EPISODE 8 TRANSCRIPT: "OPTIMIZING METRICS, DESTROYING SYSTEMS"

THINGS OVERHEARD AT THE COFFEE BAR

Episode 8: Optimizing Metrics, Destroying Systems

Runtime: ~49 minutes

[COLD OPEN - 0:00]

[AMBIENT SOUND: Coffee shop, morning rush]

DANA: So my boss just implemented a new productivity tracking system.

FRIEND: Oh no.

DANA: Yeah. It measures keystrokes per hour, time in each application, response time to emails. And here's the thing—my productivity score went way up.

FRIEND: That's good, right?

DANA: My score went up. My actual work got worse. Because now I'm optimizing for keystrokes instead of quality. I'm responding to emails instantly instead of thinking about good answers. I'm switching between applications constantly to show activity instead of doing deep work.

FRIEND: But the metric says you're more productive.

DANA: The metric is measuring the wrong thing. But my boss only sees the metric. So I'm getting praised for becoming worse at my job.

FRIEND: That's insane.

DANA: It's worse than insane. Everyone on my team is doing the same thing. We're all gaming the metrics. We're all becoming worse at our actual jobs while appearing more productive. And the whole organization is celebrating how much more efficient we've become.

FRIEND: What happens when the work quality drops?

DANA: They'll measure that. Create a quality metric. Then we'll game that one too. And around and around we go, destroying everything we're trying to improve.

[SOUND FADES]

[INTRO - 1:45]

HOST: I'm Alex Chen, and this is Things Overheard at the Coffee Bar.

There's a principle in economics and sociology called Goodhart's Law. It states: "When a measure becomes a target, it ceases to be a good measure."[1]

Here's what that means:

You want to improve something—let's say education quality. So you measure it—standardized test scores. Then you make the measure a target—tie funding to test performance. And immediately, the measure becomes useless. Because now everyone optimizes for test scores instead of actual education. They teach to the test. And test scores go up while education quality goes down.

This happens everywhere. In every domain. At every scale.

Hospitals optimize for short patient stays and readmission rates—so they discharge people too early, leading to worse outcomes.[2]

Police departments get less-lethal weapons to reduce fatalities—so officers use force more liberally, leading to more deaths.[3]

Recycling programs install high-tech sorting robots that achieve 99% purity—so people stop sorting their own recycling carefully, total capture rate drops, and more waste goes to landfills.[4]

Cities pay bounties for dead cobras to reduce the cobra population—so people start breeding cobras for the bounty, and when the program ends, they release them, making the problem worse than before.[5]

This is the pattern: **We optimize the metric. We destroy the system.**

Today we're exploring:

- Why we've lost the capacity to see whole systems
- How optimization of parts destroys wholes
- What traditional knowledge knew about system thinking that we've forgotten
- Why every solution creates new problems
- And whether we can learn to think in systems again

This is "Optimizing Metrics, Destroying Systems."

[THEME MUSIC - 4:00]

[ACT ONE: THE GOODHART PROBLEM - 4:30]

HOST: Let me start with the cleanest example of Goodhart's Law in action.

The Soviet nail factory story.[6]

A Soviet factory produces nails. Central planners want to measure productivity. First metric: number of nails produced.

Result: Factory produces millions of tiny, useless nails. Technically nails. Useless for construction. But the metric is met.

Planners adjust. New metric: weight of nails produced.

Result: Factory produces enormous, heavy nails. Technically nails. Also useless for construction. But the metric is met.

The problem isn't the metrics themselves. Weight and quantity are both relevant to nail production. The problem is that **no single metric can capture the whole system goal**: producing useful nails of various sizes appropriate for different construction needs.

But we keep trying to find the one perfect metric that will optimize everything.

I called Dr. Amara Thompson—the labor economist from episode six—who also studies organizational systems.

DR. THOMPSON: Goodhart's Law is everywhere once you start looking for it. Every time you create a target metric, people will optimize for that metric at the expense of the actual goal.

HOST: Why can't we just measure the actual goal?

DR. THOMPSON: Because actual goals are complex and unmeasurable. "Produce useful nails" isn't a number. "Provide good education" isn't a number. "Deliver quality healthcare" isn't a number. So we pick proxies—test scores, patient throughput, readmission rates. And then we optimize the proxy while destroying the thing we actually cared about.

HOST: Why don't we just measure multiple things?

DR. THOMPSON: We try. But the more metrics you add, the more complex the gaming becomes. People figure out how to optimize all the metrics simultaneously while still destroying the actual goal. Or they focus on whichever metric has the strongest incentive attached.

HOST: So there's no solution?

DR. THOMPSON: Not within a measurement framework. If you're using metrics to manage, you'll always get Goodhart effects. The only solution is qualitative judgment by people who understand the whole system. But that doesn't scale. And it can't be automated. So we keep trying to measure our way out of it.

[MUSIC TRANSITION - 8:00]

[ACT TWO: THE RECYCLING ROBOT - 8:30]

HOST: Let me give you a real, current example.

