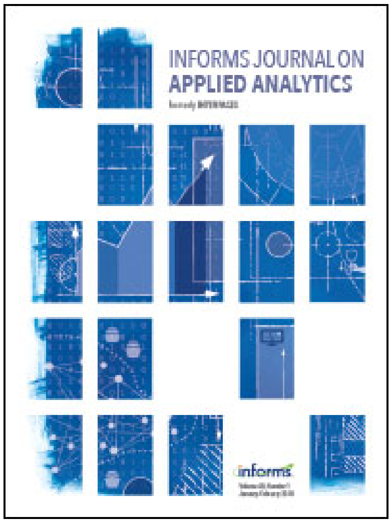

INFORMS Journal on Applied Analytics

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

The Fifth Column: Getting to the Obvious Answer the Hard Way

Adam DeJans Jr.

To cite this article:
Adam DeJans Jr. (2026) The Fifth Column: Getting to the Obvious Answer the Hard Way. INFORMS Journal on Applied Analytics

Published online in Articles in Advance 30 Jul 2026

. https://doi.org/10.1287/inte.2026.0326

INFORMS JOURNAL ON APPLIED ANALYTICS
Articles in Advance, pp. 1–4
https://pubsonline.informs.org/journal/ijaa
ISSN 2644-0865 (print), ISSN 2644-0873 (online)

The Fifth Column: Getting to the Obvious Answer the Hard Way

Adam DeJans Jr. [a]

[a] Gurobi Optimization, Canton, Michigan 48188
Contact: adam.dejans@gurobi.com, https://orcid.org/0009-0000-0844-347X (ADJ)

Received: February 23, 2026
Revised: May 2, 2026
Accepted: May 9, 2026
Published Online in Articles in Advance: July 30, 2026

https://doi.org/10.1287/inte.2026.0326

Abstract. One of the world's largest manufacturers relied on dozens of production planners using spreadsheets to coordinate production across multiple continents, hundreds of suppliers, and dozens of product families. The process absorbed contradictions through politics and manual overrides, producing plans that were feasible only because no single person ever saw all the constraints at once. The core optimization problem was textbook-simple on paper but resisted solution for years because the organization was not ready to trust what the math had to say. Mixed-integer linear programming was deployed incrementally, with capabilities chosen for what they let planners see rather than for raw algorithmic power: surfacing the smallest set of mutually inconsistent commitments behind an infeasibility, attaching a price to constraints that had been treated as absolutes, presenting a menu of near-optimal plans rather than a single prescription, and encoding planner priorities in a clear order of precedence. A natural-language interface and a Monte Carlo robustness layer were also built on top of this system; both warrant their own separate treatment, which we leave for future work. Over a two-year period, the optimization-driven production planning system freed up production capacity that the prior process had been leaving on the table, unlocking more than 40,000 additional finished units per year and approaching $400 million in realized profit improvement across one of the world's largest supply chains. Planning cycles that once consumed weeks ran in hours, and disruptions were met with structured scenario analysis rather than emergency meetings. The key to adoption was not algorithmic sophistication but tooling that made the model's reasoning legible to planners, preserving their agency and enabling trust-based organizational change.

Keywords: mixed-integer programming • production planning • decision support systems • infeasibility analysis • organizational change management

1. Introduction

Gene Woolsey spent decades reminding the operations research community that the real world does not read textbooks, and that it punishes those who assume it should behave like one. Every generation of practitioners rediscovers this the hard way. This is a story about how an obvious optimization problem at one of the world's largest supply chains resisted being solved for years, not because the math was hard, but because the organization was not ready to hear what the math had to say. Once it was, the same model quietly unlocked more than 40,000 additional finished units per year and approaching $400 million in realized profit improvement. The interesting part is what it took to get there.[1]

2. The Job

I joined a major manufacturer as an operations research practitioner, part of an internal group charged with rethinking how products moved from production commitment to customer delivery. The mandate was broad: modernize the way the company planned what to build, where to build it, and how to get it where it needed to go. The company had production facilities spread across multiple continents, hundreds of supplier relationships, and dozens of product families.

