E2 MUD · MANUFACTURING

The Suggestion Gap: A Million Ideas a Year vs. Silence

Essay 3 of "The Evidence They Can't Ignore" — What Toyota built, what GM couldn't copy, and the invisible cost on every Taylorist balance sheet.

The first two essays in this series established two facts. First, the NUMMI experiment proved that the same workers, in the same building, under a different operating system, went from GM's worst plant to GM's best. Second, Frederick Taylor's published writings showed that the dominant Western manufacturing model was explicitly designed to remove intelligence from the people doing the work.

This week I want to show you what happens when you flip the switch — when you actually build a system that asks frontline workers to think. The data comes from the most sustained experiment in deployed frontline intelligence in industrial history: the Toyota suggestion system.

The Numbers

Toyota receives more than one million improvement suggestions from its workforce every year.

Let that register. Not one million complaints. Not one million grievances. One million

ideas

— discrete proposals for how to do the work better, submitted voluntarily by the people doing the work, processed through a formal system, and implemented at a rate that has historically hovered around 90 percent or higher.

That implementation rate is the key. Toyota doesn't just collect suggestions. It acts on them. The system is designed so that the vast majority of ideas — particularly small, localized improvements submitted by individual workers or small teams — are evaluated and implemented quickly, often within the work group itself. The worker sees the result of their thinking reflected in the actual work process, usually within days.

The math works out to roughly one suggestion per worker per week. Not per year. Per week. This means the average Toyota production worker is generating, articulating, submitting, and seeing implemented approximately 50 improvement ideas every year. Over a career, that's thousands of contributions to the operating system — each one a small proof that the worker's intelligence matters.

The Contrast

Now consider what the same metric looked like at General Motors during the decades when Toyota was building this system.

GM had a suggestion program. Most large Western manufacturers did. But GM's system required that each suggestion demonstrate a minimum calculated return on investment — reportedly around $20 per suggestion — just to justify the administrative cost of processing it. The system was designed to filter out "small" ideas, on the theory that only ideas with measurable financial impact were worth management's time.

Think about what that filter communicates to the worker. It says: your idea has to be worth at least $20 to us before we'll even read it. If your suggestion is about ergonomics, or sequencing, or a better way to position a part that saves two seconds per cycle, or a safety observation that doesn't have a dollar sign attached — don't bother. The system doesn't want it.

The result was predictable. Participation rates at GM and most Western manufacturers were a fraction of Toyota's. Not because the workers had fewer ideas. Because the system told them their ideas weren't wanted.

What the Suggestions Actually Contain

The common misconception about Toyota's suggestion system is that it produces a torrent of trivial ideas — a million sticky notes saying "put the trash can closer to the workstation." That misconception reveals the bias. It assumes that because the ideas come from "operator-level people," they must be small.

Some of them are small. Many of the best ones are. A two-second cycle time reduction, applied across three shifts running 500 units per shift, 250 days per year, compounds into thousands of hours of recovered capacity annually. Toyota understood something that Taylor's system was designed to prevent workers from demonstrating: the people doing the work see things that no one else can see, because no one else is in that position, at that angle, with that frequency, for that duration.

But the suggestions aren't only small. The system produces ideas across a spectrum — from workstation-level ergonomic adjustments to process redesigns that affect entire production lines. The critical insight is that the

system itself

develops the worker's problem-solving capability over time. A worker who starts by suggesting a better hook placement in year one is, by year five, identifying root causes of quality variation and proposing countermeasures. The suggestion system isn't just a collection mechanism. It's a development engine.

This is why Toyota's cumulative suggestion value over seven decades runs into the billions. Not because any single idea was worth billions, but because the system produced millions of ideas that compounded — and, more importantly, it produced millions of

workers who got better at thinking

with every suggestion they submitted.

The System Behind the Suggestions

The suggestion system doesn't operate in isolation. It works because it sits inside an operating model that supports it at every level:

Training to see:

Toyota workers are trained in standardized problem-solving methods — the ability to observe a process, identify deviation from the standard, analyze the root cause, and propose a countermeasure. This isn't natural talent. It's taught capability. The system invests in making workers better at thinking before it asks them to think.

