

Can you measure usefulness without watching people?

Abstract. A private-agent company cannot honestly use conventional behavioral telemetry as its north-star instrument. The open problem is to estimate useful daily agents while treating the operator as an adversary and making individual behavior structurally difficult to observe.

THE PROBLEM

Can receipt-derived, edge-local sketches with calibrated noise and secure aggregation produce a useful population estimate without reconstructing a person’s behavior or a trajectory over time?

HARD CONSTRAINTS

- 01

The operator, network metadata, and reporting timing are part of the threat model.
- 02

Usefulness must be outcome-anchored, not a count of taps or impressions.
- 03

Privacy composition, bias, battery, bytes, and inflation resistance are all first-class costs.
- 04

A proposal must publish error bars and unsolved failure modes before it is a KPI.

A FIRST CONTRIBUTION

Implement a toy, auditable simulation: receipt predicate → local randomized response → batched shuffle → secure sum → population estimate with published error bounds.

BUILD IN THE OPEN

github.com/hushh-labs/hushh-research
hushh.ai/discord