Municipal recycling programs are installing high-tech sorting robots powered by AI and machine vision. These robots can identify and sort materials with 99% accuracy—way better than humans.[7]

This should be good, right? Better sorting means more material actually gets recycled instead of contaminating batches and ending up in landfills.

But here's what actually happened.

I talked to Marcus Chen, who works for Richmond's waste management department.

MARCUS C: When we installed the robots, we thought contamination rates would drop. The robots are amazing—they can identify a yogurt cup versus a butter tub versus a plastic bag, sort them correctly, achieve 99% purity in each stream.

HOST: So contamination dropped?

MARCUS C: Contamination of what makes it to the facility dropped. But total capture rate—the percentage of recyclable material that actually gets recycled—also dropped. By about 15%.

HOST: Why?

MARCUS C: Because people stopped sorting. Before the robots, we had education campaigns: "Rinse your containers. Remove lids. Don't bag recyclables." People knew their sorting mattered. They were part of the system.

Then we installed robots and marketed them as "sorting for you." And people thought: "Great, the robots will handle it. I don't need to be careful." So they stopped rinsing containers. They started throwing everything in one bin. They included non-recyclables thinking the robots would catch them.

HOST: But the robots did catch them.

MARCUS C: The robots caught what made it to the facility. But greasy pizza boxes don't make it through the collection truck compactor properly—they contaminate everything they touch. Bagged recyclables can't be sorted—the robots can't see inside. Liquids left in containers leak during transport.

So even though the robots achieve 99% purity on what reaches them, less material reaches them in usable condition. We optimized the sorting metric and destroyed the recycling system.

[COFFEE SHOP AMBIENCE - 11:30]

HOST: This is a perfect Goodhart example. The metric was "sorting accuracy." The goal was "maximize recycling." They're not the same thing.

Sorting accuracy optimizes for purity at the facility. But recycling is a whole-system process: waste generation → separation → collection → transport → sorting → processing → remanufacturing.

The robots optimized one node in the system. And broke the whole flow.

Marcus C:

MARCUS C: The lesson we learned—painfully—is that when you automate part of a system, you change the whole system in unpredictable ways. People's relationship to recycling changed. They went from active participants to passive consumers. And that shift broke things the robots couldn't fix.

HOST: Could you fix it?

MARCUS C: We're trying. Re-educating people that they still need to sort carefully even though there are robots. But once people outsource responsibility to technology, it's really hard to get them to take it back.

[MUSIC TRANSITION - 13:00]

[ACT THREE: THE LESS-LETHAL WEAPONS PARADOX - 13:30]

HOST: Here's a darker example.

In the 1990s and 2000s, police departments across the U.S. adopted "less-lethal" weapons—tasers, pepper spray, rubber bullets—as alternatives to firearms.[8]

The goal: reduce police shootings and deaths.

The metric: deployment of less-lethal options instead of firearms.

The result: More people died.

[PAUSE]

I talked to Dr. Jennifer Hartwell—the anthropologist from episode two—who studies policing and use of force.

DR. HARTWELL: The logic seemed sound. Give officers an option between talking and shooting. They'll use the intermediate option, reducing fatalities.

HOST: But that's not what happened?

DR. HARTWELL: No. What happened is officers used force more frequently overall. Because they had a "less-lethal" option, they deployed it in situations where previously they would have just talked.

HOST: Give me an example.

DR. HARTWELL: Okay, so pre-taser, an officer encounters an agitated person who's not complying with verbal commands but not posing an immediate threat. The officer has two options: keep talking, or shoot. Obviously they keep talking. Almost always.

Post-taser, same situation. Now they have three options: keep talking, taser, or shoot. And suddenly tasering seems reasonable. It's less-lethal, right? So they taser.

HOST: And then what?

DR. HARTWELL: Sometimes the person has a heart condition and dies. Sometimes the taser doesn't incapacitate them and the situation escalates to shooting anyway. Sometimes multiple officers taser the same person repeatedly, causing death through cumulative shock. The overall death rate went up in many jurisdictions that adopted tasers.[9]

HOST: Why?

DR. HARTWELL: Because the presence of less-lethal options lowered the threshold for using force at all. Officers felt they had a "safe" option, so they used it liberally. But nothing's actually safe. Tasers kill people. Pepper spray kills people. Rubber bullets kill people. Just at lower rates than firearms.

And worse—deploying those weapons often escalates situations that could have been de-escalated verbally. So you get more overall force encounters, and even though each individual encounter is "less-lethal," the aggregate effect is more deaths.

HOST: So the metric went up—more less-lethal deployments—while the goal—fewer deaths—got worse.

DR. HARTWELL: Exactly. They optimized for weapon diversification and destroyed the actual goal of harm reduction.

[MUSIC TRANSITION - 17:00]

[ACT FOUR: TEACHING TO THE TEST - 17:30]

HOST: Let's talk about education. Because this might be the most widespread example of Goodhart's Law destroying a system.

