

The Authority Axis

Enterprise governance, not model capability, decides which AI does the work. And four competing labs already agree on where the line falls.


Carlson BCM™ scored 1,079 business and functional capabilities across the Life & Annuity value chain. Two attributes in that scoring were authored independently, and months apart. One asked a market question: for each capability, how well does each frontier model family fit, scored across OpenAI, Google, Anthropic, and Meta. The other, added much later, asked a governance question: must a human retain authority, because the capability creates a financial commitment or a policy obligation.

Neither was derived from the other. That matters, because when the two are placed side by side, they line up. **The preferred model for a capability tracks how much authority that capability retains.** This is not a definition reasserting itself. It is an out-of-sample result, the structure speaking back, across a thousand capabilities that were never scored with this comparison in mind and that arranged themselves along an axis nobody drew.


01

The vendors sort themselves by governance level

Order the four model families by the average authority of the capabilities each is preferred for, and they form a clean ladder. Commodity open-weight models anchor the floor. The frontier reasoning model sits at the ceiling. The human sits just below it.



The fit scores themselves move the same way. As authority rises, OpenAI fit climbs and Meta fit falls, a near-perfect mirror image. Claude crests in the middle: the augmentation band where AI works beside a human but commits nothing.



Mean fit score (1 to 3) by authority tier. OpenAI $r = +0.43$ with authority; Meta $r = -0.44$; Claude peaks in the Moderate band.

02

Three structural laws

If the ladder is the surface, these are the mechanics underneath it.

Authority and autonomy are near-opposites ($r = -0.74$)

As a capability retains more human authority, its ability to run as an autonomous agent collapses. High-authority capabilities are 90% Autonomize-Low. Low-authority capabilities are 70% Autonomize-High. The correlation is strong but not total, and the gap is the point: a small set of high-authority capabilities can still be delegated within a fixed mandate, which is why authority and autonomy are distinct axes rather than the same measurement. This is the quantified form of a single claim. Capability is not authorization.

The commitment tier is rule-bound, not generative

At high authority, fit for rules-based AI scores 2.38 while fit for generative AI scores 1.35. The work a human must own is not the creative, open-ended work. It is the deterministic, policy-governed, auditable work: bind under these terms, disburse under that schedule. Authority concentrates exactly where the logic is most explicit, which inverts the common assumption that the hardest capabilities demand the most generative model.

You cannot buy past authority

Of the 72 high-authority capabilities, 63 carry the top fit score for OpenAI. A capable model exists. Yet 18 of those are assigned to no model at all and kept with the human. The ceiling is not a capability gap that a stronger model closes next year. It is a legitimacy reservation. The vendor ladder runs out precisely at the commitment boundary, because that boundary is not about what AI can do.

03

Two value currencies

Business value was scored across five levers. Two of them move with authority, in opposite directions, and together they explain why both halves of the enterprise matter.

2.47 →

RISK & COMPLIANCE

value rises with authority (1.59 → 2.04 → 2.47). The commitment tier is where compliance value concentrates. This is the work that protects the license to operate.

2.46 →

CUSTOMER EXPERIENCE

value rises as authority falls (1.89 → 1.89 → 2.46). The shiftable 932 is where experience value concentrates. This is the work that wins and keeps customers.

The high-authority half of the enterprise creates value by protecting. The low-authority half creates value by delighting. They trade in different currencies, and the model assigned to each tier follows the currency: Meta serves the experience work, OpenAI serves the compliance work.

The workforce map.

Authority is owned by people. When each accountable role is measured by the average authority of the work it owns, its dominant AI partner falls out automatically. Who you are in the organization determines which model you inherit, because the authority of your work does.

ACCOUNTABLE ROLE	CAPABILITIES	MEAN AUTHORITY	DOMINANT MODEL
Customer Service Representative	305	1.34	Meta · 62%
Insurance Claims & Policy Processing	72	1.35	Claude · 36%
Computer Systems Analyst	94	1.84	Claude · 44%
Business Operations Specialist	81	1.86	Claude · 56%
Information Security Analyst	65	1.92	Claude · 31%
Insurance Underwriter	44	2.02	mixed
Financial Manager	177	2.21	OpenAI · 33%

The largest role in the enterprise, Customer Service, owns the lowest-authority work and pairs with commodity models. The analyst and operations roles in the middle pair with Claude. Financial Manager, who owns the commitment-heavy billing work, pairs with the frontier model. The workforce transition is not random. It concentrates, and it concentrates along authority.

Why deployments stall, and how to unSTALL them.

Enterprises stall because they treat model selection as one decision, applied uniformly, gated by one trust conversation. The data says the right model is a function of the authority tier of the capability. Recognition and retrieval go to commodity open models. Augmentation goes to Claude. Commitment-adjacent, rule-bound judgment goes to a frontier model under deterministic rules. The commitment itself stays with a human, no matter how capable the models become.

Route by authority and the 932 capabilities that carry no commitment move into software now. The 70 that bind risk, move money, or change what is owed stay human, by design, not by hesitation. The governance structure of the enterprise is the model-selection architecture. It was there before anyone asked the question.

Capability is nearly universal. Authorization is scarce. Once a carrier separates the two, scaling AI becomes an architecture problem rather than a trust problem.