Dynata's responses to ESOMAR's 20 questions for AI-based research services

Version 1.0 | 01.14.2026

dynata™

dynata

Table of contents

dynata

Company profile

1. What experience and know-how does your company have in providing AI-based solutions for research?

Dynata combines more than 40 years of market research experience with advanced AI and data science. AI/machine learning has been used for over 6 years in our participant (respondent) quality systems. In 2022, Dynata formed our Global Research and Data Sciences Center of Excellence, which unifies research, panel/sampling science, and AI/data science teams to drive innovation across insights delivery.

Currently, AI/ML is integrated into multiple aspects of Dynata's fraud detection approach, participant quality scoring, participant matching, automated analytics, and conversational research tools. Dynata's AI systems enhance survey accuracy, efficiency, and deeper insights, enabling deeper insights at scale. Recognized as the "Most Innovative Field Services Supplier" in the 2024 Greenbook GRIT Report and quality-certified by Neutronian's NQI in 2023, Dynata demonstrates consistent leadership in trusted, AI-powered research solutions.

2. Where do you think AI-based services can have a positive impact on research? What features and benefits does AI bring, and what problems does it address?

AI brings a number of positive benefits to the research process, enhancing research by improving data quality, efficiency, and insight depth.

Some examples of these benefits at Dynata include:

- Dynata's AI-powered QualityScore™ analyzes 175+ behavioral and response factors to detect fraud and inattentiveness, removing poor data in real time and reducing manual cleaning by up to 85%.
- AI improves survey matching, boosting qualification and completion rates.
- Dynata's Quali-Quant AI uses generative AI as a virtual moderator, asking adaptive follow-ups to uncover motivations at scale—producing 2.8× more detailed feedback and 99% higher engagement.
- AI automates translation, data coding, and survey scripting, accelerating delivery while freeing researchers for higher-level analysis.

Overall, AI addresses data accuracy, timeliness, and engagement challenges, enabling faster, richer, and more actionable insights.

3. What practical problems and issues have you encountered in the use and deployment of AI? What has worked well and how, and what has worked less well and why?

Deploying AI in research has revealed both new strengths and new challenges. One of the most pressing is AI-assisted participant fraud, where bots or automated tools generate fake open-ended answers. Dynata combats this with our multi-layered QualityScore™, which analyzes behavioral signals such as typing pace, copy-paste activity, and mouse movement to detect non-human patterns. We revised QualityScore in anticipation of ChatGPT's release, and we continue to retrain it as GenAI evolves. Our quality system is intentionally holistic and multi-layered, combining AI detection with non-AI methods and human verification, to protect data integrity.

Another challenge is bias and inaccuracy in generative models. Quali-Quant AI requires fine-tuning to keep questions neutral and relevant, achieved through domain training and expert review. Balancing automation with human oversight is critical: AI scales qualitative probing efficiently but cannot replace human moderation for sensitive topics.

Ultimately, the most effective approach blends the best of both worlds. By integrating human expertise into AI feedback loops, continuously refining models, and using AI where it excels—speed, scale, and detection—we preserve the nuance, trust and rigor that high-quality research demands.

AI capabilities and services

4. Can you explain the role of AI in your service offer in simple, non-technical terms, in a way that can be easily understood by researchers and stakeholders? What are the key functionalities?

AI is the intelligent engine powering Dynata's research services, helping researchers reach the right people, ensure data quality, and uncover deeper insights. It operates seamlessly in the background, much like a smart research assistant that improves speed, accuracy, and insight depth across every stage of the process.

First, AI improves how we find and engage the right participants. Using advanced machine learning, Dynata's systems analyze hundreds of attributes and past behaviors to match participants with the surveys they are likely to qualify for and complete. This intelligent matching improves response speed, boosts completion rates, and ensures studies reach the intended audiences. Participants benefit from receiving more relevant surveys, and clients gain high-quality responses from the right individuals without delay.

Second, AI continuously ensures data quality. Dynata's QualityScore™ acts as an automated referee, reviewing each response in real time. It evaluates behavioral cues—such as typing rhythm, completion speed, and interaction patterns—to detect inattentive or fraudulent participants. This proactive process removes unreliable data before it reaches clients, resulting in cleaner, more trustworthy insights.

Third, AI deepens understanding of the "why" behind behaviors. Our Quali-Quant AI capability functions as a virtual moderator, engaging thousands of participants in dynamic, personalized interactions. It asks relevant follow-up questions based on a participant's open-ended answers, providing richer qualitative insight at quantitative scale. The result is more detailed, emotionally resonant feedback in less time.

Beyond these areas, AI components enhance efficiency by automating translation, text analytics, data coding, and audience activation.

In short, AI enables Dynata to deliver faster, higher-quality, and more insightful research. It streamlines the process for clients and researchers alike—turning complex, time-consuming tasks into simple, automated, and reliable steps that lead to smarter, data-driven decisions.

5. What is the AI model used? Are your company's AI solutions primarily developed internally, or do they integrate an existing AI system and/or involve a third party and if so, which?

Dynata employs a suite of AI models and methods, each chosen or developed for the specific task in our research workflow. Many of our AI solutions are proprietary models developed internally by our data science team, often augmented with leading third-party AI technologies for certain functions.

For example, our *QualityScore™* system is powered by a custom machine learning model that we built and trained in-house on Dynata's vast trove of survey data. This model uses algorithms (such as gradient-boosted decision trees and other ML techniques) to evaluate participants in real-time. It was developed by Dynata's data scientists to be highly specialized for market research fraud detection and participant scoring. All the logic – from what data features are examined (e.g., response times, consistency, device info) to how to weight them – was crafted internally based on years of participant management experience. This proprietary ML model is continuously refined with new training data as we observe emerging behaviors, ensuring it stays effective against evolving fraud patterns.

We also leverage state-of-the-art third-party AI where it adds the most value, particularly for Natural Language Processing. Our Quali-Quant AI solution leverages advanced large language models (LLMs) – the same class of AI behind tools like OpenAI's GPT or similar – to enable human-like conversational abilities. Rather than reinventing the wheel, we harness these proven AI engines (via secure partnerships/APIs) as a foundation for understanding and generating text. We then layer in Dynata's best-practice approaches. In practice, the generative AI engine may come from a third-party model

trained on large-scale language data, but it is guided by Dynata's proprietary framework to ensure it asks relevant follow-up questions and uses research-appropriate language. While we do not publicly disclose specific vendor names in all cases, we partner only with reputable providers for LLM capabilities, ensuring they meet our data-privacy, quality, and research-grade standards.

Overall, our approach is a hybrid one: build in-house when it's a core competency or requires our unique data (as with QualityScore™ and our sampling algorithms) and buy/partner when an external AI technology can be better leveraged (as with some natural language and translation services). Even in the latter case, we often customize or fine-tune those models. For instance, when we use a third-party language model for Quali-Quant, we will train it further (or instruct it) on market research wording so it fits our purpose. All integrated AI components are vetted thoroughly. The result is that Dynata's AI solutions use the "best of both worlds" – our own proprietary AI innovations where we lead the industry, seamlessly combined with best-in-class AI systems from trusted partners. This ensures clients benefit from robust, up-to-date AI technology that is optimized specifically for Dynata's insights platform.

6. How do the algorithms deployed deliver the desired results? Can you summarize the underlying data and the way in which it interacts with the model to train your AI service?

Our AI algorithms work by learning from and acting on the rich data flowing through Dynata's research ecosystem. In essence, each model has been trained on relevant, high-quality data and is designed to use new input data in intelligent ways to produce accurate results. We can break this down with two key examples:

QualityScore™ ML model (data quality algorithm) – the underlying data consists of thousands of survey records and behavioral signals collected from our panelists over time. We feed the model a wide array of features about participant behavior: e.g., how fast they answered, if they straight-lined answers, the coherence of their open-end text, whether their device or geo-location was suspicious, and so on. During training, the model saw examples labeled as "high quality" (genuine, attentive participants) and "low quality" (fraudulent or disengaged cases) as determined by our research experts. Through this training, the algorithm learned patterns that distinguish good data from bad.

When deployed, it interacts with incoming survey data in real time – each new participant's behavior metrics are run through the model, which then outputs a QualityScore. If the score falls below a threshold, our system will automatically remove and replace that participant during fieldwork. The model design ensures it is considering the interplay of many factors (for instance, a slightly fast completion alone will not flag someone, but fast completion *plus* inconsistent answers *plus* copy-paste behavior would). This holistic approach, learned from the training data, is how the algorithm consistently delivers the result we want: a clean, reliable dataset for our clients. We continually validate and update the model with new data – for example, once generative AI became a possible source of fraudulent answers, we retrained QualityScore to recognize the subtle signals of AI-generated text (such as lack of variance in response patterns and certain metadata). Under the hood,

it is a sophisticated statistical model, but to the end user its impact is simple: survey results with the bad data already filtered out.

