

Embedded Judgment

*A two-axis analysis of cognition and commitment in
Commercial Property & Casualty New Business
underwriting.*

ABSTRACT

Debates about the division of labor between humans and artificial intelligence usually rest on a single intuition: that hard work is human work. This paper separates that intuition into two independently measured constructs, how much cognitive judgment a capability demands (Augmentation) and whether the capability commits the enterprise to an obligation (Authority Retention), and applies them, under a governed verb-based scoring doctrine, to a complete commercial property & casualty value chain.

Across 560 business capabilities in Commercial New Business underwriting, cognitive demand is broad while binding authority is narrow and concentrated. The two axes diverge: among capabilities requiring high judgment, roughly half stop short of commitment (103 capabilities of embedded judgment) and roughly half cross into it (98 capabilities of commitment-bearing judgment). The divergence is reported here as a descriptive result for a single value chain. Critically, it is a lower bound: every refinement to the scoring doctrine moved in the one direction that should have erased it, and it persisted regardless.

01

The question, made measurable

The argument over what humans should do and what AI should do is almost always conducted philosophically. This paper attempts to conduct it arithmetically inside an actual operating model.

The prevailing framing treats cognitive difficulty as the dividing line: routine work automates first, complex work stays human. That framing collapses two different questions into one. The first is how much thinking a unit of work requires, synthesis, evaluation of trade-offs, formation of a recommendation. The second is whether the work binds the enterprise, moves money, changes a policy obligation, establishes a commitment that did not previously exist. These are not the same property, and there is no a priori reason they should travel together.

Carlson BCM™ scores every authored capability on both. This paper holds the two apart deliberately, measures them across one complete value chain, and reports what happens to the relationship between them. The finding is that they come apart, and the shape of how they come apart is the contribution.

The two-axis instrument.

Each business capability is scored independently on two governed dimensions, expressed on a three-level scale (High / Medium / Low).

Authority Retention: **does the work commit the enterprise?**

High authority binds risk, changes a policy obligation, moves money, or establishes a commitment: binding and unbinding, writing and issuing policies, cancellation and reinstatement, reinsurance establishment, coverage construction, manuscript contract language, premium overrides, and mid-term contractual modification. Medium authority recognizes, configures, or prepares without committing. Low authority consumes, view, retrieve, present, notify.

Augmentation: **how much cognitive judgment does the work demand?**

High augmentation marks underwriting judgment, pricing and risk evaluation, and authority-sensitive decision support. Medium marks workflow, synthesis, and calculation. Low marks display, tracking, and passive presentation.

THE GOVERNING DISCIPLINE

The governing discipline of the instrument is that the verb in the capability statement carries the score, not the object. “Capture the underwriting rationale” is not deciding; it records a decision already made. The object repeatedly tries to overpower the verb. Holding that line is the entire method.

This verb-over-object rule resolved into a small, reproducible set of distinctions that held across the corpus:

- **Capture ≠ Decide.** Recording requested limits, received loss data, or a stated rationale is Medium augmentation, not High.
- **View ≠ Analyze.** Viewing a precomputed rating worksheet is Low; generating one is High.
- **Validate ≠ Judge.** Rules-checking a submission is Medium; evaluating the risk is High.
- **Supported attachment ≠ Discretion.** Automatic, rule-driven form attachment is Medium authority; manual non-standard or manuscript wording is High.

Corpus and method.

The instrument was applied to the complete Commercial New Business underwriting value chain of the Carlson BCM™ commercial property & casualty model: a governed set of authored business capabilities spanning submission intake, risk assessment, pricing, coverage construction, reinsurance, binding, and issuance. Each capability was read in the context of its full descriptive statement, not from the verb in isolation.

Scoring proceeded in twenty-three controlled batches against a fixed queue. Each batch was independently proposed, then validated capability-by-capability against the source description and the accumulating precedent table, corrected where a proposed score departed from doctrine, and committed to a single persistence layer with a parallel audit ledger. Every committed score reconciles code-for-code against that ledger.

560	2	23	0
CAPABILITIES SCORED	GOVERNED AXES	CONTROLLED BATCHES	UNRESOLVED EXCEPTIONS

04

Results

The full joint distribution of Authority Retention against Augmentation is given below. Marginally, binding authority is scarce, only 107 of 560 capabilities (19%) are High authority, while cognitive demand is roughly twice as common: 201 capabilities (36%) are High augmentation.

