

esomar

10 Questions Buyers Should Ask Suppliers of AI-Based Research Services

Due diligence for AI-enabled market, opinion and social research and data analytics

How this document fits

This checklist replaces the buyer-facing part of Esomar's existing 20 Questions. It is designed to be used with the broader 10 Questions for Using AI in the Research Process, the Esomar AI Glossary and the Companion Guide on Using Research Data to Develop and Improve AI. It is an introductory due diligence framework, not a complete specification for every AI service. More detailed ICC/Esomar Code Applied guidance is being developed for particular research methods and uses.

Purpose and audience

The primary audience is the research buyer or commissioner. Depending on the service and level of risk, the buyer may need support from procurement, Legal, Privacy, Security, IT, data science or other specialists. The research buyer remains responsible for assessing methodological suitability and the implications for the research purpose.

The questions apply to external suppliers, embedded platform features, self-service systems and services built or configured for the client. Not every evidence item will be relevant in every case. The depth of due diligence should reflect the sensitivity of the data, the role of AI, participant impact, the autonomy of the system and the consequences of error.

Relationship to the ICC/Esomar Code

The ICC/Esomar Code places responsibilities on researchers, clients and other parties according to their activities, expertise and control. Clients commissioning research must ensure that contractors are aware of the Code and comply with it. These questions help the buyer obtain the information and commitments needed to meet that responsibility.

The 10 buyer questions at a glance

No.	Question
1	What AI is being used in the service, and for what purpose?
2	Who provides, operates and controls the AI system?
3	What data is used by the AI system, and where does it come from?
4	How will our data, materials and outputs be accessed, retained, reused or used to improve AI systems?
5	How do you show the service is valid, reliable and fit for purpose?
6	How do you manage bias, hallucinations, synthetic outputs and other poor-quality results?

7	What professional oversight, review and escalation are built in?
8	How do you protect privacy, confidentiality, security, intellectual property and data across locations and jurisdictions?
9	How do you disclose AI use, synthetic data, synthetic personas and limitations?
10	What evidence, audit trail and contractual commitments can you provide?

1. What AI is being used in the service, and for what purpose?

Why this matters: Buyers need a plain-language explanation of the AI capability, what it does, what it does not do, and why it is appropriate for the research service.

Points to consider

- Describe the AI components and their role in non-technical language.
- State which parts of the service depend on AI and which do not.
- Explain the expected benefit for the research objective, quality, access, speed, analysis or reporting.
- State any use of generative AI, automated analysis, AI moderation or interviewing, synthetic data, synthetic personas or agents.
- Explain whether the service produces a fixed deliverable or remains interactive after handover, and whether it is experimental or validated for production use.

Useful evidence to ask for

- Service description, methodology note, user guide, intended-use statement and limitations statement.

Code connection: Article 7 requires the method, data sources, quality controls, analysis and limitations to be explained sufficiently for the client to assess validity.

2. Who provides, operates and controls the AI system?

Why this matters: The buyer should understand the organisations and systems in the AI value chain, who can access the data, and who is accountable for the service.

Points to consider

- Identify the research supplier, model provider, platform provider, hosting provider and relevant subcontractors.
- Explain who configures the service, accesses data, approves outputs and can change model or system settings.
- State whether the buyer has a choice of model, deployment region, hosting arrangement, retention setting or model-improvement setting.

- Explain responsibilities where the buyer operates a self-service tool or client-controlled instance.
- Confirm who is accountable if the service fails, changes materially, produces poor-quality output or breaches agreed requirements.

Useful evidence to ask for

- Supplier and subcontractor map, accountability statement, governance policy, model or platform documentation and change-management policy.

Code connection: The Code requires responsibilities to reflect each party's activities, expertise and control, and requires appropriate obligations throughout the supply chain.

3. What data is used by the AI system, and where does it come from?

Why this matters: The quality, legality and risk of an AI service depend heavily on the data used for training, reference, retrieval, prompts, inputs and validation.