Quali-Quant AI (generative language model) – the algorithm is powered by a large language model that has been pre-trained on vast amounts of text data (books, articles, conversations) by our technology partners, giving it a broad base of "knowledge" about language. We then fine-tune and configure it with Dynata's own research data and guidelines. Concretely, during survey-taking we provide the AI with context such as the survey question text, the participant's initial answer, and any relevant background (like product category or scenario). The model interacts with this input by analyzing it using its learned language patterns, and then it generates a follow-up question or probing prompt that is on-topic and clear. For instance, if a participant answers an open-end about why they prefer a certain brand with a brief statement, the AI model takes that answer as input and generates something like, "*Thank you for your answer. Could you tell us a bit more about what you particularly like about that brand's products?*" – effectively continuing the conversation. This is possible because the model was trained on how dialogues flow and how to ask clarifying questions. We ensure it stays *"fit for purpose"* by training it with example Q&As from market research interviews and by imposing rules (via prompt engineering) so it focuses on the participant's data and not outside information. The underlying data guiding this is both the general language data (from pre-training) and Dynata's specific dataset of past survey responses and probes which we use to calibrate it.

At Dynata, data is the fuel for our AI strategy. We carefully curate training datasets from our panel (for ML models like QualityScore™) and use our extensive library of participant data to inform generative models. These algorithms then interact with live incoming data in defined ways: scoring models compute probabilities based on input behaviors, while generative models produce natural-language output based on input text. Throughout, we document and monitor how the data flows. For example, we know the "lineage" of what data went into training QualityScore™ and how it's applied, and we ensure that any personal data is handled in compliance with privacy and legal requirements.

It's this synergy of high-quality data and purpose-built model design that ensures our AI services reliably produce insights aligned with clear research objectives. In practice, that means faster answers, cleaner datasets, and richer insights, all achieved by the model continuously learning from and reacting to the data it's given.

Trustworthiness, ethics and transparency

7. What are the processes to verify and validate the output for accuracy, and are they documented? How do you measure and assess validity? Is there a process to identify and handle cases where the system yields unreliable, skewed or biased results? Do you use any specific techniques to fine-tune the output? How do you ensure that the results generated are 'fit for purpose'?

Dynata employs robust validation processes at multiple levels to ensure our AI outputs are accurate and fit for purpose. Each AI-driven feature is rigorously tested and documented before and after deployment. For example, our QualityScore™ model was validated against known truth data: we benchmarked its flags versus manual data quality checks by our research experts to make sure it was catching bad cases and not falsely flagging good participants. We continually monitor the performance of our overall quality system and associated AI models using key metrics such as the sample acceptance rate (the percentage of survey responses clients deem high quality).

Our sample acceptance rate is currently very high, at 95-96%. These metrics are documented in our quality reports. If any drift in accuracy is detected (e.g., if acceptance rate drops or an uptick in undetected issues is found via auditing), we immediately investigate and fine-tune the model. QualityScore™'s performance is also independently audited in some cases – notably, Dynata's data quality processes (including QualityScore™) were examined as part of our Neutronian certification, which attests to the rigor of our validation methods.

For generative AI outputs like those from Quali-Quant, our validation involves a human-in-the-loop review during development. We test the AI on pilot surveys and have researchers evaluate the follow-up questions it asks and the summaries it produces. These are checked for accuracy (e.g., does the summary correctly reflect the responses?) and bias (e.g., is the follow-up question neutral and not leading?). We document these test cases and outcomes. If the AI ever produces an off-target or inappropriate question, we adjust the underlying prompts or model parameters (this is part of our fine-tuning process). In practice, we've found the AI's probes to be on-point and beneficial – for instance, we measure that the AI-generated probing leads to significantly longer and more detailed responses (an indication that the UX is effectively eliciting more usable insights). We also keep logs of AI queries and responses to audit them for quality over time, which serves as ongoing documentation of performance.

When the system identifies unreliable or skewed participant/result, we have clear handling procedures. For our survey quality AI, any participant that triggers certain thresholds (e.g. too many poor-quality indicators) is flagged and either removed or set aside for additional verification. Those rules are codified in our system and documented in our internal knowledge base. Similarly, if a

generative AI component were to produce an obviously irrelevant or biased output, our researchers can override or exclude that output. In such cases, we analyze why it happened – whether the training data had a gap or the model needs an adjustment – and implement a fix. One example was early in development of Quali-Quant, when the AI occasionally asked a redundant question; we identified the cause (the model giving too much weight to certain phrases in the prompt) and refined the prompt structure, eliminating the issue. All these refinements are part of our continuous fine-tuning cycle.

To ensure results are fit for purpose, we align our AI outputs with known research standards and client expectations. For instance, we validate that an AI-generated summary of open-ends matches the themes a human analyst would derive by doing side-by-side comparisons on test data. We also rely on external frameworks and benchmarks where available. Our approach to validation is very much *document, measure, adjust*: each AI solution has a technical documentation outlining how it's tested (e.g., what sample data was used, what performance metrics are monitored). We measure validity through outcome metrics (data quality scores, response rates, insight depth) and even client feedback. If a client ever questions an AI-derived result, we have the information to trace its origin (data lineage) and double-check it.

In summary, multiple safeguards ensure accuracy. These include automated checks within the AI (like fallback rules if a model is uncertain), regular accuracy measurements (such as the sample acceptance rate and other quality KPIs), and human oversight to catch anything the AI might miss. Thanks to this, we are confident that the AI's contributions – whether cleaning data or generating insights – maintain the high standards of quality that Dynata promises. Our validation processes are living protocols; they are reviewed and updated as needed, and they have been key in earning industry trust in our AI-based services.

8. What are the limitations of your AI models and how do you mitigate them?

While AI is a powerful tool in our research arsenal, we recognize its limitations and take steps to mitigate them. One limitation is that AI models are only as good as the data and rules they are built on. For instance, a machine learning model might not perform well on scenarios it hasn't "seen" before. We've observed this in edge cases of data quality detection – e.g., if a completely new kind of participant behavior emerges (perhaps due to a new fraud tactic or an unusual survey design), the model might not immediately flag it. Our mitigation strategy is to continuously retrain and update our models with fresh data. For example, Dynata's AI team operates a schedule of periodic reviews where we feed the latest panel behavior data into QualityScore™, ensuring it evolves with the times. By doing this, we keep the model's knowledge current and minimize blind spots. We also employ ensemble approaches – multiple models/checks – so that even if one model has a blind spot, another mechanism can catch an issue. For example, if an AI text analysis misses that an open-end response is nonsensical, a simple business rule (like a keyword blacklist or a minimum character count rule) might catch it. Having layered defenses is one way we mitigate AI limitations.

Another inherent limitation is that generative AI, like the kind used in Quali-Quant, does not have true understanding – it can occasionally produce irrelevant or generic output if not guided properly. We mitigate this by giving the AI very specific context and instructions (in technical terms, prompt engineering and fine-tuning). Essentially, we fence in the AI to talk only about the participant's answers. It won't drift off-topic because we program it to focus on what was said and the research objective at hand. Additionally, we acknowledge that AI cannot fully replicate human judgment or creativity in certain areas. For example, while Quali-Quant AI can probe for deeper insights at scale, it doesn't entirely replace a live moderator for very nuanced or sensitive discussions. We mitigate this by clearly positioning the tool for appropriate use cases – it's used to augment our research, not to handle tasks it isn't suited for. Our team remains involved in reviewing AI outputs, especially in critical studies, to ensure nothing is misinterpreted or lost in nuance.

Bias is another potential limitation: if the data used to train an AI contains bias, the output could reflect that. Dynata is very mindful of this risk. We mitigate bias in our AI models firstly by using diverse, high-quality training data. Our panel is extremely diverse demographically and geographically, which means the data feeding our algorithms comes from a wide range of populations, reducing skew. For our synthetic data or any AI-driven content creation, we apply checks – for instance, we test that our Virtual Audiences (an AI-based data generation product) produce outputs consistent with real-world data distributions, and we correct any biases detected in those tests (by adjusting the model or input balance). In practice, when creating AI-generated insights, we measure things like topic coverage to ensure the AI isn't overly focusing on one aspect due to some bias in its training.

