

Rebuilding Ratios: AI Distortion in Value Measurement

Revision record

Canonical version: jamesemurphy.com/ai-economics/measurement · SSRN: ssrn.com/abstract=7262118 · doi:10.2139/ssrn.7262118

Substantive changes since the first version posted to SSRN on 10 August 2026, newest first. Typographic, formatting, and version labeling changes are not recorded.

v1.2 · 25 August 2026

Section	Kind	Change
1	Addition	A projection framing opening the paper: every ratio renders some dimensions of an economic object in full and collapses the rest, which is what a measure is for, and it misleads only when the collapse goes undisclosed and the dimensions move. Adapted with attribution from an image offered for the relationship between two adjacent practices.
3.2	Addition	Production side sorting: requests routed to different models by task type and by requester mean two units bearing the same label were produced differently, so the denominator holds still while the numerator is drawn from a mixture the ratio does not disclose.
3.3	Addition	Drift beneath the version label: an endpoint's effective identity can change without a version bump, so no unit event exists to log. Endpoint level evidence shows the same nominal model differing by up to 12.5 points in accuracy and by a factor of five in cost per correct answer across providers, which makes the economically stable unit the endpoint rather than the model.
3.3	Addition	Rate non-uniqueness: committed capacity and on demand consumption of the same model, with traffic overflowing silently between them, mean cost per outcome is not single valued for identical outcomes. A per request figure derived from a period level blend is an allocation rather than an observation. The structure recurs wherever a resource is priced by time rather than by what is drawn from it.
3.4	Addition	The transitional form of trade unit migration reported in practice: the seat fee has become a floor rather than a price, with a variable component driving total spend.
4.1	Addition	Comparators carry unstated workload assumptions, and the identity of a purchase can be stable while its wrapper, effective price, and unit of account are not.
4.2	Addition	The case where no counterfactual exists in principle, the emerging practice of substituting a time bound hypothesis for it, and the hazard that criteria set by the sponsor inherit the structure section 7.1 warns about.

Section	Kind	Change
6.1	Addition	Where a change is not announced the reference must be measured rather than declared, with periodic fingerprinting as the instrument, and two limits stated: consensus comparison is blind to industry wide change, and a closed model served only by its owner has no external reference at all.
6.3	Addition	Attribution coverage as the unstated denominator: a ratio computed over an unknown share of consumption and presented as though computed over all of it.
6.4	Addition	A model judged outcome inherits the stability of its judge. Separately, the quality floor: a cost per outcome ratio is defined only above the bar below which output is worthless at any price, so a comparison spanning the floor is not a comparison, and a firm level intensity figure aggregates incommensurable denominators.
6.6	Addition	A seventh integrity question on whether the rate is single valued for the unit.
7.2	Addition	A scope condition: these ratios detect waste, which binds under abundant supply, and a firm that improves its cost per unit while losing access to the units it needs has optimised a number that was no longer the operative one.
7.4	Addition	A venue now exists for the reference distribution problem, which converts the custody question into the operative one.
3	Figure	The four layer architecture is now a diagram rather than a table, with connectors showing that a single mechanism reaches several matching failures and a single failure arises from several mechanisms.
3.2	Figure	A two panel figure showing one population with a single average, then the same population routed into two pools with two correct averages that straddle the original, alongside a worked illustration of two support teams measured the same way.
1	Addition	A vocabulary rule fixing base, population, unit, and reference to distinct jobs.
5.3	Addition	The incentive asymmetry: a distortion that embarrasses the reporter tends to be investigated while one that flatters them tends to be defended.
Abstract	Rewrite	Opens on the projection framing and the paper's thesis in plain terms, separates the definition into its own paragraph, and carries the rate and routing cases, the quality floor, and the attribution condition.
References	Addition	Five references added on endpoint benchmarking, behavioural fingerprinting, model equality testing, documented behaviour change over time, and the projection framing.

Section	Kind	Change
Front	Removal	The table of contents is removed; the paper runs from the title page into the abstract.

v1.0 · 10 August 2026

Section	Kind	Change
All	Posting	First version posted.