

Share of Model Measurement Worksheet

The executive tool for measuring whether AI engines recommend you to your buyers — not just whether they mention you.

What Share of Model is

Share of Model is the percentage of relevant buyer questions where AI engines — ChatGPT, Perplexity, Gemini, and Google AI Overviews — cite or recommend your company versus competitors. It is the answer-engine-era successor to share of voice: a measure of presence inside the machine that now filters the market, not presence in the market itself.

Why it matters

Buyers increasingly delegate the first pass of vendor research to AI. The shortlist is assembled before a form is filled or a rep is alerted. If you cannot measure whether the model recommends you, you cannot manage the most consequential moment in the modern buying cycle. **Mentions are vanity.**

Recommendation is pipeline.

How to use this worksheet

- Work through Sections 1–5 with your marketing, web, and product-marketing leads in one 90-minute session.
- Score honestly. The value is in the gaps you find, not the total.
- Use Section 6 to locate your maturity level, and Section 7 to commit to immediate, 30-day, and 90-day actions.
- Re-run the full worksheet on a monthly cadence and track the movement, not the absolute number.

Scoring methodology

Each item is scored 1 (absent / weak) to 5 (strong / consistently true). Section subtotals roll into an overall Share of Model Readiness score out of 100. A rising score on high-intent buyer questions, paired with falling Description Drift, is the signal that the program is working.

1 · Recommendation Ladder Assessment

AI presence has five rungs. Only the top two move pipeline. For your priority buyer questions, mark how often each rung is the *best* position the engines give you today.

Ladder rung	What it means	Frequency (1–5)
Rung 5 — Selected	The model names you as the best fit for this buyer	__ / 5
Rung 4 — Recommended	You are put forward as a strong option, with reasons	__ / 5
Rung 3 — Compared	You appear in a side-by-side with named rivals	__ / 5
Rung 2 — Cited	Your content is quoted or linked as a source	__ / 5
Rung 1 — Mentioned	Your name appears, with no endorsement attached	__ / 5

Scoring guidance: Score 5 only when the engines consistently place you on that rung for your high-intent questions. Most teams discover they are strong at Rungs 1–2 (over-counted by tools) and weak at Rungs 4–5 (where deals begin). Your Ladder score is the sum of Rungs 4 and 5, out of 10.

2 - Qualified Share of Model (QSoM) Calculator

QSoM = the share of high-intent buyer questions where AI engines actively recommend you, weighted by the buying influence of each question.

Formula

$$\text{QSoM} = (\sum [\text{intent weight} \times \text{ladder multiplier}]) \div (\sum \text{intent weight}) \times 100$$

Inputs & weighting

- **Intent weight** per question: High = 3, Medium = 2, Low = 1.
- **Ladder multiplier** by best position: Recommended or Selected = 1.0, Compared = 0.5, Cited = 0.3, Mentioned = 0.2, Absent = 0.

Worked example (ten questions)

Buyer question	Intent weight	Best rung	Weighted
"Best platform for a regulated buyer"	High (3)	Recommended	3.0
"You vs a named competitor"	High (3)	Compared	1.5
"Who should a regulated buyer trust"	High (3)	Absent	0.0
"Top tools in the category"	Low (1)	Mentioned	0.2
"List of category vendors"	Low (1)	Mentioned	0.2
Five more low-intent prompts	Low (5)	Mentioned	1.0
Total	16	—	5.9

Raw Share of Model = 9 of 10 appearances = 90%. **Qualified Share of Model = 5.9 ÷ 16 = 37%**. The gap between the two numbers is the shape of the problem: ubiquitous on low-stakes questions, absent on the high-intent one that decides deals.

Your calculation

Buyer question	Intent	Best rung	Weighted
_____	_____	_____	_____
_____	_____	_____	_____
_____	_____	_____	_____
_____	_____	_____	_____
_____	_____	_____	_____
_____	_____	_____	_____
Your QSoM = (Σ weighted) ÷ (Σ intent weight) × 100			_____ %

3 - Description Drift Audit

Visibility without accuracy is a liability. Where the engines name you, score how accurate and consistent the description is (1 = frequently wrong, 5 = consistently accurate).

Dimension	What you are checking	Score (1–5)
Messaging consistency	Category, product, and ideal customer are described the way you describe them	__ / 5
AI interpretation consistency	The engines agree with each other about what you do and who you serve	__ / 5
Positioning consistency	Your differentiator survives the model's summary instead of being averaged away	__ / 5
Claim accuracy	Capabilities, pricing tier, and proof points are current and correct	__ / 5

Read it: A rising recommendation rate with rising Description Drift is not a win — it is a louder version of the wrong story. Treat any dimension scoring 1–2 as a priority fix at the source.

4 - Share of Model Measurement Stack Assessment

Rate the maturity of each measurement capability (1 = not in place, 5 = systematic and owned).

Capability	Standard to score against	Score (1–5)
Visibility	We know, monthly, which engines surface us for our top buyer questions	__ / 5
Citations	We track when our content is quoted or linked as a source	__ / 5
Comparisons	We monitor head-to-head comparisons the engines generate about us	__ / 5
Recommendations	We measure how often we reach Rung 4–5, not just Rung 1	__ / 5
Conversion indicators	We tie AI-shaped research to assisted pipeline and influenced deals	__ / 5

5 - Share of Model Maturity Model

Locate your organization. Movement up one level per quarter is a strong, realistic pace.

Level	Characteristics
Level 1 — Invisible	The engines rarely name you. Competitors own the shortlist and you find out by accident.
Level 2 — Mentioned	You appear in lists, but never with endorsement. High raw share of model, near-zero qualified.
Level 3 — Considered	You show up in comparisons on high-intent questions, but the recommendation usually tips elsewhere.
Level 4 — Recommended	The engines put you forward with reasons on the questions that precede real deals.
Level 5 — Preferred	You are the default answer in your category, described accurately, hard for rivals to dislodge.

Our current level: ____ Our target next quarter: ____

6 - Executive Action Plan

Immediate (this week)

- Assemble 20–40 real buyer questions from live deals and sales calls.
- Run them across four engines and record your ladder position and Description Drift. This is your baseline.
- Name one accountable owner for AI visibility.

30-day

- Rebuild your highest-intent pages as citable substance: answer-first, specific, current, with real data.
- Fix the worst Description Drift at the source.
- Publish or sharpen one original framework or proprietary benchmark only you can offer.

90-day

- Stand up a corroboration program: reviews, analyst and press mentions, communities the model trusts.
- Re-run the question set and report the movement in Qualified Share of Model to leadership.
- Set the ongoing rhythm: measure monthly, invest in substance quarterly, corroborate continuously.

Read the full article: [Share of Model — The AI-Era Metric Every B2B Team Should Track](https://erikmiller.com/blog/share-of-model-ai-visibility-metric) ·

erikmiller.com/blog/share-of-model-ai-visibility-metric

Related frameworks: [AI Visibility Architecture](#) · [AI Buying Committee](#) · [the ERM Revenue Execution System](#) ·

erikmiller.com/frameworks

Research & Supporting Evidence

- Forrester, *Buyers' Journey Survey (2025)* — generative AI / conversational search now a leading information source for B2B buyers.
- Gartner (2026) — a majority of B2B buyers prefer a rep-free, self-service research experience.
- McKinsey, *B2B Pulse (2024)* — buyers use roughly ten channels across the journey.
- Google (2025) — AI Overviews and AI Mode reached billion-plus user scale.