We also design our AI to not introduce bias where none existed: as Yabble noted in a similar context, if the input data is biased, the AI will reflect that bias but not amplify it. For Dynata, this means our AI won't magically fix a non-representative sample (that's handled by good sampling practices), but it also won't skew a representative sample – it works within the data given. Knowing this, we mitigate by ensuring upstream processes (like panel recruitment and sampling) yield balanced data, so the AI's output is built on a solid foundation.

Finally, transparency itself can be a limitation if not addressed – AI can be a "black box." We mitigate this by maintaining transparency with our clients about how the AI works and where its limits are. For instance, we inform clients that QualityScore™ identifies likely poor-performing participants but isn't an indictment of a person – there's a small false-positive rate, which we counteract by not over-purging data (we only remove when multiple indicators concur, as a safeguard against erroneously removing a good participant) . In effect, we design our usage policies to lean on the side of caution. If there's uncertainty, we might hold that data for manual review rather than automatically trusting the AI. This human fail-safe mitigates the limitation of AI confidence levels.

In summary, our mitigation strategies include retraining models regularly, combining AI with human oversight, enforcing strict usage boundaries, bias testing, and transparency about AI limitations. By confronting these limitations head-on, we ensure our AI remains a reliable asset rather than a risk. The result is that clients get the benefits of AI – speed, scale, consistency – while we proactively control the downsides.

9. What considerations, if any, have you taken into account to design your service with a duty of care to humans in mind?

Dynata designs all AI-enabled services with a strong duty of care to both research participants and clients, guided by robust governance, privacy, and ethical oversight. Our approach ensures that every innovation enhances research without compromising human welfare, dignity, or trust.

We prioritize privacy and data protection. All AI systems comply with GDPR, CCPA, and other global privacy standards. AI is used solely for legitimate research purposes, never to retain or expose personally identifiable data. Every participant participates with informed consent, and responses are handled confidentially. Governance reviews each AI use case to confirm that it upholds participants' rights, and features that pose privacy risks are modified or rejected.

We minimize participant burden. Our Quali-Quant AI, for instance, integrates conversational probes directly into surveys, reducing the need for long or repetitive follow-ups. Engagement metrics show 99% of participants remain active through AI-led sessions, confirming that the experience is respectful and non-intrusive. Content guardrails prevent biased, personal, or inappropriate questions, ensuring AI behavior follows the same ethical standards as human researchers. Participants always retain autonomy to skip questions or withdraw.

Third, we ensure the ethical use of AI outputs. Clients are informed whenever AI contributes to data analysis or predictions, maintaining transparency and context. We prohibit AI uses that could lead to discrimination, unfair targeting, or misuse of data. Our Business Code of Conduct explicitly links AI deployment to fairness, integrity, and compliance with research ethics.

Human oversight remains central. Every AI process includes accountable staff monitoring performance and outcomes. AI decisions can be reviewed, audited, and overturned when necessary. Our governance framework mandates documentation of AI logic, escalation protocols, and client transparency.

Dynata aligns with industry association standards (such as ESOMAR and Insights Association), embedding participant welfare and data subject rights into the AI lifecycle. In essence, our AI systems are designed to serve people—protecting participants, empowering researchers, and maintaining trust through responsible innovation and rigorous human governance.

10. Transparency: How do you ensure that it is clear when AI technologies are being used in any part of the service?

Dynata views transparency in AI use as essential to building trust with clients, participants, and the broader research community. We make it clear whenever and wherever AI is part of our services through consistent labeling, documentation, and open communication.

We start with explicit product naming/messaging. Offerings such as QualityScore™ and Quali-Quant AI are marketed to indicate their AI-driven nature. When we launched Quali-Quant AI, we published clear explanations and product sheets describing it as an AI-powered capability for qualitative insights at scale. Similarly, QualityScore™ is presented as an AI/ML-based quality assurance system. By naming and describing products this way, we ensure clients understand exactly when AI is at work.

Transparency also extends to documentation and client communication. In proposals, user guides, and methodology reports, we identify features like “AI Open-End Quality Checks” or “AI-Based Probing.” Deliverables generated or enhanced by AI include clear notations such as “This summary was generated using Dynata’s AI analytics.” This helps users interpret results appropriately. If clients ask how an insight was derived, we can explain the AI tool or model used in plain, accessible language. Internally, detailed methodological documentation supports this transparency and is shared with clients upon request.

Our teams are trained to discuss AI openly and accurately. Sales, account, and research staff are encouraged to highlight AI as a value-adding capability, not as an invisible process. For example, we might explain, “Our AI-driven QualityScore™ identified and removed 5% of low-quality responses, improving data reliability.” This reinforces both the utility and accountability of AI.

We also engage publicly through webinars, white papers, and industry events. Sessions like “Inside the Machine: Dynata’s Quality System Revealed” explain how our AI/ML technologies detect and prevent fraud. This openness extends to our platform itself—AI-driven functions are denoted with icons or notes like “Auto-coded by Dynata AI.”

Participant transparency is equally important. While most AI operates seamlessly in the background, our panelists are informed via terms and conditions about automated processing. If a survey involves direct AI interaction (for example, a chatbot), we introduce it clearly within the experience.

Finally, any reports and deliverables that use AI include methodological disclosures about AI’s role. This approach aligns with ESOMAR’s transparency principles and ensures clients understand, trust, and value the sophistication of Dynata’s AI systems.

11. Do you have ethical principles explicitly defined for your AI-driven solution, and how in practice does that help to determine the AI's behavior? How do you ensure that human-defined ethical principles are the governing force behind AI-driven solutions?

Yes, Dynata operates under a set of explicit ethical principles that guide all our research services – including those powered by AI – and we actively ensure these human-defined values govern our AI's behavior. While we haven't published a separate "AI ethics manifesto" to the public, our AI development is firmly rooted in the broader ethical frameworks and codes of conduct that Dynata abides by as a company.

Here's how our principles are defined and implemented in practice:

Core ethical principles:

1. **Respect for persons and data privacy:** We uphold the principle that participants' rights and privacy come first. This is reflected in strict compliance with data protection laws and our internal privacy standards. For AI, this means we design our systems never to override consent or misuse personal data. For instance, any personal data used in model training is anonymized or aggregated, and we do not allow AI to infer personal identities or sensitive attributes beyond what is permissible.
2. **Beneficence and quality:** We aim to maximize the benefit and quality of insights while minimizing any potential harm (such as misinformation or bias). This translates to AI behavior that strives for accuracy and fairness. If an AI model were ever to suggest a course of action that conflicts with participant welfare or data integrity, our human-led ethics would intervene to correct that.
3. **Justice and non-discrimination:** Our solutions must be fair and not systematically disadvantage any group. We embed this by testing our AI outputs for bias and by using diverse training data. Moreover, our Global Data Quality Pledge (Insights Association) and industry commitments (adherance to the Insight Association and ESOMAR code of ethics) ensure we focus on inclusivity and representation in research data. An AI-driven insight that seemed to unfairly skew against a demographic would trigger a review under these principles.
4. **Transparency and accountability:** As discussed in Question 10, transparency is a key ethical stance. Additionally, accountability means a human is always responsible for the outcome, not "the AI" in isolation. We have internal governance (committees and leadership oversight) that hold our AI use accountable to our code of conduct.

These principles are explicitly defined in documents like Dynata's Business Code of Conduct and company values, which stress integrity, quality, and trust. They may not mention "AI" by name, but they absolutely apply to any technology we use. We also align with external ethical guidelines: Dynata team members are active in industry associations (ESOMAR, Insights Association, etc.) that have published guidelines on AI ethics in research, and we incorporate those into our policies. For example,

ESOMAR's guidelines on artificial intelligence emphasize human oversight and participant welfare, which we have adopted as standard operating procedures.

Ensuring principles govern AI behavior in practice:

- Ethics by design: When developing an AI feature, we have checkpoints to consider ethical implications. For instance, during design and quality assurance of our Quali-Quant AI product, we reviewed the question prompts to ensure they were unbiased and appropriate (reflecting our non-discrimination and respect principles). Any prompt that might lead to an uncomfortable or misleading interaction was reworded or removed.
- Human oversight and intervention: Our ethical principles manifest through human control. If an AI's decision or output might breach an ethical norm, a human in the loop will catch and correct it. For example, if QualityScore™ consistently flagged participants from a particular country at a higher rate, our team would notice and investigate to ensure there's a valid reason and not a hidden bias. We actively monitor for such scenarios. We also have clear escalation paths – for example, our Data Privacy Officer and legal team are involved in vetting any new AI that touches personal data, ensuring it aligns with our privacy principles.
- Training and awareness: We train the teams who work with AI (product managers, data scientists, operations) on our ethical standards. They are encouraged to question "Should we do this?" not just "Can we do this?" with AI capabilities. This culture means that ethical considerations are not an afterthought; they're part of the feature requirements. A practical example: when implementing an AI-based translation service to help with multilingual surveys, the team decided that it must not translate open-end responses that contain personally identifying information (to avoid any inadvertent exposure of PII). That decision was driven by our principle of data privacy.
- Auditing and validation against principles: We periodically audit our AI systems for compliance with ethical expectations. This might include reviewing a sample of AI-generated content to ensure tone and substance meet our standards. If we find an anomaly (say an AI summary that could be misconstrued), we refine the system. Moreover, external validations like the Neutronian data quality certification also indirectly audit us on ethical handling of data, reinforcing that our processes (AI included) meet high ethical benchmarks.

