers feel there are challenges to ensuring effective data security?

Please answer below:

10 Do you have any other comments about data security?

Please answer below:

King Yew Choo's work on payments, AI agents and financial regulation has been published by the House of Lords, Federal Reserve and Financial Stability Board. He has authored payments use cases for the United Nations' ITU focus group on trust and identity for agentic AI. He serves on the British Standards Institution's financial services committee. Before founding True Primary, which builds payment infrastructure for agentic commerce, he traded emerging-market credit at Citigroup.

He responds in a personal capacity, from London. The House of Lords Financial Services Regulation Committee published his written evidence on testing consumer payments made by AI agents in regulatory sandboxes (RSB0009, October 2026). Of his two payments use cases for the ITU-T Focus Group on Trust and Identity for Humans and Agentic AI, one is public. The British Standards Institution committee he serves on is IST/12, which leads the UK's input to ISO/TC 68. His individual IETF Internet-Draft on agentic payments (work in progress) models the provider's side: a payment or entitlement gate before fulfilment, for agents retrying after uncertain outcomes. His SSRN working paper sets out a related gate's architecture and a formal admission model. These answers take the same failure, a retry after a lost receipt, to the payer's side, where the payer's bank decides a dispute on evidence other firms hold.

Interest: he founded True Primary, which is focused on the agent-initiated purchase of primary research and has a commercial interest in how agentic payments are regulated. In tests of its payment gate build, a forged key, a replayed proof and a request over a $5 daily cap were each refused.

Links: Lords written evidence (RSB0009): https://committees.parliament.uk/writtenevidence/170144/html/ · FSB: https://www.fsb.org/uploads/King-Yew-Choo-independent-response.pdf · Federal Reserve Board: https://www.federalreserve.gov/apps/proposals/comments/FR-2026-0013-01-C86 · ITU use case: https://github.com/fg-tida-wg0/use-cases/issues/35 · BSI IST/12: https://standardsdevelopment.bsigroup.com/committees/50001771 · Internet-Draft: https://datatracker.ietf.org/doc/draft-king-yew-choo-agentic-payments/ · SSRN working paper: https://doi.org/10.2139/ssrn.7543686 · Submissions and papers: https://kingyewchoo.co.uk · ORCID: https://orcid.org/0009-0002-5244-6082

Section 3: Transparency

11 How do developers or deployers of an agentic AI system explain whether they handle personal information on a person's device, in a remote environment or through a combination of both?

Please answer below:

12 How do organisations provide timely information to people when an agent: accesses a new data source; uses an external tool; or shares personal information with another agent or organisation?

Please answer below:

13 How do organisations use pre-authorisation, prompts, approval steps, live status indicators or other techniques / approaches to improve transparency for users?

Please answer below:

14 How do organisations allow end users to set limits on the websites, tools, data sources and actions available to an agent?

Please answer below:

15 How do post-task execution summaries and audit trails explain: what an agent did; what personal information it used; which tools and recipients were involved, and what happened as a result?

Please answer below:

Answers 22 and 29 set out such a trail for agent payments under the Agent Payments Protocol (AP2). The code signing for the agent logs each mandate it signs, with the agent's key and the task and permission then in force; each party receiving a mandate logs every arrival under its own reference; and the signed receipts record each outcome. The log records which personal data the agent read, from which source and under which permission; for each disclosure, its recipient, purpose and lawful basis; and, at each stage, the payment authority, the lawful basis and the human involvement. The signatures make a later change to a mandate or receipt detectable, and the logs sit outside the agent, which AP2 treats as a potential attacker, because a signature cannot show a record is missing.