The core question sounded deceptively simple: how many of each product to build, where, when, and using which suppliers, subject to the usual capacity and throughput limits and to demand forecasts that everyone acknowledged were wrong but nobody could agree on how to replace. A textbook mixed-integer program. I could sketch the formulation on a napkin. The first prototype took a few weeks and solved in a minute or two. I remember watching it finish and thinking, naively, that the hard part was over.

3. The Planning Floor

The production planning process did not look broken, and that was precisely why it survived for so long. When I first entered the planning cadence, what I encountered was not chaos but a system that had evolved to accommodate decades of organizational scar tissue. Meetings were calm and procedural. Spreadsheets were opened before conversations began, already filtered and annotated, each reflecting the worldview of the person who owned it. Every plan emerged feeling negotiated rather than designed.

Each production planner managed risk on behalf of their function. A product-line planner carried regional demand commitments that the sales team had promised months earlier, and the planner, not the sales team, would have to explain any shortfall. A capacity planner worried less about which products were built than about avoiding expensive changeovers. A supply planner tracked contracts that included minimum purchase obligations, penalty clauses, and informal understandings that never appeared in writing. A regional demand planner fought for allocations to keep customers stocked. Each perspective was legitimate. None was complete.

Excel functioned as the common language because it allowed each production planner to remain sovereign. A spreadsheet did not challenge assumptions; it obeyed them. If two assumptions conflicted, the conflict surfaced socially, in side conversations and compromises that were rarely documented. I once watched a supply planner spend 45 minutes honoring a supplier minimum that, had anyone checked, was no longer contractually binding. The agreement had expired two quarters earlier, but the planner had been burned once and was not about to let it happen again. The spreadsheet did not know. It just held whatever numbers were typed in.

On any given cycle, dozens of production planners would independently update their slices of the world, email revised spreadsheets to a central coordinator, who would reconcile them by hand and convene a meeting where people argued about color-coded conflicts. The coordinator was, in every meaningful sense, a human solver: holding constraints in their head, managing the objective function through intuition, producing feasible plans through sheer force of will. The process worked. It was also extraordinarily fragile. When the senior coordinator took vacation, the cycle did not stop; they stayed reachable on personal time, while a backup who had been ramped up over months absorbed the day-to-day reconciliation. Even with that hand-off rehearsed, plans drifted out of sync, the meeting slipped, and the room rarely converged on the first attempt.

The cost of that fragility was felt most acutely under disruption. Even an uneventful planning cycle ran on the order of weeks: planners pushed proposals, the coordinator reconciled, scenarios were evaluated by hand, and feasibility was verified through arguments rather than arithmetic. When a key supplier announced a multiweek capacity cut midcycle, the response was not a reoptimization but a series of emergency meetings: planners pulling their spreadsheets back open, renegotiating commitments by phone, and producing a manually rebalanced plan days later. The plan that finally emerged was defensible, but no one could say with confidence that it was anywhere near the best feasible response. There was no time, and no shared model, to ask.

4. When the Solver Said No

When the model was first introduced, the response was polite and conditional. Optimization could be used as long as it respected all existing rules and confirmed what everyone already believed. The implication was that the model should reflect reality as production planners understood it, not challenge it.

Constraints accumulated rapidly. Some represented genuine physical limits: line rates, labor headcount, tooling windows. Others were ghost constraints, haunting the model long after the conditions that created them had changed. Still others were political compromises that had calcified into policy. Each made sense in isolation. Together, they formed a system that could not breathe.

One meeting exposed this in a way no abstract discussion could. A product-line planner stated the plant needed at least 600 units of a high-demand product. A manufacturing planner added that no single program could exceed 40% of plant output. A capacity planner noted the plant was capped at 1,000 units. The room fell silent. Six hundred is more than 40% of 1,000. In the spreadsheet process, this tension had always been handled implicitly. Someone would shade a forecast. Someone else would lean on a supplier. The math never objected because it never saw all the numbers at once. The solver saw them all, and returned `INFEASIBLE`.