An example of principles in action is our stance on data usage. We contractually and technically ensure that any data or participant data processed through our AI is not used to train third-party models or for any purpose outside the client's project. This is a direct result of our ethical commitment to client confidentiality and privacy. We ringfence client-specific data, so if we use an external AI service, we do so in a way that they cannot harvest the data (often through enterprise agreements or opting out of data sharing) – this practice was inspired by ethical guidelines and ensures the AI behaves in a way aligned with our principles, not just its creators' intentions.

In summary, human-defined ethics steer our AI at every turn. We explicitly define what is acceptable and what is not, and those rules become part of how the AI is configured. Our leadership and governance structures keep a watchful eye, and we foster a company-wide awareness that

technology must serve our values, not override them. This ensures that when you use a Dynata AI-driven solution, you can trust that a strong moral compass is influencing how that solution was built and how it operates.

12. Responsible Innovation: How does your AI solution integrate human oversight to ensure ethical compliance?

Dynata's AI solutions are built with a "human in the loop" philosophy at their core. We firmly believe that human oversight is essential for responsible AI, and we have embedded multiple layers of human checks and balances to ensure ethical compliance throughout the AI lifecycle – from development to deployment to ongoing use.

Here are the key ways we integrate human oversight:

- **Expert review committees:** Before an AI-driven feature is rolled out, it is reviewed by relevant experts and stakeholders within Dynata. For example, during the creation of QualityScore™, a committee of senior researchers (who understand participant behavior) and data scientists (who build the model) met regularly to review the model's parameters and outputs. They asked critical questions: *Are we inadvertently excluding any demographic unfairly? Are the thresholds we set reasonable?* This collaborative oversight ensured that the model's design decisions were not made in a vacuum. Similarly, for Quali-Quant AI, researchers reviewed the types of follow-up questions the AI would generate, ensuring they align with our ethical standards and survey best practices (e.g., not leading or coercive). This pre-launch oversight acts as an ethical gatekeeper.
- **Human-monitored performance and intervention:** Once an AI system is live, we don't put it on autopilot. Our operations and data quality teams continuously monitor the AI's performance metrics and sample outputs. For instance, the data quality and client service teams monitor the rates and profiles of QualityScore™ removals within and across studies. If anything anomalous appears (such as a spike in rejections from a particular source or time), a human investigates. We have tools and dashboards internally that surface AI decisions, precisely so a human can review them. In cases where a judgment call is needed, we err on the side of involving a human. For example, if a participant's score is borderline (near the "poor performer" threshold), we might include their data as a complete but mark it for a manual check - rather than letting the AI alone decide to throw it away. This ensures that participants are not unjustly excluded without a human double-checking borderline cases.
- **Hybrid AI-human workflows:** Many of our processes intentionally combine AI and human steps. For example, our fraud detection and participant cleaning process incorporates human review (PII patterns, open-end responses, etc.) of flagged participants (flagged by AI, client teams, or reconciled by clients). We use AI to scale up what humans can do, not to replace human judgment entirely. On QualityScore™, this looks like: 1) AI flags a behavior that is borderline, 2) a human reviews the flag, 3) appropriate action is taken (removal or marked as good complete). This ensures fairness and reduces false positives. On the analysis side, if an AI text analytics tool summarizes sentiment, our research team may spot-check a few verbatims against the summary. They can override or adjust the summary if needed before it goes into the final report. By structuring workflows this way, we keep humans in control of the final output.
- **Training and empowering staff:** All Dynata staff working with AI systems are trained to understand not just how to use them, but also how to question and correct them. We cultivate a mindset that it's not only allowed but *expected* for team members to intervene if an AI output seems off or potentially problematic. For example, a project manager noticing an odd pattern in data (even after AI cleaning) is encouraged to escalate it. Our culture doesn't treat AI results as

infallible; instead, we treat them as one input to decision-making. This empowerment is crucial – oversight is effective only if the humans involved feel responsible and authorized to act. We also designate specific roles, like a product owner or a data steward, who are accountable for the ethical performance of each AI feature. They serve as points of contact for any concerns and regularly audit the outputs.

- **Governance and compliance oversight:** On a higher level, our Information Security and Data Privacy teams provide oversight to ensure compliance. Any new AI that deals with personal data goes through a privacy impact assessment by our Privacy team. Our General Counsel's office and compliance team are involved in reviewing AI uses, especially in sensitive areas like data linking or activation. This means that, beyond the immediate project team, a second set of eyes in the organization with compliance expertise oversees AI deployments. For instance, if we start using a new third-party AI translation API, our Information Security team will vet it for data handling practices, and our Legal team will confirm it aligns with participant consent terms. We won't proceed unless these human overseers are satisfied.
- **Client transparency and feedback:** In a way, our clients also contribute to oversight. By being transparent (as per Q10) about using AI, we invite clients to ask questions and give feedback on those AI-generated results. If a client ever says, "This insight doesn't look right to us," we take that seriously, investigate whether the AI might have erred, and if so, correct it. This feedback loop ensures we remain accountable to the people ultimately using the insights.

In summary, human oversight is woven into every stage of our AI solutions. We strongly believe AI should augment human expertise, not replace it, especially when it comes to ethical considerations. Our AI systems thus operate under the guidance of humans – much like an autopilot that can handle routine tasks but always has a pilot (our team) ready to take control if needed. This ensures that ethical compliance is continuously maintained.

We can confidently say there's no "rogue AI" at Dynata: every algorithm behaves within bounds set by humans, and there's a real person ready to step in if it strays or if a judgment outside the AI's capability is required.

Data governance and privacy

13. Data quality: How do you assess if the training data used for AI models is accurate, complete, and relevant to the research objectives in the interests of reliable results and as required by some data privacy laws?

Ensuring high-quality training data is a foundational step for us, as it directly affects the reliability of our AI models. We have several practices in place to assess and maintain the accuracy, completeness, and relevance of training data:

- **Leveraging Dynata's first-party data advantages:** Because Dynata has one of the industry's largest first-party panels, we have access to a vast and well-characterized dataset for training our models. All training data we use from our panel comes from real, verified, and consented participants who

participate in surveys. This means our models start off learning from accurate and authentic responses rather than synthetic or scrapped data of unknown quality. For instance, to train QualityScore™, we pulled historical survey data where the ground truth was established (e.g., cases known to be fraudulent vs. genuine). These were labeled using our rigorous data quality checks and human expert judgment. By training on this curated, labeled dataset, we ensure the model learns what "good" and "bad" data look like under real research conditions.

