# THE CASE FOR SIMULATION IN THE EMERGENCY DEPARTMENT

*Why 189 Hospitals Are Already Failing the Test — and Why the Fix Doesn't Require New Beds, New Staff, or New Money*

**Patrick Dugan, LSSMBB — Discrete-Event Simulation for Emergency Department Flow**

Every emergency department in the country runs on the same handful of numbers: how long a patient waits to see a doctor, how long they are in the department in total, how many walk out before being seen, and how long an admitted patient waits for an inpatient bed. CMS publishes all four, for every hospital in the country, every quarter.

A screen of CMS's own published Care Compare data across 2,558 hospitals found that 189 of them — nearly one in thirteen — are flagged as underperforming on two or more of these measures at once, and thirty of those on all three. This is not a story about a handful of struggling outliers. It is a system-wide pattern, sitting in plain sight inside data CMS already collects and already publishes.

The standard responses — hire more nurses, build more beds, buy new patient-flow software — are expensive, slow, and, as this paper will show, frequently do not fix the actual problem, because the emergency department is not the kind of system those fixes were built for.

The fix is not more capital. It is a different method for deciding what to do with the capital already there: discrete-event simulation, run as its own management discipline rather than a one-off IT project, that tests an intervention against a department's real behavior before a dollar is spent or a decision is announced. And because of how AI has changed the cost of building these models, that method is available now — not as a future aspiration, but as something a small community hospital can afford to start this quarter.

| | |
|---|---|
| **189**<br>of 2,558 hospitals screened flag risk on 2+ of 3 CMS ED measures | **~$4**<br>of released value sits unconverted for every $1 currently booked as savings |

## I. The Evidence Is Not Anecdotal

Every year, CMS publishes hospital-level medians for the metrics that describe how an emergency department actually functions for the patients who walk through its doors: OP-18a, the median time from arrival to departure for all ED patients; OP-18b, the same measure restricted to discharged patients, whose gap from OP-18a is a reasonable proxy for how much time boarding admitted patients is adding to the department; and OP-22, the percentage of patients who leave without being seen at all.

A screen of the most recent full year of this data found 189 hospitals — nearly 1 in 13 of everything CMS scored — flagged as being in the worst 10% nationally on two or more of these three measures simultaneously. Thirty of those are flagged on all three. The hospitals on this list run the full range of size, ownership, and geography: rural critical-access facilities and large urban systems alike.

The point of this data is not to single out any one hospital. It is to establish, with numbers CMS itself already collects and already makes public, that ED flow failure is common, measurable, and largely invisible to the people trying to fix it until it shows up on a public scorecard. Whatever the actual constraint is inside these 189 departments, it usually is not obvious from the outside — and in most cases, it is not obvious from the inside either.

## II. Why Intuition and Simple Math Can't Fix It

Every ED faces the same underlying mathematical problem, whether or not anyone on staff would describe it that way: patients arrive randomly, need random combinations of tests and beds, and compete for a fixed, shared set of resources. Two properties of that problem — being stochastic and being nonlinear — are why ordinary planning tools consistently get the wrong answer, even when the person using them is experienced and careful.

Stochastic simply means random, in a specific sense: you can know the average number of patients who will arrive on a Tuesday and still be unable to say what will happen at 2pm on any particular Tuesday. Averages describe the whole distribution badly. A department staffed to its “average” daily volume will, by definition, be overwhelmed roughly half the time, because half of all days run above average — and the busy half is exactly when boarding, walkouts, and diversion actually happen. Planning to the average is planning to fail on your worst days, systematically.

Nonlinearity is the second property, and it is the less intuitive one. In a linear system, a small additional load produces a small additional strain. Emergency departments do not behave that way. Alexander Kolker's own discrete-event research on ED patient flow found that the relationship between length of stay and ambulance diversion is close to flat — genuinely negligible — right up until admitted patients cross roughly a six-hour length of stay and discharged patients roughly five hours. Past that line, diversion accelerates sharply. Two hospitals running at what looks like “almost the same” length of stay can be living in completely different worlds — one comfortably below the cliff, one just past it — and nothing about their average numbers says which side they're on.

This is precisely the kind of pattern a spreadsheet, a staffing ratio, or a comparison to a “similar” hospital cannot find, because those tools describe averages and straight lines. A discrete-event simulation is built to reproduce the actual random arrivals and the actual shared resources, run forward thousands of times, and show exactly where a department's cliffs are — before anyone finds them the hard way, on a real Tuesday afternoon.

The same blindness shows up sideways as well as at the extremes. A hospital struggling with boarding or walkouts will almost always look first at ED capacity — more beds, more ED staff — because that is where the symptom is visible. But CT scanners, ICU beds, OR block time, and imaging reads are frequently shared across demand streams nobody is tracking together, and the actual constraint often lives somewhere else entirely. Kolker's own modeling of ICU bed diversion found the real driver was not ICU capacity at all — it was how elective surgery scheduling was loading the same bed pool emergency demand also depended on; capping elective volume at four cases a day nearly eliminated the diversion, without adding a single bed. Add capacity to the department everyone is staring at, and the wait simply relocates to wherever the real shared constraint lives — invisible until someone models the whole system, not just the part with the visible symptom.