Please answer below:

16 How can people challenge, correct or reverse an agent's actions?

Please answer below:

My response asks the ICO's guidance on agentic AI to cover six points:

1. Each controller should decide and write down who keeps each record of a person's instruction and of the agent's actions, for how long, and how she and anyone deciding her dispute can obtain it (answers 20 to 22).
2. The Theme 3 list of considerations should include one more, for any agent acting on a person's delegated authority: who decides how long the evidence of her instruction, and of what the agent did under it, is kept, and to whom it is released in a dispute (answer 20).
3. The guidance should say whether the data subject's own approval can count as meaningful human involvement, and on what terms (answer 27).
4. Logs of an agent's actions should be written outside the agent, by the code signing for it and by each party receiving what it signs (answers 22 and 29).
5. Developers should state the limits of the protocols their systems rely on, such as dispute retention left outside a protocol's scope (answer 24).
6. The mandate, the payment and the data should each have their own screen, so one approval is not read as covering all three (answers 29 and 30).

One assumed card case runs through the later answers. A UK consumer's AI shopping agent pays by debit card, from her current account in credit, in the human-not-present flow of the Agent Payments Protocol (AP2, v0.2, ap2-protocol.org), the specification Google contributed to the FIDO Alliance in April 2026. Her bank issues the card only and receives no mandate with the authorisation request, and a separate firm runs the wallet (both assumed; AP2 lets one entity play several roles). She approves open checkout and payment mandates for one pair of shoes under £120, bought online from a trader under a distance contract outside the excepted kinds. In AP2's terms, the open mandates hold the limits she approves. For each purchase the agent signs, with its own key, closed mandates for the exact checkout and payment, and the merchant and its payment processor return signed receipts accepting or rejecting them. The receipt for the first payment never reaches the agent, which presents the same open mandates again for a second order. This breaches AP2's rule against any subsequent presentation "without receiving a rejection receipt from the previous one", a rule whose stated purpose is "to prevent an Agent approving multiple different Checkouts using the same open Mandate". She disputes the second payment.

The evidence of her instruction, and of what the agent did, is in the mandates and receipts. AP2 lists the agent, the credential provider (her wallet), the network and the merchant's payment processor as possible sources of the payment mandate and receipt in a dispute. Its overview says the payment mandate "may be shared with the network & issuer for transaction authorization", but AP2 requires delivery of neither to the card issuer, and it leaves

dispute retention, retrieval and any sharing by agreement outside its scope.

Her right of access (UK GDPR Article 15) reaches the personal data each controller holds about her, which can include those mandates and receipts. Where a firm's case file records as fact that she approved the second order, the entry is inaccurate. She can have it rectified without undue delay (Article 16) and its use restricted while the firm verifies it (Article 18(1)(a)), and the firm passes the correction, and any restriction, to each recipient of the entry, such as the wallet, merchant or processor, unless this proves impossible or involves disproportionate effort (Article 19). Guidance should have the firm still keep the signed record of what the agent did, with the correction linked to it, because whether the payment was authorised is settled in the payment dispute.

That dispute runs with her bank. If she denies authorising the second payment, the bank must prove, among other things, that it was authenticated and accurately recorded, and the recorded use of her card is not in itself necessarily enough to prove she authorised it (regulation 75(1) and (3) of the Payment Services Regulations 2017). The bank must refund an unauthorised payment as soon as practicable and at the latest by the end of the business day after it becomes aware of it (regulation 76(1) and (2)). Only reasonable grounds to suspect her of fraud, notified as regulation 76(3) provides, lift that deadline, and the refund waits neither for other firms' records nor for her access requests. On this distance sale she bears no loss from an unauthorised payment unless she acted fraudulently (regulation 77(4)(d) and (5)). She can also cancel the second order within the cancellation period for distance contracts (regulation 29 of the Consumer Contracts Regulations 2013) and ask the agent provider to suspend her mandate.

Please answer below:

17 How can agent identification, authentication, permissions and tracing support transparency to people?

Please answer below:

18 How do organisations provide meaningful transparency when personal information moves between agents in a multi-agent system?

Please answer below:

19 Do you have any other comments about Transparency?

Please answer below:

Section 4: Accountability

20 How do organisations allocate controller, joint controller and processor accountabilities in agentic AI systems?

Please answer below:

In an agent payment, allocation is worked out operation by operation on each record of the person's instruction (keeping, verifying, reusing or releasing it). The controller of each operation is whoever determines its purposes and means (UK GDPR Article 4(1)(7)), and those parties can differ from the bank, which must refund an unauthorised payment and needs the records to investigate.