- **Relevance to research objectives:** We are careful to use training data that matches the context of the model's intended use. For example, if we're developing an AI to summarize open-ended responses in product feedback surveys, we train and validate it on similar open-ends (not, say, unrelated social media comments). This domain-specific approach ensures the model isn't learning irrelevant patterns. Our researchers define the objectives and then work with data scientists to select the right training corpus. We also involve subject-matter experts to verify that the chosen data embodies the research constructs we care about (topics, language style, etc.). This makes the AI's learned behavior aligned with what researchers and clients expect to see.
- **Data accuracy checks:** Before using any dataset for training, we run it through quality control. This includes removing or correcting any errors in the data labels, outliers, or inconsistent entries. For panel data, accuracy is high because responses come directly from engaged panelists and our survey scripting often has validation (e.g., preventing impossible values). But we always double-check. For example, if we use response time as a feature, we ensure our data doesn't include any corrupted timestamps. Another angle to accuracy is ensuring the data truth we train on is solid: for QualityScore™, we double-validated the "bad" cases by multiple methods (digital fingerprinting, open-end review, etc.) to be confident those truly were fraudulent or poor-quality participants. We essentially want the training set to be as "truthy" as possible so the model isn't learning from noise or mistakes.
- **Completeness and representativeness:** We strive to cover the breadth of scenarios the AI will encounter. During training data selection, we ask: Is the dataset large and varied enough? For example, when training our AI open-end coding tool, we included responses from various industries (CPG, tech, healthcare, etc.) and question types, so the model can handle the range of content it might see in study production. We also ensure representation of different demographics and geographies in training when relevant, mirroring our global panel diversity. This completeness helps the AI avoid bias and perform robustly. If any subset of data is under-represented (e.g., we have fewer non-English responses in a dataset but the model will encounter them in studies), we take steps to gather more data from that subset or augment what we have (using translation or additional collection) until we're satisfied the training data is sufficiently comprehensive.
- **Ongoing validation of data suitability:** After training, we validate the model on separate test datasets and compare outputs to human judgement or known benchmarks. This not only evaluates the model but also back-checks the training data's adequacy. If a model performs poorly in a certain area, it can signal that our training data wasn't rich or accurate enough in that aspect, prompting us to refine the dataset and re-train the model. For example, when we first built an AI model for sentiment analysis, we noticed it struggled with sarcasm in responses. That indicated our training data didn't have enough examples of that nuance, so we went back and enriched the training set with more such cases (with human-labeled ground truth), which

improved performance. In doing so, we're essentially testing how "complete" the training data was relative to real-world data.

- **Compliance with privacy laws in training data:** This is crucial – laws like GDPR require that personal data used for purposes like model training have a legal basis and that individuals' rights are protected. All our training data usage falls under either consent or legitimate interest that is compatible with the original survey purpose. We ensure that if we use third-party LLMs, the data we send isn't absorbed into their public training sets – we protect training data from leaving our controlled environment.
- **Third-party audit and certification:** To bolster our confidence, we subjected our data practices to external scrutiny. Dynata achieved the Neutronian Quality Index (NQI) certification in 2023 after a comprehensive third-party audit of our data policies, procedures, and datasets. This audit looked across five categories of data quality (including accuracy, representation, and compliance). Passing such an audit is a strong validation that the data feeding our processes (and by extension, our AI training) meets high standards. It's an external check that we're not just marking our own homework. Moreover, Dynata is ISO 20252 certified in the US, UK, and Australia (the international standard for market research), which explicitly covers data quality and data protection. This standard influences how we manage training data (e.g., requiring documented procedures for data handling and validation).

In conclusion, we put a lot of effort into "getting the data right" before it ever trains our AI. We choose relevant and reliable datasets, clean and complete them, and continuously verify that the model outcomes align with research objectives, adjusting training data as necessary. By doing so, we ensure that our AI models are learning from the best possible examples of truth, which in turn produces trustworthy and legally compliant results for our clients.

14. Data lineage: Do you document the origin and processing of training or input data, and are these sources made available?

Yes, Dynata maintains thorough documentation of the data lineage for both the training data used in our AI models and the input data that our systems process. We recognize the importance of data provenance – knowing where data comes from, how it's been handled, and what transformations it undergoes – for transparency, quality control, and compliance purposes.

Here's how we approach it:

- **Documentation of training data origin:** For each AI model or automated solution we develop, we keep records of the datasets used in training. This includes details such as dataset name/description, time period of collection, data sources, any filters applied, and size/scale of the data. In practice, since we often use our own panel data, the origin is typically "Dynata panel responses from X type of surveys (e.g., feedback surveys in 2020, global English responses)" or something of that nature. We note if the data was internally generated (which it usually is) or if any external data sources were integrated. For example, if we incorporated an external linguistic

dataset to improve an NLP model (e.g., a public sentiment lexicon or an open-source language model which was pre-trained on Wikipedia), we document that as part of the model's lineage. We also log any pre-processing steps (e.g., "text responses were lowercased and tokenized" or "outliers beyond 3 standard deviations were removed from response time feature.")

- **Data processing pipelines:** We employ versioned pipelines for data handling, meaning each step of processing (from raw data to training-ready data) is recorded and can be reproduced. This is often done in our data engineering workflows and annotated in our internal wikis or data catalogs. For instance, we can trace that for QualityScore™ version 1.2, the training data came from a certain data lake table (with a reference ID), was processed by script version 1.2.5 which did XYZ transformations and resulted in a training file on a certain date. This level of detail ensures that we know exactly what went into the model. It's not typically client-facing information, but it's available internally for audits or if clients inquire deeply.
- **Availability of source information:** While we might not publish all training data sources publicly (due to volume and confidentiality), we are transparent with clients about the nature of the data sources. If a client asks, we can tell them, for example, "The AI that coded your open-ends was trained on thousands of prior open-ended responses from similar surveys in our database, primarily in the retail sector, collected over the past two years." We can also confirm if any third-party data was used. In most cases, the sources are first-party (our panel), which we openly say is the case. If external benchmarks or libraries are used, we can cite those too (e.g., "we fine-tuned the model on the OpenAI GPT-3 base model" or "we used a public sentiment dataset from [source] as additional training").
- **Input data lineage in projects:** For the data we collect and analyze for clients (the input data to our AI during a project), we also track its journey. Our systems log, for each participant, metadata such as when they were invited, what sample source they came from, any data merges or enrichments applied (with Dynata's Connected Data or third-party appends, for example), and how the data moves into analysis. If data goes through an AI, that is part of its processing history. We make sure that if needed, we can provide a methodological trace. For example, if using Virtual Audiences (our AI-generated data tool) together with first party respondent data, we document what external data fueled the AI generation and how it was combined. For Dynata's AI data augmentations, we document sources - if we pulled social media trends or syndicated data as part of an AI model's input, we list those sources (e.g., "social data from [platform], time range Q1 2025") in our methodology notes.

- **Client access to lineage Info:** We are willing to share data origin information with clients to the extent it's useful and not proprietary. For instance, if a client is evaluating our AI offering, they might ask: "Where does the data come from that trains your models?" We can confidently answer that it's from our robust first-party repository with broad coverage and explain any other sources. We also share that our data is permission-based and research-grade – essentially communicating lineage in terms of being ethically sourced from our own panel. If a formal documentation is needed (some clients have diligence questionnaires about AI), we can provide written responses on data lineage and governance.
- **Auditable trails:** In compliance with standards like ISO 27001/27701 (info security and privacy) and ISO 20252 (market research), we maintain auditable records of data flows. That means if an

auditor or regulator ever needed to trace data from collection to analysis, we have the internal logs and documentation to do so. For example, GDPR's principles of accountability and data minimization mean we should know what data we used and ensure we didn't use more than necessary. Our lineage docs help demonstrate that; for example, we can show we only used survey response data (no extra personal data) to train a model, and we can pinpoint when and how it was anonymized.

In summary, we do document the origin and processing of the data powering our AI, and while the full technical lineage might be mostly for internal governance, we extract and share the relevant source information with clients and stakeholders. This transparency gives clients peace of mind that our models are trained on appropriate, high-quality data and that we're not pulling in mysterious or questionable sources. It's part of building trust – clients know that *Dynata's AI is fueled by Dynata's data* (and clearly identified supplementary sources when applicable), and all of that is managed under strict data governance protocols.

15. Please provide the link to your privacy notice (sometimes referred to as a privacy policy). If your company uses different privacy notices for different products or services, please provide an example relevant to the products or services covered in your response.

Dynata's Privacy Policy, which covers our data practices for all products and services (including AI-based research services), is available on our website. You can find our comprehensive privacy notice at the following link: Privacy Policy - Dynata

This privacy policy outlines how we collect, use, store, and protect personal data across our platforms. It applies to panelists (survey participants), corporate website users, and clients, and it addresses specifics such as data subject rights, data retention, security measures, and our compliance with frameworks like GDPR and CCPA.

For convenience, here's a brief summary of what the Dynata Privacy Policy covers:

- **Description of the business and data collected:** It describes Dynata's role in providing sampling solutions and the types of personal data we handle (e.g., contact information, demographic info, responses to surveys, etc.).
- **Use of personal data:** The policy explains the purposes for which we use personal data (e.g., to invite and manage survey participation, to perform research analysis, to improve our services).
- **Legal basis:** It outlines the legal grounds for processing data (such as consent for panelists, legitimate interest for certain internal analytics, compliance with legal obligations, etc.).
- **Sharing of data:** The notice details with whom we may share data (such as clients receiving survey results in anonymized form, or service providers under strict agreements).