## III. The Missing Discipline: Management Engineering

Hospitals already have a rigorous, evidence-based discipline for clinical decisions. Medicine does not run on intuition and educated guesses — it runs on trials, protocols, and continuously updated evidence, and no hospital would accept “this seemed like it should work” as the standard for a treatment decision.

Operational decisions inside the same hospital — how many beds to run, how to schedule elective surgery against emergency demand, how to staff a shift, whether to add a discharge lounge — are, in most hospitals, still made the way medicine itself was practiced before evidence-based methods existed: by experience, by intuition, and by copying whatever a similar-looking hospital seems to be doing.

Alexander Kolker — the same researcher whose ED and ICU findings anchor the mechanics described above — has spent much of his career arguing that this gap is the actual problem, not a footnote to it. His book, Healthcare Management Engineering: What Does This Fancy Term Really Mean?, makes the case that hospitals need Management Engineering recognized and staffed as its own function: the disciplined application of operations research — queueing theory, discrete-event simulation, optimization, forecasting — to exactly the operational questions listed above, held to the same standard of rigor clinical decisions already are. Not a one-off analysis run before a strategic planning retreat. A standing capability, the operational equivalent of a hospital's own clinical quality function.

Kolker is blunt about how wide the current gap is: by his own assessment, relatively few healthcare administrators are equipped to “think analytically about healthcare delivery as a system.” That is not a criticism of administrators' competence — it is a statement about training and organizational design. Hospital leaders are trained in clinical care, in finance, or in general administration; almost none are trained to reason about queues, variability, and shared-resource contention, because nothing in a typical career path teaches it. The result is a hospital that would never let an unproven treatment near a patient, making capacity and scheduling decisions on exactly that basis every day.

Kolker's own prescription is realistic rather than idealistic: hospital leaders do not need to become operations researchers themselves. They need to understand why the old way — intuition, simple ratios, benchmarking against a “similar” hospital — is unreliable for precisely the stochastic, nonlinear reasons described above, and to treat Management Engineering as a function worth having on staff or on retainer, not as an occasional luxury.

## IV. From Minutes to Dollars Finance Can Trust

A simulation that reports “boarding time fell by 45 minutes” is true and useless in the same breath, if it stops there. Finance departments do not budget in minutes. The gap between an operational finding and a number a CFO will actually put in next year's budget is where most improvement projects quietly die — not because the improvement was not real, but because nobody could show where it landed in the ledger.

The Institute for Healthcare Improvement's Martin, Neumann, Mountford, Bisognano and Nolan built the accounting discipline that closes that gap. Every dollar of claimed benefit is sorted into one of two categories: dark green dollars, savings traceable to an actual line in the P&L — a shift not filled, an agency contract not renewed, a capital purchase deferred, a walkout who is instead treated and billed — and light green dollars, real efficiency nobody converts into any of those things, which therefore evaporates into a marginally calmer waiting room and an unchanged budget next year.

Applied honestly, this discipline is uncomfortable, which is exactly why it is credible. In one applied case study — a calibrated simulation of a 113-bed community hospital's emergency department — an intervention scenario that looked strong in raw simulation output produced roughly $1.02 million in identifiable released capacity, of which only about $233,000 could honestly be booked as dark green in year one. That is not a disappointing result. It is the single most useful number in the analysis: for every dollar this kind of work can currently promise a CFO, there are roughly four more dollars of real, measured value sitting on the table, waiting on specific staffing, scheduling, and capital decisions that belong to hospital leadership, not to the model. The simulation's job is to make those decisions visible and quantified. The finance discipline's job is to make sure whatever gets claimed survives contact with an actual audit.

## V. Why This Is Solvable Today, Not Someday

None of the methods described above are new. Kolker has published this kind of ED and ICU work for over fifteen years. What has changed — recently, and substantially — is the cost of doing it.

Building and calibrating a discrete-event model of an emergency department used to require a specialized analyst, an expensive software license, and months of build time. That is exactly why almost none of the 189 hospitals flagged above, particularly the smaller and rural ones among them, have ever had this capability in-house or been able to afford it as an outside engagement. Simulation has functioned, in practice, as a large-hospital, large-budget tool.

AI-assisted development changes that math directly. A calibrated baseline model — built first from a hospital's own published CMS data, then sharpened with its actual bed counts, staffing, and volume —

can now be built and a battery of already-validated intervention scenarios run against it inside days rather than quarters. None of it requires new construction, new headcount, or risk to a live patient: every scenario runs first inside the model, at zero clinical risk, before a single shift changes or a single dollar is spent. A hospital does not need to guess whether a discharge lounge, an added inpatient bed, a bed-coordinator role, or a scheduling change to a shared resource like a CT scanner will help — a scenario library already built and validated on a real reference hospital can test dozens of these ideas against that hospital's own recalibrated numbers before any of them is tried live.