In the card case (answer 16), the roles in the Agent Payments Protocol (AP2) carry protocol or payment-law duties, which settle no one's controller status. The agent provider runs the agent, which signs the closed mandates with its key; the merchant verifies the checkout mandate; the wallet, any network and the merchant's payment processor verify the payment mandate; and the bank executes the payment and carries the refund duty (regulation 76 of the Payment Services Regulations 2017). A firm playing several AP2 roles takes each role's duties. Verifying and keeping are separate. AP2 has the agent store the mandates once a successful receipt arrives, and lists "Receiving and storing the Payment Receipt" among the credential provider's implementation considerations, a second copy the bank can ask for by agreement. AP2 sets no retention period for any of them.

The same record can carry two uses. A firm storing a mandate only on the merchant's documented instructions is the merchant's processor for that storage (Articles 4(1)(8) and 28(3)(a)). If it feeds the same record into its own fraud model, outside those instructions, it becomes a controller for that use (Articles 4(1)(7) and 28(10)), with its own lawful basis to show (Article 6(1)). Firms jointly determining the purposes and means of keeping a mandate are joint controllers (Article 26(1)), and the person can exercise her rights against each of them (Article 26(3)).

Theme 3 lists considerations which can influence the purposes and means of processing, such as who decides the agent's goals and what permissions it has. One more belongs on that list for any agent acting on a person's delegated authority: who decides how long the evidence of her instruction, and of what the agent did under it, is kept, and to whom it is released in a dispute. Payments are one example; an agent giving her landlord notice to end her tenancy is another. If she later says she never asked for it, the record of her instruction is the evidence, so how long it is kept, and who may see it, should be settled before any dispute.

21 How do organsiations show their allocation of accountability in practice and in their contracts?

Please answer below:

An organisation can show its allocation in practice by testing the dispute path end to end before launch and recording a pass or fail. Its contracts show the allocation in terms setting out, for each mandate and receipt, who keeps it, for how long, to whom it is released, on what request and how fast, and who decides the purpose of each use. Those terms sit in the joint controllers' arrangement or the processor terms (UK GDPR Articles 26 and 28(3)) or, between independent controllers, in evidence-sharing terms.

In the card case, the bank's refund of the second debit can fall due before other firms' records arrive. The records still matter afterwards: they can support a later reversal, and recovery from "another payment service provider or an intermediary" to which the bank's liability "is attributable" (regulation

95 of the Payment Services Regulations 2017). The test therefore records when the bank first held a verifiable mandate record, and which records were still outstanding when the refund fell due.

The reconstructed dispute starts from the bank's own distinct reference for each debit, because merchant, amount and time can be the same for both. The Agent Payments Protocol (AP2) defines no receipt field for that reference: its payment_id is defined only as "A unique identifier for the payment", and the PSP and network confirmation identifiers are optional. The first link is therefore a processing record tying the bank's debit reference to a card transaction identifier that the merchant's payment processor (which signs the receipt) or, where it receives the receipt, the network can match to the payment receipt for that debit. The contract names who supplies and attests each link. Signature and hash checks start at the receipt, whose reference is the hash of the closed payment mandate, and run back through the signed checkout to the open mandate the consumer approved. A link which cannot be made is recorded as unresolved, with the party responsible for closing it. The test should include a swap: give two same-merchant, same-amount debits each other's receipts, so each points at the wrong checkout under the same open mandate. Every signature and hash still verifies, so only the processing record can catch the swap, and without it both matches stay unresolved and the test fails.

22 What mechanisms should organisations use for logging, tracing and monitoring AI agent actions (including data access) to ensure controllers retain control and accountability for the data processing that serves the purposes they set?