- **Data subject rights:** It informs individuals of their rights to access, correct, or delete their data, and how to exercise those rights.
- **International transfers:** Dynata is a global company, so the policy covers how we lawfully transfer data across borders (e.g., standard contractual clauses, Privacy Shield legacy commitments).
- **Data security and retention:** It describes the safeguards we have in place to protect data (technical, administrative, physical security measures) and how long we keep personal data.
- **Contact Information:** It provides contact details for our privacy office and how to lodge inquiries or complaints.

All of Dynata's research offerings, including the AI-enhanced services discussed in this document, are governed by this central privacy policy. We do not maintain separate privacy notices for individual AI products; instead, we ensure our core privacy policy is broad and clear enough to encompass new technologies like AI. In some cases, we might have supplemental notices or consent language for specific use cases – for example, if we launch a new mobile app for surveys or a specific panel community, there might be an in-app privacy summary – but ultimately all point back to the main Dynata Privacy Policy linked above.

For any client or participant seeking details on how Dynata upholds privacy in the context of AI and all other research activities, the Privacy Policy is the definitive resource. It's also accessible via the footer of our website (labeled "Privacy Policy"), ensuring transparency to anyone who wants to review our practices in depth.

16. What steps do you take to comply with data protection laws and implement measures to protect the privacy of research participants? Have you evaluated any risks to the individual as required by privacy legislation and ensured you have obtained consent for data processing where necessary or have another legal basis?

Dynata takes compliance with data protection laws and the privacy of research participants extremely seriously. As a global company, we adhere to regulations like the GDPR in the EU, CCPA/CPRA in California, and similar laws in other jurisdictions. Here are the key steps and measures we have in place:

- **Privacy by design and default:** We incorporate privacy considerations at the design stage of our research services (including AI-driven ones). This means we only collect data that is necessary for research purposes and use it in ways that are compatible with those purposes. For example, if our AI needs certain data points, we ensure those are justified, and we don't pull in extra personal data that is not explicitly needed. Systems are configured to default to non-identifiable data wherever possible. Survey results are analyzed in aggregate, and our tools (AI included) focus on aggregated patterns, not on identifying specific individuals.

- **Consent and legal basis:** We always ensure we have a proper legal basis for processing participant data:
    - For our panel members, we obtain informed consent at the time of enrollment and before any data collection. They agree to our Panelist Terms & Conditions and Privacy Policy, which outlines what data will be collected and how it will be used for research. For sensitive data or new uses, we seek explicit consent as required. For example, if we were to collect health information or geolocation in a study, we'd ask permission specifically for that.
    - We also operate under legitimate interest for certain processing that is integral to research quality, such as fraud detection and data quality improvements, ensuring these interests are not overridden by participants' rights (and indeed serve to protect the overall research integrity for all participants). In all cases, participants can opt out, and we honor such choices promptly.
    - If we ever process data from EU residents, we adhere to GDPR requirements such as being able to demonstrate consent or legitimate interest assessments. Our records of processing (RoPA) document each processing purpose and its legal basis.
- **Data protection impact assessments (DPIAs):** We evaluate risks to individuals through formal assessments. When we introduce a new technology or process that could impact personal data (for instance, implementing a new AI-based tool involving personal data), our Data Privacy team conducts a DPIA or similar risk analysis. This involves identifying potential risks (re-identification, bias, security breaches, etc.) and ensuring we have mitigations in place. For example, before rolling out an AI chatbot for qualitative research, we would assess risks such as participants being able to inadvertently share personal data with the bot and how to prevent that, or whether the bot's logic could result in intrusive questions. By evaluating these risks, we implement safeguards (e.g., filtering out PII that a participant might type, restricting question types). We maintain documentation of these assessments as required by law and update them as the service evolves.
- **Participant rights and control:** We have processes to enable participants to exercise their rights under privacy laws. If a panelist requests access to their data, or correct it, or delete it, we have an infrastructure to respond within mandated timeframes. For deletion requests (right to be forgotten), we scrub their personal data from our systems and also ensure that none of our AI models contain identifiable data of that person (note: our AI models typically do not store personal identifiers, but we include model data in our deletion workflows if ever applicable). We also provide easy opt-out mechanisms from research at multiple touchpoints (in survey invites, within surveys, on our panelist portal, etc.).
- **Security measures:** Protecting privacy goes hand-in-hand with strong data security. We have a comprehensive Information Security program that aligns with industry standards (we maintain SOC 2 Type II compliance and follow frameworks akin to ISO 27001 and the NIST Cybersecurity Framework). We implement measures such as encryption of data in transit and at rest, access controls (only authorized personnel can access personal data, on a need-to-know basis), network security defenses, and regular security audits and penetration tests. Our

systems are monitored and designed to prevent, detect, and respond to any unauthorized access or data leakage. We also require our vendors who might process data on our behalf (such as cloud service providers or third-party AI services) to have strong security and privacy commitments, reflected in our contracts (DPAs, etc.).

- **Minimization and anonymization:** We practice data minimization – only collecting what we need for the research. Where possible, we use anonymized or pseudonymized data in our analytics. For instance, when our AI analyzes survey responses, the data typically doesn't include names or contact info – it's just participant ID codes with their answers. Those ID codes are separate from personal identifiers. In deliverables to clients, data is aggregated or anonymized (clients see responses but not, say, the email or full names of participants). If we supply unit-level data (microdata) to clients for their own analysis, it's without direct identifiers and under strict usage terms. This ensures participants' privacy is respected even in what clients receive.
- **Compliance team and training:** Dynata has a dedicated data protection officer (DPO) and a privacy team who oversees adherence to regulations. We also have a legal team that keeps track of evolving laws worldwide. We provide regular training to our employees on data privacy and security, as well as specific training for those handling personal data (like reminding them of protocols for handling subject access requests, or how to spot a potential social engineering attempt). Our company culture emphasizes respecting participant privacy as a core value, not just a legal duty.
- **Audits and certifications:** We often subject ourselves to third-party audits to verify compliance. For example, we hold certifications like the Neutronian NQI (which, as mentioned, includes a privacy compliance component). We also comply with the Insights Association's Privacy Code and ESOMAR's guidelines. We've publicly committed to principles like the ESOMAR Data Protection Checklist, which further assures that our practices are up to the mark. Internally, we perform routine audits of data processes. If we find any gaps, we remediate them swiftly.
- **Consent for AI and new uses:** If we introduce AI in a way that changes how participant data is used beyond their original expectations, we are careful to handle consent. For example, if we were to repurpose some survey responses to train a new AI model that has broader uses, we'd consider whether the original consent covers that. If not, we might anonymize the data or seek additional consent. However, generally, our panelist consent includes the improvement of research methodologies, which covers internal AI training. And notably, we ensure that data provided to AI providers cannot be used to train their models – this is both a privacy and proprietary measure. So, our participants' data isn't inadvertently "leaked" into some external AI.

In essence, Dynata's approach is multifaceted: **get consent, be transparent, enforce security, minimize data, and rigorously respect all individual rights.** We maintain documentation (Data Processing Agreements, DPIAs, etc.) as required by laws and are always prepared to demonstrate our compliance with clients or regulators. By weaving privacy protection into our daily operations and technological design, we aim to uphold the trust of our research participants – they are, after all, the lifeblood of our business, and respecting their privacy is paramount to maintaining that trust.

17. What steps do you follow to ensure AI systems are resilient to adversarial attacks, noise, and other potential disruptions? Which information security frameworks and standards do you use?

Dynata takes a proactive approach to make sure our AI systems (and the infrastructure they run on) are resilient against adversarial attacks, noise, and disruptions. We recognize that as we integrate AI, we must also anticipate intentional attacks (like someone trying to trick the system) and unintentional disruptions (like data noise or system failures). Here are the key steps and measures we have in place:

- **Robust model training and testing:** From the outset, we train our models to be robust to noise in data. For example, when developing QualityScore™, we included a variety of real-world "messy" data in training – participants who exhibit odd but not fraudulent behavior – so the model learns not to overreact to benign anomalies. We also test our AI with simulated adversarial inputs. In the context of survey data, an adversarial attack could be a participant (or bot) deliberately trying to circumvent our quality checks (perhaps by randomizing answers in a clever way, or by mimicking patterns of good participants). We stay ahead by continually updating our detection algorithms to new fraud patterns. Dynata's report on generative AI and fraud highlights how we adjusted our model to catch AI-generated responses by focusing on passive behavior signals. That's an example of thwarting an adversarial tactic (fraudsters using AI) by enhancing our system's resilience.
- **Adversarial testing:** We sometimes conduct red-team style exercises on our AI systems. Internally, we might assign someone to try to "beat" the AI – for instance, attempt to create a survey bot that slips past QualityScore™'s radar – and then use those findings to strengthen the model. In cases of text analysis AI, we examine how the model handles inputs with typos, slang, or deliberately misleading text to ensure it doesn't break or produce inaccurate results. If we find vulnerabilities (such as a certain pattern of gibberish that wasn't caught), we patch the system rules or retrain the model accordingly.
- **Multi-layered security controls:** On the systems side, our AI runs on platforms that are secured under Dynata's extensive Information Security program. We comply with and certify to high security standards – for example, we maintain SOC 2 Type II compliance, which means our controls for security, availability, and integrity are audited annually. We also align with frameworks like ISO 27001 and the NIST Cybersecurity Framework for best practices in risk management. These frameworks ensure that we have controls like intrusion detection, incident response plans, regular penetration testing, and access controls. If there's a potential disruption like a network outage or a DDoS attack, we have measures in place (redundancies, traffic filtering, etc.) to keep our systems running or restore them quickly.
- **Data encryption and integrity:** We protect data in transit and at rest through strong encryption (TLS for data in motion, encryption at rest on databases and storage). This prevents attackers from tampering with or injecting noise into data streams feeding our AI. We also use checksums and validation steps to ensure data hasn't been corrupted. For example, when moving a dataset for

training, we verify its integrity. These practices ensure that what the AI ingests and outputs isn't secretly modified by an adversary or random corruption.