Rather than explaining the infeasibility myself, which would have positioned me as the adversary, I asked the solver to explain itself. Most capable commercial MILP solvers can isolate a minimal set of constraints that cannot all hold simultaneously and report them back; remove any one of them and the rest become feasible. Think of it as a neutral referee that points to exactly which two or three commitments in the room are mutually impossible, without taking sides. In our case, the solver identified exactly three: the product minimum, the mix cap, and the plant capacity. Choose any two. The third must give. The conflict was no longer between production planners. It was between beliefs. No one was being called unreasonable. The inconsistency had been made explicit by an algorithm with no agenda. A product-line planner eventually said the 600-unit minimum might have flexibility if customer commitments could shift to the following month. The referee had done something I could not: it made it safe to reconsider.

5. From Absolutes to Tradeoffs

Even after the referee surfaced contradictions, progress was slow. Many constraints reflected real risks. A supply planner who had personally managed a supplier crisis was not going to abandon a minimum purchase constraint just because a model said it was inconvenient. These were scar tissue, and scar tissue exists for a reason.

What changed the conversation was attaching a price to every constraint that planners insisted on protecting, a technique known in the literature as feasibility relaxation. The math is almost embarrassingly simple. For each protected constraint, you allow it to be violated by some amount s and add a penalty $p \cdot s$ to the objective, where p is a user-chosen weight reflecting how important the constraint is. The solver then maximizes profit minus total penalty across all relaxed constraints. The effect is that every constraint gets a price tag. Constraints the solver respects for free are not binding anyone. Constraints the solver pays heavily to satisfy are the ones worth discussing.

One procurement manager stared at the output and said, "I had no idea that constraint was costing us that much." It was a minimum purchase volume negotiated three years earlier under different demand conditions, never revisited, silently draining hundreds of thousands of dollars per cycle. It was removed that afternoon, not because anyone overruled the manager, but because the manager chose to remove it. The tool provided information. The decision was theirs.

Over weeks, constraints that had been treated as sacred were quietly retired when their costs became visible. New constraints were added more carefully, with an awareness that every hard boundary carried a price. The model became feasible not because it was simplified, but because it was aligned with what the organization actually valued.

6. The Myth of the Single Best Plan

The first globally optimal solution was met with hesitation, not celebration. A manufacturing planner flagged a problematic changeover sequence. A supply planner noted the solution assumed smooth ramp from a volatile supplier. A regional demand planner observed it would require an uncomfortable conversation with a sales director who had been promised specific volumes. These were not objections to optimization. They were reminders that mathematical optimality is necessary but not sufficient for operational viability.

We configured the model to return not a single prescription but a set of plans whose costs were all within a small tolerance of the optimum, so that planners could choose between operationally distinct alternatives rather than be handed a verdict. Each plan in the set achieved an objective within a few percent of the optimum but differed in operationally meaningful ways: different changeover sequences, different supplier loading patterns, different regional splits. Production planners could see that moving to the third-best plan cost 0.3% in margin but eliminated a problematic changeover. They chose plans they believed the organization could execute, even when those plans left theoretical optimality on the table. The plan that gets executed well beats the plan that gets executed poorly.

Political resistance softened because the fight had never been about rejecting math. It had been about preserving agency. Production planners were accountable for outcomes long after the solver finished. A system demanding blind trust in a single solution was incompatible with that accountability. A system offering structured choice was not.

7. What Matters Most, and What Matters Next

The initial single-objective formulation (maximize contribution margin) was defensible but incomplete. A profit-maximizing model will chase margin wherever it finds it, swinging production between programs month to month. Mathematically correct. Operationally catastrophic. Plants cannot ramp like a faucet. Suppliers assume stable volumes. Customers cannot plan around wild swings. Production planners had always smoothed plans by hand. They articulated it not as a second objective but as common sense: "You cannot jerk a plant around like that."

The fix was to restructure the model around prioritized objectives. Maximize profit first; then, among plans that achieved near-maximal profit, minimize month-over-month production swings across plants and programs. The order is the decision; the rest is bookkeeping. Production planners immediately recognized the value. The model was finally respecting something they had always known: a good plan is not just profitable, but executable.

8. Letting Planners Drive

One critical need remained unmet. Production planners did not simply want answers; they wanted to ask questions. What if a supplier lost 20% of capacity? What if a tariff changed cost structures overnight? What if a plant ran an extra shift for one month? Traditional optimization workflows required modeling expertise to adjust parameters and rerun solves, recreating the bottleneck optimization was supposed to eliminate.