Please answer below:

Each payment needs a dispute-grade record, keyed on the bank's debit reference. It links to the receipt and mandates of the Agent Payments Protocol (AP2) through a processing record tying the debit reference to a card transaction identifier matched to the payment receipt (answer 21), or is marked unresolved, and it keeps four things apart:

1. The person's instruction: the open mandates she approved and their constraints, such as the item and the price cap.
2. The agent's action: the closed mandates signed with the agent's key, each presentation of the open and closed mandates to the wallet or merchant, and whether a rejection receipt preceded it.
3. The payment authority: the consent, under her bank's agreed procedure, covering the debit (regulation 67 of the Payment Services Regulations 2017).
4. Data access and disclosure: which personal data the agent read, from which source and under which permission, and each disclosure's recipient, purpose and lawful basis.

AP2 v0.2 signs each mandate and receipt, so a later change to a kept record's signed content is detectable. A signature cannot show a record is missing, though, and AP2 treats the agent itself as a potential attacker, so item 2 is written outside the agent, at two points. The first is the signer, deterministic code which checks each closed mandate before signing it with the agent key, and which AP2 gives as one way to implement the agent key. It records each mandate it signs, with the task and permission in force, and holds an overlapping one until a rejection receipt for the earlier one arrives. The second is each recipient. AP2 has the agent provide the open and closed mandates to the verifying parties, and re-sending the same signed mandates needs no new signature, so each recipient records every arrival under its own reference, identical signed bytes included: the wallet (the credential provider holding her card) for payment mandates, the merchant for checkout mandates. A recipient rejects an overlapping purchase where the earlier one succeeded and holds it while the earlier outcome is unresolved (AP2 lets the credential provider reject overlapping mandates); a named owner reconciles each hold.

Retention is set per controller and record, with dispute evidence named as the purpose when the record is made (UK GDPR Article 5(1)(b)). Each controller keeps the signed mandates exactly as it sent or received them (AP2's "compact serialization"), with the parts disclosed to it, so their hashes can be recomputed, and nothing further without a purpose of its own (Article 5(1)(c)). The period follows that purpose, the 13-month window in regulation 74(1) and any open dispute (Article 5(1)(e)); the bank's record of the complaint itself runs three years from receipt (FCA Handbook, DISP 1.9.1R). Access is limited to dispute handling, and records go when the need ends. Regulation 74(2) can keep a claim open past 13 months with no set period, so the schedule carries a review point.

23 What information do developers of agentic AI systems provide to organisations deploying them about how the systems work, how they process personal information and what risks could emerge through use, to help them demonstrate their data protection compliance?

Please answer below:

24 What information should developers of agentic AI systems provide to set out clear controller and processor accountabilities?

Please answer below:

Developers should set out, for each record the system creates about the person's instruction and the agent's actions, which party keeps it, for what purpose and for how long, who can retrieve it and by what key, and which records cannot be reconstructed once deleted. For access controls, they should say which the developer fixes and which the deployer can change.

They should also state the limits of the protocols the system relies on. The Agent Payments Protocol (AP2) v0.2 puts "dispute resolution, retention, and retrieval requirements" "outside the scope of this specification", and calls delegation of mandates from one agent to another "outside the scope of the current specification". A deployer carries the retention gap until a named party fills it and the deployer can show it. The delegation gap binds any deployment delegating mandates between agents, whether the deployer built, bought or switched on that step, and that deployer names who sets and controls the handoff.

25 Do you have any other comments about Accountability?

Please answer below:

Both checkouts in the card case trace to the one instruction the consumer gave, so one question tests a deployer's allocation: which party can show the bank which checkout produced each debit, and on what terms does it release that record?

The card case is an assumed worked example; every source is cited as it stood on 15 November 2026.

Section 5: Automated decision making (ADM)

26 How do organisations identify whether ADM provisions apply?

Please answer below:

The ICO's draft guidance on automated decision-making (updated 31 March 2026) says the provisions apply only where three factors are present, to be considered "in tandem": a system making a decision about a person, a significant decision, and no meaningful human involvement. On the first, "Decisions need to involve some kind of evaluation or analysis of personal information"; where a human sets a simple rule in advance and the system only applies it, the system makes no separate decision of its own. The draft's example is a business deciding which payment cards it accepts, with the system applying that rule.