- **Noise handling in AI:** We design our AI algorithms to tolerate a degree of noise. For example, our models often use aggregated behavior scores or multiple indicators, so that one noisy signal doesn't throw off the result. If one input feature is unreliable, others compensate. In natural language tasks, we may deploy filters to clean text (to remove garbled characters, etc.) before analysis. This makes our AI more fault-tolerant. Additionally, we cap the influence of any single factor in some models – for example, in QualityScore™, rather than one metric deciding the ultimate fate of the participant, it's a combination of 175+ factors. This diversity means random noise in one factor is unlikely to cause a big disruption.
- **Real-time monitoring and alerts:** We have monitoring in place for our AI systems and their outputs. If an AI service starts behaving oddly – suddenly flagging 30% of participants as fraud when it is normally 5% -- our monitoring will catch that anomaly and alert our engineers. This could indicate either an adversarial onslaught (perhaps a flood of bad actors, or someone found a loophole) or a system glitch. In either case, we can react swiftly: investigate the root cause, block malicious actors (if that is the issue), or roll back a model update if it introduced instability. Our incident response process, honed under frameworks like NIST, guides how we isolate issues and recover.
- **Secure development lifecycle:** Our AI models and software are developed following secure coding practices. We threat-model new features (including AI ones) to foresee how they might be abused. For example, we ensure that any API endpoints exposing AI functionalities are authenticated and rate-limited to prevent abuse. We also sanitize inputs rigorously – an adversary might try SQL injection or code injection via input fields in a survey, but our systems neutralize such attempts.
- **Third-party components vetting:** We often integrate third-party AI libraries or services (like open-source ML frameworks or cloud AI APIs). We vet these components for security. We keep them updated to patch any known vulnerabilities. For cloud AI (if used), we rely on providers that meet high security standards, and we configure them to our needs (e.g., using private endpoints, not exposing data to public training). Essentially, we extend our security posture to any external piece we use.
- **Resilience to outages:** Beyond attacks, disruptions can be mundane (hardware failure, etc.). We host our systems in robust, cloud-based environments that have redundancy. Data is backed up. For AI model serving, we often have fallback models or rules if the main model fails. For example, if an AI coding service is unavailable, we might have a rules-based coding system as backup or at least queue the task until service resumes. This prevents a single point of failure.

By following these steps and adhering to industry-leading security standards, we make our AI systems **resilient and secure**. A concrete example tying these together: When ChatGPT and similar GenAI emerged, we anticipated that survey fraud might evolve (adversaries using GenAI to create human-like answers). We revised our QualityScore™ algorithm in late 2022 to put more emphasis on passive behavioral cues that GenAI cannot easily fake. This preemptive step ensured resilience against that new threat. Similarly, our information security program ensures that even if someone

tried to sabotage or manipulate our data (an adversarial attack on our infrastructure), they would face multiple barriers and detection at multiple levels.

In summary, we blend sound engineering, continuous monitoring, and adherence to top security frameworks (SOC 2, ISO 27001, etc.) to keep our AI systems robust against attacks and disruptions. Our clients and participants can trust that we are vigilant about safeguarding the integrity of the data and the insights that depend on our AI.

18. Data ownership: Do you clearly define and communicate the ownership of data, including intellectual property rights and usage permissions?

Yes, Dynata clearly defines and communicates data ownership and usage rights in our contracts, policies, and client engagements. We understand that data ownership can be a complex topic – involving the client's data, Dynata's data, and any output or insights generated – so we make sure all parties know who owns what and how it can be used.

Here is how we handle it:

- **Client data:** When a client brings their own data to a project (for example, a customer list for sampling or proprietary survey questions/results), the client retains ownership of that data. Our contracts (Master Services Agreements or Data Processing Agreements) explicitly state that any client-provided data remains the intellectual property of the client. We act as a service provider/data processor for that information, using it only for the purposes the client has specified. We do not reuse or resell client's proprietary data in any way unless given explicit permission. This is clearly communicated during project scoping and in our terms – essentially, *"your data stays yours."*
- **Dynata's data (panel and derived data):** Dynata owns and maintains one of the largest first-party panels, and the data we collect from our panelists (their profile attributes, responses, etc.) is owned by Dynata (while, of course, subject to the privacy rights of the individuals). We are essentially the stewards of that data and have intellectual property rights over the aggregated panel dataset, our methodologies, and any proprietary metrics (such as our QualityScore™ algorithm outputs). That said, when a client commissions research, they typically obtain rights to use the *results* of that research. We communicate that the participant-level data and insights delivered from a study are for the client's use, but the underlying panel and any methodologies remain Dynata's property. For example, the fact that we have demographic info on a participant is part of our panel asset, but the output dataset of survey responses for the client's project is theirs to use under the agreement.
- **Survey responses and insights:** Generally, the client owns the research results (survey responses, reports, etc.) that we deliver to them, while Dynata retains ownership of any of our templates, tools, or analytics methods used to produce those results. This means the client can use the data and insights for their internal purposes, decision-making, etc., as agreed. We typically restrict clients from re-contacting participants directly (since panelists are our asset), unless a specific

arrangement is made through us. But aside from such restrictions (which are in place to protect participant privacy and Dynata's panel relationship), clients have broad usage rights to the study data. We make this clear in our agreements: for example, "Client shall have a license to use the Deliverables (survey findings, etc.) for its business purposes." If they want to publish the findings, that is usually allowed, sometimes with attribution to Dynata for data collection if appropriate.

- **AI-generated content or models:** In cases where our AI might generate content (such as a summary or a "synthetic" data set via Virtual Audiences), we address ownership as well. Typically, if it is part of a client project, the output is treated as part of the deliverable to the client – hence their property to use. For example, if Quali-Quant AI produces a thematic report of reasons why consumers like a product, that report is the client's to use, just as if a human analyst wrote it. However, the underlying AI model and algorithms remain Dynata's intellectual property. We might explicitly state that any improvements to our AI models from doing the project do not confer ownership to the client (so they cannot claim rights over our algorithms), while conversely we cannot claim ownership over the client's specific data trends. Everyone retains ownership of their "secret sauce."
- **Communication in proposals and terms:** We communicate these ownership details in plain language as part of our client onboarding or proposal process. Many clients are accustomed to market research norms where, by default, they own the survey data they commissioned. We confirm that understanding and put it in writing. We also address usage permissions: for example, we often request that if a client publishes the research results externally (such as in a press release or white paper), they consult with Dynata as to whether they need to credit us as the data provider (this is more about usage rights and recognition than ownership per se). If we want to use a project as a case study, we ask the client's permission, respecting that the data/results belong to them.
- **Panelist data ownership:** It is worth noting, from the participant perspective, when they join our panel, they agree that their responses can be used in aggregate for research. They "own" their personal data in the sense they have rights to it (access, deletion, etc.), but they transfer to Dynata the permission to use their responses for research purposes. We clearly explain this in our panelist terms/privacy policy. This ensures we have the right to compile their answers into the datasets we provide to clients. However, we do not claim ownership of their answers or any information they give us – panelists can request deletion, and we honor it (meaning we do not use their data going forward). So, in that micro-sense, data ownership is also addressed vis-à-vis the participant.