It is worth being precise about what AI actually contributes here, because it is not the whole answer by itself. AI accelerates the building, calibrating, and testing work — the part that used to take months — but it does not replace the underlying method. Pattern-matching AI, trained on historical data, is very good at recognizing what has already happened; it has no reliable way to tell a hospital what would happen under a change nobody has tried yet, because the historical record contains no examples of that hypothetical. A discrete-event model can answer exactly that question, because it is built from the department's actual mechanics — its beds, its queues, its shared resources — not from a correlation buried in past data. The realistic picture is a stack, not a substitute: AI is what now makes the mechanistic model fast and affordable to build; the mechanistic model is what makes the answer trustworthy once it is built.

## VI. What's Actually Stopping ED Administrators

If the data is public, the method is proven, and the cost has fallen this far, the honest question is why more of the 189 hospitals are not already doing this. The barriers left are not really technical or financial anymore. They are cultural and organizational, and naming them plainly is the first step to getting past them.

The disciplinary gap comes first, because it is upstream of everything else. Most ED and hospital leaders were trained clinically or administratively, not analytically — Management Engineering, per Kolker's own diagnosis above, looks like someone else's specialty rather than a core management tool, so it never makes it onto the list of options considered.

The reflex fix comes second. When boarding or walkout numbers spike, the intuitive response is to ask for more beds or more staff, because that is where the symptom is visible — and asking that question pre-empts the harder, more useful one: where is the actual constraint, and is it even inside this department at all.

Finance skepticism is third, and it is earned rather than irrational. Improvement teams have a long history of claiming savings finance could not trace to an actual budget line, so finance has learned to discount improvement projects by default — which means operations leaders stop bringing forward proposals they cannot back with numbers finance will accept, even when the underlying idea is sound.

Fourth is a genuine fear of being publicly wrong. A model does not only validate popular ideas — it can just as easily show that a change everyone already assumed would help, like a dedicated bed coordinator, actually makes boarding worse once tested against the whole department. That is a real risk to a leader who has already floated the idea internally, and it discourages some administrators from running the test at all rather than risk an uncomfortable, public answer.

Fifth is the outdated belief that this still requires a large budget and a long timeline — true five years ago, no longer true today, and worth saying explicitly rather than assuming administrators already know it.

And sixth is simply not knowing where to start, often compounded by the belief that a hospital needs perfect internal data before any of this is worth attempting. It does not. Publicly available CMS data is enough to build a directionally useful first model; a hospital's own numbers only make it sharper from there.

None of these six is a reason to wait. Every one of them is addressed by doing the work rather than debating it — building the model, running the scenarios, and letting the results, not the argument, make the case.

## Closing

The emergency department is not a simple system, and it was never going to yield to simple tools. It is stochastic, it is nonlinear, and it runs on shared resources that nobody in any one department can see the whole of. Treating it as if intuition, spreadsheets, and comparison to a similar-looking hospital were adequate has produced exactly the result CMS's own data shows: 189 hospitals, and likely many more not captured by any public measure, quietly losing ground on the numbers that describe whether their emergency department is actually working.

The fix has existed in the academic and operations-research literature for well over a decade. What is new is that it no longer requires a large hospital's budget or a large hospital's timeline to use it.

***Model first. Spend second.***

**Patrick Dugan, LSSMBB**
Discrete-Event Simulation for Emergency Department Flow
patrick.dugan99.9@gmail.com

## Sources

Kolker A. Process Modeling of Emergency Department Patient Flow: Effect of Patient Length of Stay on ED Diversion. Journal of Medical Systems. 2008;32(5):389-401.

Kolker A. Process Modeling of ICU Patient Flow: Effect of Daily Load Leveling of Elective Surgeries on ICU Diversion. Journal of Medical Systems. 2009;33:27-40.

Kolker A. Healthcare Management Engineering: What Does This Fancy Term Really Mean? The Use of Operations Management Methodology for Quantitative Decision-Making in Healthcare Settings. Springer; 2012.

Martin LA, Neumann CW, Mountford J, Bisognano M, Nolan TW. Increasing Efficiency and Enhancing Value in Health Care: Ways to Achieve Savings in Operating Costs per Year. IHI Innovation Series white paper 19. Cambridge, MA: Institute for Healthcare Improvement; 2009.

Schreyer KE, Martin R. The Economics of an Admissions Holding Unit. Western Journal of Emergency Medicine. 2017;18(4):553-558.

Centers for Medicare and Medicaid Services. Timely and Effective Care — Hospital. Provider Data Catalog, Care Compare (screened dataset, 2,558 hospitals, most recent full reporting year).