Standards as baselines:

As I discussed in the Taylor essay, Toyota's standardized work is a floor, not a ceiling. The current standard is the documented best-known method — and it exists precisely so that someone can identify a better method and improve it. Without the standard, you can't see the deviation. Without seeing the deviation, you can't generate the suggestion. The standard and the suggestion system are two halves of the same design.

Rapid feedback:

Ideas are evaluated quickly and implemented quickly. The worker sees the impact. This closes the loop and reinforces the behavior. In most Western suggestion systems, ideas disappear into a bureaucratic review process and emerge months later with a form letter. The worker learns that submitting ideas is a waste of time. The behavior extinguishes.

Recognition, not just reward:

Toyota's system emphasizes recognition — the acknowledgment that the worker's thinking contributed to the organization's improvement. Financial rewards exist but are modest. The primary currency is respect: the organization saw your idea, valued it, and changed the work because of it.

The Suppression Tax

I want to introduce a concept that will recur throughout this series: the

suppression tax.

The suppression tax is the economic cost of operating a system that prevents frontline workers from contributing their intelligence. It's not a line item on any balance sheet. No CFO reports it. No consultant measures it. But it is real, it is enormous, and it compounds every day the system runs.

Consider: if Toyota gets a million suggestions per year and implements the vast majority, and if the average suggestion produces even modest value, the cumulative economic contribution of deployed frontline intelligence over decades is staggering.

Now consider the inverse. Every Western manufacturer running a Taylorist operating model — where the suggestion system is either nonexistent, bureaucratic, or designed with filters that suppress participation — is

forgoing

that value every year. The ideas exist. The workers have them. The system doesn't ask, or asks and doesn't listen, or listens and doesn't act.

That's the suppression tax. It's the delta between what your workforce knows and what your operating system allows them to contribute. At Toyota, that delta is small — the system is designed to minimize it. At most Western manufacturers, that delta is enormous — and invisible, because you can't measure ideas that were never submitted.

The skeptic says: “Our workers don’t have ideas.” The data says: Toyota’s workers have a million of them per year. The difference isn’t the workers. It’s the system.

The GM Paradox

Here is the part that should make every executive uncomfortable.

General Motors had a front-row seat to Toyota’s suggestion system for 25 years through the NUMMI joint venture. GM managers visited Toyota City. They studied the system. They wrote reports about it. They saw the data.

And they couldn’t replicate it.

Not because the system was secret — Toyota was remarkably open about sharing its methods. Not because the workers were different — NUMMI proved the same American workers could operate inside the Toyota system. But because replicating the suggestion system required changing everything underneath it: the relationship between management and labor, the purpose of standardized work, the speed of feedback loops, the willingness to invest in training workers to think, and — most fundamentally — the belief that frontline workers

have

intelligence worth deploying.

GM couldn’t import the suggestion system as a standalone program because it wasn’t a program. It was an output of an entire operating philosophy. Grafting a suggestion box onto a Taylorist operating model produces what every Western manufacturer has experienced: initial enthusiasm, declining participation, management frustration, and eventual abandonment. The suggestion box didn’t fail. The operating model underneath it made success impossible.

The Lesson for What Comes Next

The suggestion gap is not an anecdote. It is a controlled experiment, running continuously for over 70 years, comparing what happens when an operating system asks for frontline intelligence versus what happens when it doesn’t.

The results are unambiguous. When you build a system that trains workers to see problems, gives them a method to articulate solutions, processes their ideas quickly, implements the good ones visibly, and recognizes the contribution — you get a million ideas a year. When

you don’t, you get silence. And then you conclude that the workers have nothing to say.

Dr. Venki Padmanabhan

has 36 years of global manufacturing leadership experience, including executive roles at GM, Chrysler, Mercedes-Benz, Royal Enfield, and Ather Energy. He holds a PhD in Industrial Engineering from the University of Pittsburgh. He is the author of the forthcoming

Already Paid For: Why Unlocking Frontline Intelligence Beats Automating Workers Away

and co-founder of the Capability Capital Institute. He writes The Long Game at thelonggameforall.substack.com.

Venki Padmanabhan is a co-founder of the Capability Capital Institute and the author of the forthcoming Built to Extract and Already Paid For (Capability Capital Press). He writes at thelonggameforall.substack.com.

Written with AI assistance. The argument, the judgments, and the floor testimony are the author’s own.

Published 2026-03-29 · thelonggameforall.substack.com · Capability Capital Institute