In the card case (answer 16), her mandate fits many possible orders, so what resolved the choice matters. Where she had picked the exact pair and the agent only placed the order, nothing was left to settle and the decision is hers, by analogy with that example. Where the agent chose by evaluating her data, for instance by ranking pairs on her purchase history, that is a decision about her: the draft calls a video service's choice of content recommended from viewing habits "a decision based solely on automated processing". It is tested under Article 22A(1)(a) and (b), with the extent of profiling weighed in judging human involvement (Article 22A(2)). Where the agent evaluated nothing about her, as when it takes the cheapest pair inside her limits, it makes no decision about her. A choice which only shortlists pairs for her to pick from has no legal effect and, like that recommendation, is "less likely to have a significant effect"; one committing her to a binding order creates an obligation to pay, which can itself be a legal effect. In the human-not-present flow of the Agent Payments Protocol (AP2) the agent assembles the checkout itself, so her cap shows the limit she set and leaves open which pair she would have chosen. The retry is a different matter: it follows a missing receipt and evaluates nothing about her, so it is no decision about her, and her protection against it lies in the payment dispute.

The bank's later decisions take the same test, each on its own: the hold on the refund under regulation 76(3) of the Payment Services Regulations 2017, and the reversal after investigation (FCA, Payment Services and Electronic Money: Our Approach, paragraph 8.210). The first question is what personal data each evaluated. A fraud score on her behaviour is profiling (UK GDPR Article 4(1)(4)), which Article 22A(2) requires the bank to weigh before relying on human involvement. The second is the effect. The ICO's draft defines a legal effect as one that "affects a person's legal status or their legal rights"; a reversal bears on her statutory right to the refund, so it can have one. The draft's automated account freeze "can be a significant decision because it impacts that person's financial circumstances"; a hold under regulation 76(3) lifts the bank's statutory deadline for her refund, so it bears on her legal rights too, and its financial effect turns on the sum and how long she waits. The third is whether a reviewer with authority to change the decision weighed the score before it took effect, or saw it only afterwards. A bank unable to confidently separate the people facing a significant effect from the rest applies the safeguards to all of them, as the draft advises.

Theme 4's complete-workflow assessment therefore names, per step, who decided, on what personal data, with what effect, and the human involvement present. Where a significant decision is solely automated, Article 22B restricts it if it rests even partly on special category data and bars it if the processing relies even partly on recognised legitimate interests (Article 6(1)(ea)), so a bank resting its fraud scoring on the crime condition (Annex 1, paragraph 5) cannot take a significant hold or reversal by automation alone; where permitted, Article 22C requires safeguards.

27 How does an organisation ensure, where they have judged that the ADM provisions do not apply, that human involvement is meaningful in agentic AI systems?

Please answer below:

It depends on who made the choice: the consumer herself, or the agent evaluating her data. Where she picked the exact pair of shoes in the card case, the evidence is the instruction identifying it and the record of the agent placing it; a log showing each step stayed inside her limits proves conformance and leaves open who chose. Where the agent chose by evaluating her data, involvement has to come before the decision is applied to her and while the recommendation can still change (ICO draft guidance). If her review is relied on, the proposed pair reaches her before commitment with the reason it was chosen; she can change or refuse it, the order waits for her affirmative choice, and the record shows what she reviewed and decided. Without that, a significant decision there is solely automated, with the Article 22B restrictions and Article 22C safeguards set out in answers 26 and 28.

The draft's tests address the organisation's own reviewer, who should "be suitably trained and qualified to understand the system’s logic, outputs, limitations, and risks"; here the human is the data subject herself. Guidance should say whether her own approval can count as meaningful involvement and, if it can, what stands in for training. My answer is yes, on the conditions above plus two more: she sees the data and factors behind the choice, which the draft's tests require a reviewer to take into account, and the alternatives passed over; and, in place of training, she is shown what approval commits her to. In judging her involvement, the agent provider weighs how far the choice rested on profiling, as Article 22A(2) requires.

