



A Pulse Is Not Enough. The Layers Are the Point.

Autonomy Platform · Public Blog

P E R S P E C T I V E

© 2026 Azirella Ltd. All rights reserved worldwide.

Read more at azirella.com/blog

A Pulse Is Not Enough. The Layers Are the Point.

PwC's 'From Process to Pulse' gets the diagnosis exactly right: the bullwhip is a latency problem. But the cure is not one fast loop. It is the right loops, at the right horizons, each acting on a calibrated band.

Steve Puricelli and his PwC colleagues [wrote this month](#) about agentic S&OP, *from process to pulse*, and the piece gets the hardest thing right. The bullwhip effect was never really a forecasting problem. It is a latency problem: distortion that compounds in the gap between when the market signals and when the organisation responds. Their structural paradox is the cleanest sentence anyone has written about it: supply chains can move materials in hours but take weeks to make decisions.

I agree with almost all of it. The shift from a calendar to a signal. The elevation of governance rather than its elimination. The explicit decision architecture about which decisions humans make, which agents recommend, and which agents execute. The warning that autonomy

layered on noisy signals or misaligned KPIs amplifies the bullwhip rather than taming it. That last point is the one most vendors skip, and PwC put it in writing.

So this is not a rebuttal. It is the part the genre of a survey article cannot reach: the mechanism. And the mechanism turns on two words the piece uses lightly and one word it discards too quickly. The two it uses lightly are *pulse* and *guardrails*. The one it discards is *forecasting*.

A single pulse reaches one horizon.

Taming the bullwhip needs three levers the single-loop framing under-reaches: layers, an execution tier, and probabilistic planning.

A nervous system is not one pulse


The article's own best metaphor is the cognitive nervous system. I want to hold it to that metaphor, because a nervous system is precisely *not* a single pulse. It is a layered hierarchy of loops running at different speeds. A spinal reflex fires in milliseconds without consulting the cortex. The cortex deliberates over seconds and minutes. Long-run learning happens over days. The genius of the design is not that everything runs fast. It is that each loop runs at the horizon its decisions belong to, and the layers coordinate without the slow layer becoming a bottleneck on the fast one.

Look at PwC's own maturity path. It tops out at *Agentic Orchestration*, decision speed measured in *minutes*, autonomous actions within guardrails. Minutes is the right cadence for the network and the S&OP horizon. It is the wrong cadence for the tier where the bullwhip actually does its damage.

Because the amplification does not happen in the boardroom. It happens at the bottom of the chain, in ordering and replenishment, where each node reacts with delay and partial information. Lee, Padmanabhan and Whang, the 1997 paper PwC cites, located the distortion in exactly those operational reactions. A pulse that lives at the planning layer, however continuous, does not reach the replenishment decision being made thousands of times a day at thousands of locations. That decision needs its own loop, at its own speed.

A PULSE IS NOT ENOUGH. THE LAYERS ARE THE POINT.

The bullwhip is a latency problem. The cure is not one fast loop: it is the right loops, at the right horizons.



The third lever PwC discards: probabilistic planning.

Every tier decides on a calibrated band (P10 / P50 / P90), not a point. The calibrated likelihood is what draws the line between Automate, Inform, and Inspect: the mechanism "bounded autonomy" actually requires. Guardrails without calibration are a wish.

Azurella · the layered decision substrate beneath agentic S&OP

Down the stack flows context, guardrails, and targets, never a plan as a directive. Each tier writes its own plan inside the envelope it inherits, and sends signals back up. PwC's continuous pulse is the tactical loop. The execution tier is the reflex arc a planning-layer pulse never reaches.

This is why the right architecture is not one agent and not one cadence. It is a layered policy stack: **Strategic** sets the network envelope weekly; **Tactical** balances demand and supply across the network daily; **Operational** runs the single site hourly; and **Execution** is a hive of narrow agents acting in real time, per product, per site, in milliseconds.

The pulse PwC describes is real, but it is the *tactical* pulse. The execution layer is the reflex arc, and a planning system without it leaves the most bullwhip-prone decisions exactly where they were: slow, local, and uncoordinated.

