

THE LEGITIMACY GAP

Eric Solomon, Ph.D. and Lambert Hogenhout

When Everyone Has an Answer

As AI gets better at producing answers, do the answers themselves become less valuable?

Consider a Monday morning at a mid-sized company. The CEO has asked her leadership team to evaluate a potential market entry. By noon, eight slide decks are waiting in her inbox, each generated overnight with AI tools, each internally coherent, each arriving at a different conclusion. Three recommend entering the market. Two recommend waiting. One recommends acquiring a competitor instead. The remaining two are hedged in ways that technically allow agreement with any of the others.

Which one is right? And whose answer counts?

As AI capabilities have spread across organizations, they have done something more disruptive than automate tasks: they have democratized the production of credible-seeming strategic analysis. Any employee with a well-crafted prompt can now produce a competitive scenario set, a business model evaluation, or a financial analysis that, according to recent research, rivals what experienced investors and entrepreneurs produce. Generating a market entry memo requires no special expertise, no senior title, and no lengthy process. It requires a good prompt and a few minutes.

The common narrative frames this as a story about efficiency and empowerment: AI flattening hierarchies, giving junior employees access to capabilities once reserved for expensive consultants. That narrative is true as far as it goes, but the deeper story is about *accountability*. When analysis was hard to produce, scarcity itself was a filter. Only well-considered views made it onto the table. When analysis is easy, that filter disappears. The question for organizations is no longer who can produce the best answer. It is who owns the decision when everyone in the room has a credible one.

We call the resulting void the *Legitimacy Gap*. And closing it requires understanding not just that accountability has become scarce, but why.

What the Research Shows

The academic literature on AI and organizational decision-making has grown rapidly, and its findings are more nuanced than the headlines suggest.

At the individual level, large language models are now genuinely competitive with human experts on structured analytical tasks. Csaszar, Ketkar, and Kim (2024) ran controlled experiments with a startup accelerator and found that LLM-generated business plans received scores and investment interest comparable to those written by the entrepreneurs themselves. AI outperformed humans most clearly among weaker plans, suggesting it functions as a floor-raiser for strategic thinking rather than a ceiling-breaker.

Doshi, Bell, Mirzayev, and Vanneste (2025) complicate the picture in ways that matter enormously for organizations. Individual AI evaluations are inconsistent and systematically biased, exhibiting order effects and preference distortions that closely mirror human cognitive limitations. A single AI analysis is no more reliable than a single human one. But when the researchers combined evaluations across multiple models, roles, and prompting strategies, the resulting rankings converged reliably with those of human expert judges. Aggregation, not any single model, produced the most reliable output.

That finding has a counterintuitive wrinkle. The reliability of the output was a function of how the human and machine ensemble was designed: who chose which models, which prompts, and who decided when the aggregation was sufficient. The AI does not adjudicate these questions. A human does. The research literature largely stops at this cognitive frontier, asking how to make AI more accurate. Organizations are asking a different question: who is in charge?

Two Disruptions, One Crisis

AI disrupts organizations in two distinct ways, and most leaders are focused on the wrong one.

The first is *access democratization*. Retrieval-augmented systems now connect workers to institutional knowledge that was previously inaccessible: the company's full repository of strategic plans, past competitive analyses, and documented decisions. The informational moat that once surrounded senior roles can now be partially replicated by a junior analyst with the right system prompt.

The second is *volume amplification*. Generative AI does not merely democratize access to existing knowledge; it manufactures new content at a pace that overwhelms existing decision processes. A leadership team that once received one carefully synthesized strategy memo now receives twelve plausible, internally consistent ones, each generated in hours rather than weeks.

Access democratization lowers the barriers to participation. Volume amplification raises the cost of judgment. Together, these two disruptions create a crisis that has little to do with speed or accuracy. It is about accountability: who decides which analyses matter, whose judgment resolves contradictions, and who bears responsibility when the decision is made.

Why AI Makes Leadership Harder

The intuitive prediction is that AI makes leadership easier by reducing analytical burden. The evidence suggests the opposite.