The retry in the card case is neither of these: the agent presented the open mandate again. The control it skipped is the rule in the Agent Payments Protocol (AP2) against a further presentation without a rejection receipt from the previous one. AP2 bars no status check, so the agent provider asks the merchant, which signs the checkout receipt, or the wallet, which receives the payment receipt, and tells her which outcomes are confirmed and which unresolved before seeking any new approval, since a fresh approval while the first payment is unresolved can produce two purchases.

A new purchase outside her limits, such as a price above the cap or a merchant outside an allowed-merchants constraint, goes to her on a trusted surface before commitment, with the change shown. The record holds what she saw, could change and did, and an unanswered prompt does not count as approval.

28 How do organisations using ADM in the context of agentic AI apply the relevant safeguards?

Please answer below:

Article 22C(2)'s four safeguards (information, representations, human intervention, contest) sit with the controller taking the decision. For a qualifying agent selection (answer 26), that is the party deciding why and how the consumer's data drives the choice, such as the agent provider (UK GDPR Article 4(1)(7)). For a qualifying bank decision it is the bank, which in the card case holds no mandate when it decides; the Agent Payments Protocol requires none to be sent to it. For the bank, that means four measures:

1. The notice identifies the debit by its bank reference; states the outcome, the grounds the bank can lawfully share, and which records were held and which unresolved; and says how to give her account, contest the decision and obtain review. Where telling her of a money-laundering report or investigation would be an offence under section 333A of the Proceeds of Crime Act 2002, only that detail is withheld; the challenge route is still given and the full reasons stay in the restricted file.
2. The bank obtains the mandate records under evidence-sharing terms saying who releases each record, on what request and how fast; its refund, due by the end of the business day after it becomes aware of the unauthorised payment, does not wait for them.
3. Her access requests (Article 15) to each controller are a separate route; the review proceeds without them.
4. Human intervention comes from a trained reviewer with authority to change the outcome, who weighs the data the decision used, any further information or evidence she provides, and whichever signed mandates and receipts the bank obtained. Where a record fails to arrive, the review completes on the evidence held and records the gap, outcome and reasons. The missing record does not discharge the bank's burden of proof (regulation 75 of the Payment Services Regulations 2017), and a reversal needs proof she authorised the payment or was otherwise liable (FCA, Payment Services and Electronic Money: Our Approach, paragraph 8.210).

29 How can technical logging and audit mechanisms support compliant multi-stage ADM through improved transparency and accountability?

Please answer below:

Theme 4 asks organisations to be able to reconstruct, where an agent behaves unexpectedly, "which agent was involved", "what authority it was operating under", "what systems or data it accessed" and "whether any human involvement was meaningful". In an agent payment the question of authority has three separate answers:

1. Payment authority: the consent the payer's bank recognises under its agreed procedure, to one payment or to a series (regulation 67 of the Payment Services Regulations 2017).
2. The lawful basis for each processing step, consent included where relied on (UK GDPR Article 6).
3. Human involvement in the decision affecting her (Article 22A(1)(a)).

A single "I agree" at card enrolment cannot answer all three, so the log records the answer to each at each stage. It links each debit to its checkout under the Agent Payments Protocol (AP2) through the processing record in answer 21, and keeps the mandate chain beside the records made by the code signing for the agent and by each party receiving a mandate (answer 22). An investigator can then find where the agent left her instruction. In the card case, that point is the wallet's record of her shoe mandate arriving a second time, for a different checkout, joined to the signer's record of that closed mandate, with the agent instance and key, task and permissions then in force. The trail then runs on through the decision she contests: the data and any score relied on, what the reviewer saw and could change, the outcome and reason, and any change on challenge.

30 How are organisations using permission setting and user interface design to support compliance with the ADM provisions?