Points to consider

- Identify training, tuning, reference, retrieval, validation and project input data where this is relevant and reasonably knowable.
- Distinguish participant data, client data, public data, proprietary supplier data, third-party data and synthetic data.
- Explain provenance, rights, quality, recency, representativeness and known gaps.
- State whether relevant languages, minority groups, regions and specialist populations are adequately represented.
- Explain how source data and reference materials are kept current and how outdated or unreliable material is handled.

Useful evidence to ask for

- Data provenance note, data lineage documentation, data quality report, population coverage statement and relevant data or system cards.

Code connection: Articles 5, 7 and 8 require appropriate reuse, transparency about data sources and limitations, and compliance with contractual and intellectual property restrictions.

4. How will our data, materials and outputs be accessed, retained, reused or used to improve AI systems?

Why this matters: The buyer needs explicit agreement about how client and participant materials are accessed and whether they may be retained, reused or used beyond delivery of the current project.

Points to consider

- State how the supplier and its platforms access client and participant data, including through connectors, uploads, APIs or support personnel.
- Clarify whether data, prompts, recordings, transcripts, coded data, reports or other outputs are retained after the project.
- Distinguish use for the current project, future projects for the same client, a client-specific model, a supplier's shared system and a general-purpose or third-party model.
- Confirm that client materials and project outputs will not improve a shared supplier or general-purpose model unless specifically authorised and contractually controlled.
- State whether and on what terms the buyer may use commissioned datasets and outputs in its own internal AI tools or with other authorised partners.
- Confirm that downstream platform terms are consistent with commitments made to the buyer and with participant permissions.

Useful evidence to ask for

- Contract clauses, data-use schedule, access description, retention and deletion schedule, platform settings, model-improvement policy and permitted-use matrix.

Code connection: Articles 5, 6, 7(f), 8 and 11 address purpose, confidentiality, secure environments, privacy and intellectual property. The companion client-data guide gives further practical guidance.

5. How do you show the service is valid, reliable and fit for purpose?

Why this matters: AI outputs can be impressive without being valid. Buyers need evidence that the service is suitable for the particular purpose, population and claim being made.

Points to consider

- Describe validation against source data, independent human judgement, established methods, external benchmarks, primary research or holdout data.
- Explain performance measures, error rates, quality thresholds, failure modes and known limitations.
- Describe parallel testing, repeated testing or stability checks where relevant.

- Explain how the service is tested before deployment and revalidated after material changes.
- State how uncertainty, incomplete evidence and inconsistent results are handled.
- Provide explainability where useful, while recognising that a plausible explanation or a model's self-critique is not independent proof of accuracy.

Useful evidence to ask for

- Validation report, quality metrics, pilot or parallel-test results, testing protocol, model or system card and limitations statement.

Code connection: Article 7 requires research to be fit for purpose and findings, interpretations and recommendations to be supported by the data.

6. How do you manage bias, hallucinations, synthetic outputs and other poor-quality results?

Why this matters: AI can reinforce bias, fabricate content, overstate certainty, omit minority views or produce synthetic material that is mistaken for human evidence.

Points to consider

- Explain checks for demographic, cultural, linguistic, geographic, accessibility and other relevant biases.
- State how hallucinations, fabricated verbatims, false summaries, unsupported claims and false precision are detected.
- Explain how synthetic, inferred and AI-generated material is labelled and separated from human-derived data.
- Describe how the service deals with plausible but incorrect output and with errors that remain within an apparently high accuracy rate.
- Describe correction and learning processes when quality problems are discovered.

Useful evidence to ask for

- Bias assessment, quality-control protocol, synthetic-data labelling policy, error log, output-review checklist and examples of corrected failures.

Code connection: Articles 7 and 8 require relevant biases and limitations to be disclosed and conclusions to be adequately supported by data.

7. What professional oversight, review and escalation are built in?

Why this matters: Professional responsibility cannot be delegated to an AI system. The buyer should understand who reviews the work, what expertise they have and when intervention occurs.

Points to consider

- Identify where qualified professionals review, approve, challenge or reject AI outputs.
- Describe the expertise and training required of reviewers, including their ability to recognise confident but incorrect output.
- Explain the review approach, including any sampling, second review or disagreement thresholds used for material outputs.
- Describe how automation bias, reviewer fatigue and excessive review volumes are managed.
- Explain how participant-facing interactions are monitored, when a person can intervene and how participants can reach a person.
- State what triggers escalation, how quickly material incidents are considered and who can pause or stop the service.