Xu, Hou, Chen, and Xie's (2025) model of knowledge-based hierarchies shows that as generative AI capabilities increase and hallucination rates decrease, organizations do not simply flatten. They tend to reduce worker-level knowledge requirements, producing a deskilling effect at the operational level, while the need for managerial oversight can actually *increase*. More reliable AI generates more AI-assisted output. More output creates more material that must be validated, adjudicated, and committed to. That validation burden can narrow the span of control rather than widening it.

AI does not reduce the need for judgment, it relocates it. The bottleneck has shifted from production to evaluation, and the scarce resource is no longer “who can produce the best analysis” but “which analysis do we act on, and who is accountable for that choice.”

There is an additional layer that the organizational economics literature does not fully capture: for most of organizational history, the person who made a decision bore a traceable relationship to the information that informed it. When strategy came from a senior leader, it was reasonable to hold that leader accountable. They had reviewed, synthesized, and vouched for the analysis behind it.

In an AI-saturated environment, that traceability is severed. When a decision is informed by an ensemble of AI-generated analyses, orchestrated by a mid-level analyst, summarized by a different team member, and presented by a third, the accountability chain becomes invisible. Organizations are already experiencing the consequences: decisions that felt widely validated at the time of making them, and for which no one is willing to be held accountable when they fail.

The Legitimacy Gap

These dynamics converge in the **Legitimacy Gap**: *the growing void between the formal authority vested in leadership positions and the substantive basis that once justified that authority.*

Legitimacy in organizations has always been partly formal (title, reporting line, budget authority) and partly substantive (demonstrable expertise, track record, analytical edge). For most of the twentieth century, these two sources reinforced each other. Senior leaders had formal authority because their organizations believed, usually correctly, that they held an informational advantage. The hierarchy was both a power structure and an approximation of epistemic structure.

AI has broken that alignment. Formal authority structures remain unchanged, while the substantive basis that made those structures seem reasonable is eroding. When a junior analyst can generate an analysis as rigorous as a senior executive's, why should the executive's view prevail?

The most dangerous organizational response is to paper over it: to behave as if senior leaders still hold the informational advantage they no longer possess, to defer decisions upward out of reflex rather than because the people at the top know more. This produces the worst of both worlds: the inefficiency of hierarchy without the epistemic benefits that once justified it.

A Framework for the New Authority

If informational advantage no longer anchors leadership authority, what does? Legitimate authority in AI-saturated organizations must be reconstructed around three distinct functions, each requiring explicit assignment and deliberate design. Together, they describe a leader who looks less like a chief answer officer and more like an architect of disagreement: the person responsible for designing how competing analyses come into the room, get adjudicated, and turn into commitments.

Layer 1: The Question. Whoever controls what gets asked controls the strategic space before any analysis runs. When generating an analysis cost weeks of analyst time, the question being asked got scrutinized before work began. When it costs minutes, the question often goes unexamined, and the framing embedded in the prompt shapes every output that follows. Power at this layer belongs to whoever sets the prompt, chooses the variables, and decides which scenarios are worth modeling. In many organizations today, that person is not the senior leader. It is whoever is closest to the AI system, often a junior analyst who has defaulted into this function without explicit authority or accountability. Reclaiming this layer requires leaders to articulate, deliberately and publicly, what questions the organization is asking and why.

Layer 2: The Synthesis. When multiple analyses exist, and in AI-saturated organizations they always will, someone must determine which perspectives were considered, how conflicts were resolved, and what the organization's official view is. In practice, synthesis often defaults to whoever writes the summary slide or chairs the final meeting, by organizational accident rather than design. The person who synthesizes is deciding, in effect, what the organization believes. If that authority is not explicitly granted, it is still being exercised, just invisibly, without accountability, and often without awareness. Effective synthesis in the AI era requires the ability to assess the quality of AI-generated reasoning and to identify where different analyses are making different assumptions rather than reaching different conclusions.