We addressed this through systematic exploration of the model's tunable parameters. Production planning models are full of coefficients and bounds that are not fixed by physics but by judgment: penalty weights on relaxed constraints, the profit tolerance used when prioritizing smoothness over margin, demand targets that reflect forecasts rather than firm orders, and supplier minimums tied to contracts with renegotiation clauses.

We automated the model so it could be resolved across a large grid of permutations of these parameters in a single overnight run, producing a landscape rather than a point answer. With that landscape in hand, planners could trace the tradeoffs between competing priorities, such as profit versus production stability, or total volume versus supplier concentration risk, and navigate to configurations that matched their operational judgment. When a planner discovered that relaxing a particular demand target by 5% or adjusting the smoothing tolerance opened up a meaningfully better plan, that insight could be taken back to the planning table as a concrete proposal. Exploration became the path to better production decisions.

We also built a natural-language interface on top of this system, so that planners could pose "what if" questions in plain English, and a Monte Carlo layer that stress-tested candidate plans against demand and supplier uncertainty. Both proved valuable, but they are sufficiently distinct in motivation and design that a fuller treatment of either belongs in its own paper, which we leave for future work; the planner-empowerment story above stands on its own.

The effect was to restore planner agency completely. Production planners could explore freely, ask "what if" on their own terms, identify better configurations through structured search, and adopt those configurations as their production plan. The iterative workflow that had always been the real craft of their profession was preserved, but now it operated within a globally coherent model rather than disconnected spreadsheets. The answers were consistent. The tradeoffs were quantified. The explanations were legible to anyone in the room.

9. The Only Path That Could Have Worked

Over two years, the transformation was difficult to overstate. The optimization-driven plans consistently freed up production capacity that conservative assumptions and coordination friction had been leaving on the table, including more than 40,000 additional finished units per year that the manual process had been quietly absorbing into safety margins. Planning cycles that once consumed weeks ran in hours. New production planners onboarded faster because the model encoded institutional knowledge that had previously existed only in the heads of senior staff. Disruptions were met with structured scenario analysis rather than days of emergency meetings. The financial impact, on the order noted at the outset, was easily large enough to fund the analytics organization many times over.

But the lesson here is not that optimization works. Of course it works. The lesson is that optimization deployed without regard for the organizational system it inhabits will fail, not because it produces wrong answers, but because it produces right answers nobody is prepared to act on. Results do not speak for themselves in large organizations. They are interpreted through trust, accountability, and political reality. People will not say "I reject your model." They will say "the model does not capture our situation" and add constraints until it agrees with them.

None of the mathematics in this story is novel. Every technique we deployed can be explained on a whiteboard in ten minutes. What mattered was not the algorithms but the design choices that made the model's reasoning legible to the humans who had to live with its recommendations and that left them, not the model, in charge of the plan. Surfacing the smallest set of mutually inconsistent commitments turned contradictions into named choices a planner could deliberate over. Pricing relaxed constraints turned absolutes into priced tradeoffs. Returning a menu of near-optimal plans turned prescriptions into operationally distinct options. Restructuring around prioritized objectives captured the fact that organizations have priorities, not a single number. Parameter exploration gave production planners back the exploratory freedom that Excel had always provided, within a framework that maintained global coherence. Each of these decisions had the same goal: not to make the model faster or more elegant, but to keep the planner accountable for the plan and equipped to defend it.

Woolsey and the practitioner tradition he championed always insisted that operations research fails most often not from bad mathematics but from bad anthropology. The obvious answer was always there. Getting the organization to a place where it could accept it, trust it, and act on it: that was the real optimization problem. It was the most difficult project I have worked on, and without question, the most valuable.

Endnote

[1] The firm and certain operational details have been anonymized in keeping with standard practice for practitioner case studies. Reported impact figures are aggregate and approximate; granular financials, supplier identities, plant locations, and product names are withheld for the same reason.

Adam DeJans Jr. is a senior research scientist at Amazon specializing in supply chain optimization, simulation, and decision-making under uncertainty. His applied work has supported large-scale planning and operational decisions across online retail, automotive, technology, defense, and pharmaceutical supply chains. He focuses on translating advanced analytics and optimization methods into practical, economically meaningful decisions that improve real-world performance and resilience.