In summary, our approach is to make ownership transparent and logical: Clients own their proprietary inputs and the specific outputs of their commissioned research; Dynata owns the tools, methods, and panel data that enable the research; and participants retain ownership of their personal data rights even as we use their responses for authorized purposes. We make sure these points are clearly written in contracts and often discuss them in kickoff meetings to ensure mutual understanding. This clarity prevents misunderstandings and allows all parties to use the data confidently within agreed boundaries.

(As this is a general explanation, if you need reference to specific contract clauses or policy excerpts, we can provide those. But the above captures our standard practice.)

19. Data sovereignty: Do you restrict what can be done with the data?

Yes, Dynata imposes certain restrictions on what can be done with data, in line with legal requirements, ethical standards, and our contractual commitments. These restrictions are in place to protect research participants, comply with regulations (which often involve data localization or purpose limitation), and safeguard the integrity of the data.

Here is how we address data sovereignty and usage restrictions:

- **Purpose limitation and usage agreements:** We contractually ensure that data is used only for the agreed research purposes. For clients, this means when we deliver survey data or insights, it is to be used for market research and insight generation purposes as stated. We prohibit using the data for purposes outside of research (for example, using participant-level data for marketing to those participants, unless explicitly allowed via a recontact agreement that respects privacy laws). This is often spelled out in our agreements or in terms of service. In other words, the data we provide cannot be repurposed arbitrarily – there are boundaries. For example, if we deliver anonymized survey responses, the client is not allowed to attempt to re-identify participants or append that data to other databases to profile individuals. These kinds of restrictions are rooted in both ethics and compliance.
- **Geographical restrictions (data localization):** In terms of data sovereignty in the geographic sense, we do accommodate and adhere to laws that restrict data movement. If a country or region requires that personal data stays within its borders (or has strict conditions for transfer), we abide by that. For example, for EU personal data, we utilize approved transfer mechanisms (SCCs, etc.) when moving data to the US, or we process data on EU servers if required by a client or regulation. In some cases, clients might request that all data for their project be stored/processed in a certain jurisdiction for compliance – we have the infrastructure to do that or have partners in a region, and we then restrict data flows accordingly. This means we won't, for example, take data collected in China and freely move it to US systems without clearance; we have protocols to comply with data sovereignty laws like China's PIPL or Russia's localization laws, usually by keeping data local or ensuring proper legal frameworks for transfer.
- **System and role-based restrictions:** Within our platform, we restrict who can do what with data. Only authorized personnel can access raw personal data, and even they have role-based limits (a data processor might see IDs and responses, whereas an analyst might only see aggregated data). These internal controls prevent misuse. On the client side, if we provide a dashboard or portal, we may impose technical restrictions to prevent data exports beyond a certain level of detail if that was the agreement. We can also watermark or audit trail the data usage. Essentially, the systems enforce the policies where possible.
- **Tool-specific guardrails:** Each of our AI tools or data products has built-in restrictions by design, as part of our commitment to proper data use. For example, Virtual Audiences (our AI-simulation tool) inherently restricts input/output to acceptable content. In Quali-Quant AI, the tool is configured to only generate probes that stay within the scope of the survey topic (it cannot

suddenly ask for personal data or off-topic information). This is a subtle form of data use restriction -- we restrict what the AI can solicit or output. Moreover, QualityScore™ automatically restricts low-quality or non-compliant data from reaching the client; for example, if a participant were to enter a vulgar or privacy-violating comment, our systems might strip that out. These content moderation steps ensure that certain kinds of data (like hate speech, PII in open text, etc.) are restricted from propagation.

- **No unauthorized sharing or selling:** We do not allow data that is collected for one purpose to be shared with third parties for other purposes without permission. For example, we do not take a client's research data and sell it elsewhere. And we do not allow ourselves or our clients to use participant data for direct marketing (unless the participant explicitly opted in for that in a special case). Participants are usually anonymous to clients, and clients get no right to contact them directly unless a re-contact study is arranged via Dynata. This preserves participants' control and aligns with data sovereignty principles (participants "own" their presence and we protect that).
- **Retaining sovereignty over panel data:** As Dynata, we consider our panel data (and any data collected from it) as something we must govern responsibly. We restrict what we share (externally and internally). For example, if a client wanted raw identifying information of our panelists, we would not provide it because that goes against our commitments to panelists. We also restrict how data flows out of our environment. Clients typically receive either aggregate reports or anonymized microdata. We ensure that no hidden personal identifiers slip through.
- **Compliance with consent and opt-outs:** If a participant has said "don't use my data for X," we absolutely restrict that. For example, if someone opts out of having their data used in predictive modeling, we tag and remove them from any such data sets. Through consent management, we implement what effectively are restrictions in data usage per individual's wishes or per law (such as not using someone's data after they withdraw consent).

In communication with clients, we clarify these restrictions. We might say, for example, "This dataset is for your internal research use only. It should not be redistributed or used to identify or target individuals." These terms typically appear in our usage guidelines.

By clearly restricting data usage in these ways, we protect all parties: participants' privacy is safeguarded, clients remain in compliance and ethical territory, and Dynata preserves the trust in our data asset. Every Dynata tool or dataset comes with "rules of the road," and we make those known up front.

20. Ownership: Are you clear about who owns the output?

Absolutely. We make it explicit who owns the outputs of our services and analyses, so there is no ambiguity for our clients or partners. In general, the client who commissions the research owns the output of that research, and we are clear about this in our agreements and communications.

Here is how we handle output ownership:

- **Client ownership of deliverables:** The reports, data tables, dashboards, and any analytic results we deliver to a client become the client's property once payment and contract terms are fulfilled. This is industry standard and we adhere to it. For example, if we provide a market segmentation analysis or a set of cross-tabulated survey data, the client can use those outputs freely within their organization – these are theirs to keep. We typically state this in contracts in language such as: "All deliverables prepared for Client shall be owned by Client." The only caveat is that we retain the right to use the knowledge generically (i.e., we do not give away our methods, and we cannot disclose the client's specific findings to others because it is theirs). But as far as the output data and insights, the client has full ownership and usage rights.
- **Dynata's ownership of methods/tools:** We ensure clarity that while the findings belong to the client, Dynata retains ownership of the underlying tools, models, and intellectual property that produced those findings. For example, if we use our AI to generate a summary, the summary is the client's, but the AI software and algorithms remain Dynata's intellectual property. This distinction is usually clarified in our terms. So, the client could not extract our AI model from our system and claim it as theirs, but the results it produced for their project are theirs. This separation is communicated, so clients understand they have the fruits of the analysis, but not ownership of, for example, the QualityScore™ algorithm itself (which is a Dynata asset).
- **Panel and participant data:** Another aspect is ownership of the panelist responses vs. output data. As discussed earlier, Dynata owns the panel data overall, but when a study is run, the client effectively owns the compiled responses for that study (minus personal identifiers). To put it simply: the client owns the survey dataset (responses) and the insights derived from those responses, and we confirm that in writing. We sometimes phrase it as the client having a perpetual, worldwide license to use the study data and results for any lawful purpose (since they paid for its collection). It is effectively ownership for practical purposes.
- **Outputs in AI context:** If AI was used to create an output (such as a synthesized report or a predictive model tailored for the client), we also clarify ownership. For example, if we generate a predictive scoring model for a client's customer churn (using our AI on their data), typically that model (and the scores for their customers) would be considered an output of the project and thus the client's asset. We would not give that same model to another client, as it's developed from their proprietary data. We clarify that in such custom AI projects, the client owns the specific model instance or output, whereas we may retain the general methodology.
- **Mutual understanding and no surprises:** At project kickoff or proposal stage, if there's any unusual aspect of output ownership, we address it. However, in most cases, it's straightforward: the client gets full ownership of the deliverables. We even often send raw data files if requested, acknowledging they are theirs to analyze further as they wish. And if we ever want to use the output (say, a particularly great insight or chart) for a Dynata case study or marketing, we ask for the client's permission – indicating we recognize it is their output, not ours, once delivered.
- **Contractual wording:** To ensure clarity, our contracts might include something like, "Client shall own all results, reports, or other deliverables provided by Dynata as part of the services. Dynata retains ownership of any Dynata proprietary materials or tools used in performing the services." This clear delineation prevents any misunderstanding.

In summary, we are clear and upfront that the client owns the outputs of our work, just as Dynata owns its platform and methods. This mutual respect of ownership rights is foundational to our client relationships. Clients invest in our services to get insights and data they can use and control, and we make sure they indeed have that ownership once the project is completed.

All data and insights generated by Dynata on behalf of a client are handed over as the client's property, to be used at their discretion. We stand by this principle and have it well-documented, so there is no confusion about who owns the results of our AI-powered research services.

dynata™