Layer 3: The Commitment. The third layer is commitment: the conversion of analysis into action, with explicit acceptance of consequences if the action fails. Commitment is what makes decisions real. Without it, you have infinite analysis and no action, which is precisely what many AI-saturated organizations are experiencing right now. This is where accountability gets manufactured. The leader who commits is doing more than approving an analytical output. They are staking their credibility on a judgment that cannot be fully derived from any analysis, however sophisticated: that this is the right course of action, for this organization, at this moment, given these stakes. That is an inherently human act. Organizations that do not assign commitment authority explicitly tend to experience analytical paralysis: the perpetual deferral of action in favor of additional analysis that never quite resolves the uncertainty.

Diagnosing Your Organization

Three questions show where the Legitimacy Gap is most acute. They can be asked of any significant decision your leadership team has faced in the past year.

Who actually controls the Question layer?

In practice, who decided what scenarios to model, what variables to include, what time horizon to use? Was that decision made consciously, by someone with accountability for it? Or did it happen by default, embedded in the choice of tool or the preferences of whoever ran the analysis?

Who performs the Synthesis, and have they been given explicit authority to do so?

Is the person producing the final synthesis the same person who will be held accountable for the quality of the analytical process? Or are they filling a vacuum that no one else chose to occupy?

Who makes the Commitment, and is that the same person who will own the consequences?

In many organizations, the commitment function has been implicitly delegated to process. The decision is said to have been made when the analysis was approved, or when the deck was presented, or when no one objected in the meeting. That is not commitment. It is deniability by consensus.

In most organizations, the answer to all three questions is either “nobody knows” or “it is fragmented across people who do not have explicit authority.” That fragmentation is the Legitimacy Gap made structural. Map the last three significant decisions your organization made and trace, honestly, where each layer sat. The gaps in that map are your leadership design problem.

The gap is not between what AI can do and what humans can do. It is between the authority structures organizations have and the authority structures they *need*.

What Leadership Is Today

The Legitimacy Gap will not close by itself, and it will not close by pretending the informational advantage has not been lost. Leaders who respond to AI by doubling down on formal authority, asserting hierarchy more loudly while delivering less analytical value, will find that the assertion rings increasingly hollow to the people around them.

Closing the gap requires replacing informational authority with process authority. The legitimate leader in an AI-saturated organization is the person who has been explicitly authorized to own the Question, perform the Synthesis, and make the Commitment, and who accepts the accountability that comes with each. Architecting disagreement, in other words, rather than producing answers.

This is a harder form of leadership than the one it replaces. Cognitively, it requires sophisticated judgment about AI-generated reasoning that goes beyond subject-matter expertise. Institutionally, it requires the explicit redesign of authority structures that most organizations have never had to think carefully about, because the old informational hierarchy handled them by default.

It is also a more honest form of leadership. The leader who says “I have reviewed the analytical landscape, I have made a judgment about which synthesis to trust, and I am committing to this course of action and accepting the consequences” is doing something qualitatively different from the leader who says “the analysis supports this decision.” The first is exercising authority. The second is hiding behind it.

The old question was “what does the AI say?” The harder one is: what is the minimum information I need to make this decision well, where should it come from, and who is responsible for ensuring it reaches me in a form I can act on? Those are design questions, and they have design answers. Answering them is, increasingly, the work of leadership.

References

Csaszar, F. A., Ketkar, H., & Kim, H. (2024). Artificial intelligence and strategic decision-making: Evidence from entrepreneurs and investors. *Strategy Science*, 9(4), 322–345.

<https://doi.org/10.1287/stsc.2024.0190>

Doshi, A. R., Bell, J. J., Mirzayev, E., & Vanneste, B. S. (2025). Generative artificial intelligence and evaluating strategic decisions. *Strategic Management Journal*, 46(3), 583–610.

<https://doi.org/10.1002/smj.3677>

Xu, F., Hou, J., Chen, W., & Xie, K. (2025). Generative AI and organizational structure in the knowledge economy (arXiv preprint No. 2506.00532). <https://arxiv.org/abs/2506.00532>