Please answer below:

The interface should keep three screens apart: one for the mandate, one for the payment and one for the data. The Agent Payments Protocol (AP2) defines the first, its trusted surface, and says the agent should offer a way to manage active mandates, with notifications to the user. The payment and data screens, and a status notice and suspend control for an unresolved attempt, are the design the card case calls for, which guidance should ask deployers to show.

On the trusted surface the person sees and approves the constraints: the item, the price cap, any merchant constraint and an expiry AP2 recommends keeping as short as the task allows. Where her review of the particular pair of shoes is relied on, the proposed pair and the reason for the choice reach her there before commitment, and the order waits for her approval (Article 22A(1)(a)). She then leaves that surface, so the first attempt's status reaches her in the agent's or wallet's notifications, with a control to suspend further use while it is unresolved.

The payment screen states whether she authorises one payment or a series (regulation 67(1) of the Payment Services Regulations 2017). Her mandate carries no reuse condition (AP2's "Agent Recurrence" constraint), so the screen says one purchase, and where her bank's procedure says the same, a second payment falls outside her consent. The data screen states what personal information goes to which recipient and on what basis.

Section 6: Fairness and purpose limitation

31 Can you provide examples of instances where organisations have assessed the compatibility (or non-compatibility) of a new processing purpose?

Please answer below:

32 What are the guardrails and control measures that support compliance with the purpose limitation and fairness principles?

Please answer below:

33 Do you have examples of where monitoring of agent actions, data access and tool interactions, alongside other safeguards can reduce the risk of AI agents using personal information for purposes that have not been appropriately assessed and authorised?

Please answer below:

Section 7: Lawful basis

34 Which lawful bases do organisations use for the different processing activities involved in agentic AI development and deployment?

Please list the processing activities and the lawful bases used :

35 What challenges do organisations face in identifying the appropriate lawful basis for their agentic AI processing?

Please answer below:

36 What analysis or evaluation supports organisations' decisions to identify a given lawful basis?

Please answer below:

37 How do organisations conduct Legitimate Interests Assessments for agentic AI data processing, including training or adapting orchestrator models?

Please answer below:

Section 8: Additional questions

38 Which best describes your organisation's current and expected use of agentic AI?

Please select one option in each column - No current plans for adoption:

Please select one option in each column - Exploring potential use cases:

Please select one option in each column - Pilot or proof of concept:

Please select one option in each column - Limited deployment:

Please select one option in each column - wide deployment:

39 Please describe any current or planned uses of agentic AI, including: - the business function or activity involved; - whether the use is internal, customer-facing or both; and - the types of decisions, actions or workflows it supports.

Please answer below:

40 To what extent have the following issues acted as barriers to developing, obtaining or deploying agentic AI systems?

Please select one option in each row - Lack of regulatory certainty over data protection requirements:

Please select one option in each row - Security concerns:

Please select one option in each row - Technical limitations:

Please select one option in each row - Organisational capability and skills:

Please select one option in each row - Cost of implementation:

Please select one option in each row - Cost of compliance safeguards:

Please select one option in each row - Supplier assurance or due diligence challenges:

Please select one option in each row - Other (please specify):

If you chose 'Other', please specify: :

Please explain any factors identified as moderate or major barriers: :

41 Has your organisation delayed or abandoned plans to deploy agentic AI because of uncertainty relating to data protection requirements?

Not Answered

If you answered "Yes", please explain your answer.:

42 In the absence of specific regulatory guidance on agentic AI, how does your organisation obtain assurance that its use complies with data protection requirements?

Other (please specify)::

Please explain your answer: :

43 To what extent, if at all, does uncertainty about data protection requirements affect your organisation's decisions about whether and how to develop, obtain or deploy agentic AI?

Not Answered

44 If applicable, what are the main data protection challenges, uncertainties or areas of difficulty you have encountered in developing, obtaining or deploying agentic AI where greater regulatory certainty can be beneficial?

Please answer below:

45 What benefits has your organisation experienced, or does it expect to experience, from deploying agentic AI?

Other (please specify)::

Please provide examples and, where possible, any evidence or estimates of the benefits experienced or expected. :