Useful evidence to ask for

- Oversight plan, reviewer competence requirements, escalation process, safeguarding policy, incident-response procedure and complaint-handling process.

Code connection: The Code requires human oversight to be stated when AI is used in data-sets, analysis, reporting or interpretation, and allocates overall responsibility to research-ers.

8. How do you protect privacy, confidentiality, security, intellectual property and data across locations and jurisdictions?

Why this matters: Participant data, client materials, confidential information and intellectual property require distinct but coordinated protection throughout the AI workflow.

Points to consider

- Describe data minimisation, anonymisation or pseudonymisation, access controls, retention, deletion and secure hosting.
- State whether data is processed in public tools, private instances, controlled environments or on-device systems.
- Explain data location, cross-border transfers, applicable jurisdictions and regional hosting options.

- Assess re-identification, inferred identity and sensitive-inference risks, including linkage with client, public or secondary data.
- Explain how participant privacy is protected separately from client confidentiality, commercial information and intellectual property.
- Confirm that subcontractors and platforms maintain equivalent protections and that relevant Legal, Privacy, Security or IT specialists can review the arrangements.

Useful evidence to ask for

- Risk or impact assessment, security certification, data-processing agreement, retention schedule, IP and confidentiality terms, transfer safeguards and breach-response process.

Code connection: Articles 3, 6, 7(f), 8 and 11 address minimisation, inferred identity, security, transfers, confidentiality and intellectual property.

9. How do you disclose AI use, synthetic data, synthetic personas and limitations?

Why this matters: Disclosure requirements differ according to whether information is being given at proposal stage, to a participant, with a deliverable, in a report or to the public.

Points to consider

- Explain what will be disclosed in the proposal or commissioning documentation.
- Provide participant-facing disclosure for synthetic personas and for material direct interaction with AI interviewers, moderators, chatbots, voice tools or avatars.
- State the extent of human monitoring or intervention where this matters to participants or the buyer.
- Explain how AI, synthetic data and AI-generated material will be identified in datasets, analysis, reports and interpretation.
- Provide appropriate wording for public methodology statements where AI or synthetic methods play a significant role.
- Disclose limitations in a form that is useful to the intended audience rather than merely technically complete.

Useful evidence to ask for

- Participant notice, proposal disclosure, client reporting template, dataset labels, publication methodology note and examples.

Code connection: Articles 4(a)(ii), 7(e) and 9(b) specify disclosure of synthetic personas, AI use and the extent of human oversight.

10. What evidence, audit trail and contractual commitments can you provide?

Why this matters: Due diligence should lead to usable evidence and commitments, not only reassuring explanations.

Points to consider

- Provide a proportionate audit trail covering the tool, material model versions or settings, data sources, validation, professional review and important changes.
- Commit to confidentiality, permitted-use limits, data deletion, security, subcontractor controls and no unauthorised model improvement.
- Agree how errors, complaints, disputes, breaches, corrections and requests for independent quality checks will be handled.
- Commit not to fabricate, falsify or misrepresent research data or findings through AI or any other means.
- Confirm compliance with the ICC/Esomar Code, applicable law, local codes and agreed client requirements.
- Agree how material changes to models, providers, data use or performance will be notified.

Useful evidence to ask for

- Supplier due diligence pack, contract clauses, audit rights, compliance statement, documentation retention policy, incident terms and change notification process.

Code connection: Articles 7 to 12 cover technical information, confidentiality, honesty, legality, independent checks, dispute resolution and compliance throughout the research chain.

Optional buyer assessment

The research buyer may use a simple traffic-light assessment for each question. The assessment should reflect the importance of AI use and should not be treated as a price-led procurement exercise.

Green	Adequate evidence and contractual commitment for the proposed use.
Amber	Potentially acceptable, but clarification, additional evidence or controls are needed.
Red	Insufficient evidence, unacceptable risk or no clear commitment.

esomar