

Panel quality and fraud controls

PollStream delivers insights from high grade proprietary data from real humans, not replaceable with AI guesswork.

Bad data is cheap to produce and expensive to act on. Survey fraud is organised, automated and increasingly written by machines. This note sets out how we keep it out, what we tell you when some gets through, and why we will not fill the gap with data that was never collected from anyone.

WHY THIS MATTERS

A fraudulent respondent is not a rounding error. Fraud clusters. The same operator completes many surveys, the same script fills many open ends, and the risk compounds as sample size grows. Controls applied after fieldwork cannot rescue a study contaminated at entry, so ours run at every stage.

01 · BEFORE FIELDWORK: SOURCE GOVERNANCE

Every respondent source is vetted against defined quality standards before it reaches the platform, and monitored continuously after it does. We sample across multiple sources rather than one, which reduces the systematic bias any single recruitment channel carries. Sources that fail on quality are removed, not managed.

02 · AT ENTRY: IDENTITY AND DEVICE VERIFICATION

Device fingerprinting, location validation and identity checks run before a respondent reaches a live survey. VPN and proxy use, duplicate entry, spoofed device environments and known fraud signatures are rejected at the door. A respondent stopped here never enters the dataset, so there is nothing to clean out later.

03 · IN SURVEY: ATTENTION, CONSISTENCY AND AI GENERATED ANSWERS

Attention traps, speeder thresholds and straight-lining flags are applied throughout the survey. Open ended answers are reviewed in real time for relevance, internal consistency and AI generated language. Behaviour is assessed across the whole session, because a respondent who passes the first screen can still fail the fifth.

04 · AFTER FIELDWORK: REMOVAL AND DISCLOSURE

Completes that fail post-fieldwork checks are removed before analysis, not after a client queries a result. Every dataset is published with its sample size, fieldwork dates, margin of error and sample composition, so the basis of a figure travels with the figure.

REAL RESPONDENTS, NOT SIMULATED ONES

Every figure PollStream publishes came from a person who answered a question on a known date. A growing share of what is sold as public opinion does not clear that bar.

Ask a general purpose AI model what the public thinks about an issue and it will answer fluently, drawing on text absorbed during training, with no sample, no fieldwork date and no way to check. That is a plausible sentence, not a measurement. Synthetic sample goes further, replacing respondents with model generated personas. It reproduces the distribution of its training data rather than the distribution of the population, regresses towards the average view, understates minority and low salience opinion, and carries no provenance. There is no respondent to verify, no consent trail and no date attached to the answer.

The harder problem is verification. Synthetic sample is validated by showing it can reproduce surveys that have already been run, which means it can only be trusted where the answer is already known. For the question nobody has asked yet, there is nothing to check it against. Research exists to tell you what you did not expect, and a model trained on the past cannot do that.

We use AI heavily, for analysis, monitoring and summarisation. We do not use it to invent respondents. PollStream analyses data collected from real people and shows its working. Claims about public opinion are checked against data we have collected, not against a model's recollection.

WHAT WE WILL NOT CLAIM

No panel is fraud free, and any provider who says otherwise is selling something. Detection rates move because fraud adapts. Our commitment is disclosure: if a study is affected, the client is told before the results are used, not after. Quality metrics are reported with every dataset, not just claimed.

STANDARDS AND ESCALATION

We operate to the MRS Code of Conduct and the ESOMAR and AAPOR codes, and hold WAPOR corporate membership. ISO 20252 is in preparation for Q3 2026. Any concern about data quality can be raised at integrity@pollstream.ai and will receive a written response.

Every number in a PollStream dataset can be traced to a person who answered a question on a date we can name. That is the standard we hold, and the one we would ask you to hold any provider to.