The layers are the cure, not the obstacle

Here is the part that matters most, and it is counterintuitive. Layering is not the legacy thing agentic S&OP replaces. Layering, done correctly, is the bullwhip cure.

The bullwhip is information distortion propagating downstream. The classic failure is that each tier passes its *plan* to the next as a directive, and every handoff adds delay, reinterpretation, and local buffering. That cascade of plans-as-orders is the distortion.

So we do not cascade plans. Down the stack flows context, guardrails, and targets. Never a plan as a command. Each tier generates its own plan of record inside the envelope it inherits, judged by its own balanced scorecard, and sends signals back up. Strategic tells Tactical the bounds, not the buffer. Tactical tells Execution the targets, not the order quantity. The execution agent decides the order quantity itself, in real time, against live state.

That is structurally anti-bullwhip. A distorted directive cannot propagate downstream as a directive, because no tier is taking a directive. What propagates is intent, and intent does not amplify the way a buffered, re-buffered, re-interpreted order does. PwC's *synchronising intent* is exactly the right phrase. The layered envelope is how you actually do it. When two tiers or two functions do collide, the conflict belongs on a surface where it is rendered visible and resolvable, not silently propagated. We call that the [Decision Stream](#).

The word PwC discards too quickly

The article says, twice, that this is *not simply better forecasting* and *never really a forecasting problem*. As rhetoric against calendar-bound forecast-and-freeze planning, fair. As architecture, it is an overcorrection, and it leaves a hole right under the load-bearing claim of the whole piece.

Because the piece asks for *bounded autonomy*, *trustworthy automation*, and *speed where the signal is reliable*. Every one of those phrases presumes a quantity the article never names: a calibrated measure of how reliable the signal is. You cannot set a bound on autonomy without knowing where the boundary is. You cannot grant speed where the signal is reliable without a number for reliability. Guardrails without calibration are a wish.

The answer is not less forecasting. It is *probabilistic* forecasting. Forecast the band, not the point. Plan against P10, P50 and P90, and carry a calibrated likelihood on every decision the agent makes. Conformal prediction gives that band a distribution-free coverage guarantee and recalibrates as outcomes arrive, which is what keeps it honest through the regime shifts and demand shocks that break a point forecast. (We made the wider case for this in [Stop Planning Against Averages](#).)

Then the likelihood becomes the thing that decides the mode. An agent automates silently when calibrated confidence is high. It informs a human when confidence is low and the stakes are high. It holds for inspection when there is time to think. That is the [AI·IO·ML operating model](#): Automate, Inform, Inspect, Override, and underneath them Measure and Learn. It is not a slider on a dashboard. It is a two-dimensional risk-appetite surface, per decision class, on the urgency-and-likelihood plane, owned by the customer and tuned by outcomes.

This is also the precise answer to PwC's own *What Can Go Wrong*. Agents amplify the bullwhip when they act on noisy signals.

Calibration is the defence: an agent that knows its band is wide does not act with false confidence, it informs. The probabilistic layer is what turns *bounded autonomy* from a governance slogan into a mechanism.

Three levers, one diagnosis

PwC's diagnosis is right and I would not change a word of it. The bullwhip is latency. Decision velocity is the prize. But a single pulse reaches one horizon. Taming the bullwhip needs three levers the single-loop framing under-reaches:

1. **Layers**, so each decision runs at its own horizon and the slow loop never bottlenecks the fast one.
2. **An execution tier**, because the amplification lives at the bottom, in real-time replenishment, not in the monthly review.
3. **Probabilistic planning**, because calibrated likelihood is the only thing that can actually draw the line between automate, inform, and inspect that *bounded autonomy* requires.

The future of S&OP is not a monthly process. PwC is right about that. But it is not a single continuous pulse either. It is a nervous system: layered loops, a reflex tier that acts in milliseconds, and a calibrated sense of its own uncertainty so it knows, at every level, how much to trust itself before it acts.

That is the architecture underneath the pulse.

See Autonomy in action

Walk through how Autonomy models, executes, monitors, and governs supply chain decisions with autonomous AI agents.

[See It Live](#)