	AUGMENT H	AUGMENT M	AUGMENT L
AUTH H	98 COMMITMENT JUDGMENT	9	0
AUTH M	103 EMBEDDED JUDGMENT	240	32
AUTH L	0	4	74

Fig. 1 · Authority Retention × Augmentation, 560 Commercial New Business capabilities. Teal: embedded judgment (high cognition, no commitment). Ember: commitment-bearing judgment.

Two features carry the paper. First, among the 201 high-cognition capabilities, the split between those that commit the enterprise and those that do not is nearly even: 98 commitment-bearing against 103 embedded. Second, the high-cognition / consumption-authority cell is empty: there is no population of detached advisory expertise sitting outside the operating system. Every instance of high judgment has an operational footprint, it classifies, assesses, prices, compares, configures, but only half of it crosses the commitment boundary.

04 · B

Where the separation lives

Grouping the corpus by underwriting stage shows the two axes are not evenly mixed across the value chain. Embedded judgment and commitment concentrate in different stages.




Fig. 2 · Embedded judgment (teal) vs commitment judgment (ember) by underwriting stage, in flow order. Counts of capabilities; stages carry multiple memberships, so totals overlap.

The separation is positional, not just statistical. Read along the underwriting flow, cognition climbs through Risk Assessment & Pricing and peaks at Quote Generation, where nine capabilities demand high judgment and none commit the enterprise. It inverts at the very next stage: Bind Decision and Policy Issuance carry the commitment and almost none of the detached judgment. The lifecycle-mutation stage behaves the same way, because changing an obligation is a commitment act. Judgment concentrates upstream; commitment fires at the bind boundary and wherever the contract changes.

Within Commercial New Business underwriting, judgment detaches from commitment **without detaching from execution.**

The 103 embedded-judgment capabilities are the interesting population. They demand real expertise, an experienced underwriter synthesizes evidence, evaluates trade-offs, and forms a recommendation, yet they never themselves change what the enterprise owes. The commitment happens elsewhere, later. They are not administrative, rote, or deterministic; nor are they detached consultation. They are high judgment paired with recognition-or-configuration authority.

COMMITMENT JUDGMENT • AUTH H / AUG H

Bind coverage directly
Cancel a policy
Reinstate a policy
Rewrite a cancelled policy
Cancel reinsurance
Define a reinsurance contract

EMBEDDED JUDGMENT • AUTH M / AUG H

Assess underwriting risk
Determine a composite rate
Pre-qualify a submission
Implement a technical pricing model
Compare premiums
Collect data to underwrite

The implication for AI strategy is direct, and corrects both extremes of the current debate. The addressable surface for AI assistance is roughly double the surface for AI autonomy, 201 high-cognition capabilities against 107 high-authority ones. The transition path is therefore not human execution giving way to AI advice giving way to AI commitment. It is AI penetrating upward through execution into judgment while the commitment layer remains human.

*Difficulty is not the boundary; **commitment** is.*

Why the result is a lower bound.

A natural objection is that the scoring doctrine manufactured the finding. The directionality of the corrections answers it. Across all twenty-three batches, every material refinement moved the same way, a capability's augmentation was revised downward, from High to Medium, never upward. Each such correction took a capture, view, validate, calculate, screen, or document capability and moved it out of the high-cognition band.

Those corrections landed almost entirely on Authority-Medium capabilities, that is, they were draining the embedded-judgment cell. No correction ever touched the commitment-judgment cell. Denoting embedded judgment E and commitment judgment C , the refinement produced $\Delta E < 0$ with $\Delta C = 0$. The doctrine was, mechanically, suppressing the very population that produces the finding.

THE ANSWER

The critique says: your doctrine manufactured the finding. The answer is: the doctrine systematically compressed the cell that produces it, and it persisted anyway. The 103 embedded-judgment capabilities are a lower bound, not an estimate inflated by permissive scoring.

From finding to program.

The embedded-judgment result turns a philosophical dispute into a falsifiable empirical program with a single decisive statistic. The hypothesis under test is whether high cognition systematically extends beyond commitment-bearing work across enterprise value chains, governed by the question: among Augment-High capabilities, does the Authority-Medium population remain materially above zero?

This is Paper II of a series. Paper I, The Authority Axis, established authority as an independent dimension of enterprise AI suitability. Paper III, Commitment, Not Difficulty, will extend the instrument across the value chain, with at least one chain scored independently to guard against single-instrument bias. This paper supplies the instrument, the worked example, and the first clean separation of cognition from commitment inside a real operating model.

Enterprises may retain people less because work is hard than because work commits. Commercial New Business underwriting provides the first measurable evidence consistent with that separation, and the foundation the series builds on.

Workforce Capacity Intelligence Working Paper Series, Paper II. Method, scoring doctrine, and scored ledger available as a companion artifact.