No Child Left Behind, passed in 2001, tied school funding to standardized test scores.[10] The logic: measure student performance, incentivize improvement, education quality goes up.

The result: Test scores went up. Education quality went down.

I talked to Sandra Williams, a teacher in Richmond who's been in public schools for 22 years.

SANDRA: I watched this happen in real time. Pre-NCLB, I had autonomy. I could teach based on what my students needed. Some kids needed more time on certain concepts. Some were advanced and needed enrichment. I could adapt.

HOST: And after NCLB?

SANDRA: After NCLB, everything became about the test. We cut art. We cut music. We cut recess. We cut social studies. Everything that wasn't on the test got eliminated. And what was on the test, we drilled endlessly.

HOST: Did test scores go up?

SANDRA: Test scores went way up. Students got very good at taking standardized tests. But they weren't learning. They were memorizing patterns. Test-taking strategies. They could bubble in the right answer without understanding the concept.

HOST: How do you know they weren't learning?

SANDRA: Because I'd have students who scored "proficient" on the math test who couldn't solve a real-world problem. They'd mastered the test format but not the underlying mathematics. Or students who scored well on reading comprehension but never read a book for pleasure because we'd made reading a chore.

HOST: What happened to curious kids?

SANDRA: Curiosity got punished. If a student asked an interesting question that wasn't related to the test, I had to shut it down. "That's not on the test. We don't have time." Year after year of

that, and kids stop being curious. They stop asking questions. They learn that education is about compliance, not exploration.

[COFFEE SHOP AMBIENCE - 20:30]

HOST: This is Goodhart's Law at civilizational scale.

The metric: Test scores. The goal: Educational quality. The result: Test scores optimized, education destroyed.

Sandra:

SANDRA: The cruelest part is that everyone involved knows it's broken. Teachers know. Parents know. Even administrators know. But we're trapped. If our scores drop, we lose funding. If we lose funding, we can't serve the kids who need the most help. So we keep teaching to the test, hating it, watching it hollow out education, because the alternative is worse.

HOST: Is there a way out?

SANDRA: Not within a testing framework. You cannot measure education quality with standardized tests. Education is about developing thinking skills, creativity, curiosity, resilience, social-emotional competence. None of that shows up on a bubble sheet.

HOST: So how do you measure it?

SANDRA: You don't. You trust teachers who know their students to make professional judgments. But that requires trusting expertise and accepting unmeasurable outcomes. And our system can't do that. Everything has to be quantified, measured, ranked. So we measure what we can and destroy what we can't.

[MUSIC TRANSITION - 22:30]

[ACT FIVE: WHY WE CAN'T SEE SYSTEMS ANYMORE - 23:00]

HOST: So here's my question: why are we so bad at this?

Why do we keep making the same mistake—optimizing metrics, destroying systems—over and over?

Dr. Thompson:

DR. THOMPSON: I think it's a combination of factors. First, we've lost the capacity to think in systems. We think in linear cause-and-effect. "If we do X, we'll get Y." But systems are non-linear. They have feedback loops, emergent properties, time delays. Optimizing X often produces not-Y.

HOST: Why have we lost systems thinking?

DR. THOMPSON: Education. We teach subjects in isolation. Math separate from science separate from history. We don't teach how things connect. How systems work. How changing one variable affects everything else.

Second, measurement bias. We can only manage what we measure. So we focus on measurable things and ignore unmeasurable things. But often the most important things—meaning, relationship quality, wisdom, resilience—are unmeasurable.

Third, time horizons. Metrics give quick feedback. Systems reveal themselves slowly. It might take five years to see that your test score optimization destroyed curiosity. But you get quarterly reports on test scores. So you optimize for the quick feedback and ignore the slow consequences.

HOST: Can we fix this?

DR. THOMPSON: We'd have to change everything. How we educate people. How we measure success. How we think about time. How we value things that can't be quantified. That's... a lot.

[MUSIC TRANSITION - 25:30]

[ACT SIX: THE HOSPITAL PROBLEM - 26:00]

HOST: Let me give you one more example. This one hits close to home because it affects people's lives directly.

Medicare started penalizing hospitals for high readmission rates.[11] The logic: if patients are readmitted within 30 days, care quality must be poor. Incentivize lower readmission rates, quality improves.

I talked to Dr. Lisa Patel—the microbiome researcher from episode three—who also works as a hospitalist.

DR. PATEL: The readmission metric seemed reasonable. If we're discharging people who aren't actually stable, they come back. That's bad care. So let's measure readmissions and reduce them.

HOST: What happened?

DR. PATEL: Hospitals became terrified of readmissions. Because the penalties are severe. So what do you do if a patient shows up at the emergency room 15 days after discharge with concerning symptoms?

HOST: Readmit them?

DR. PATEL: That's what you should do. That's good care. But that counts as a readmission. Penalty. Lost revenue. Administrators angry. So instead, many hospitals started observing patients instead of admitting them. Technically they're in the hospital, technically they're receiving care, but technically they're not admitted so it doesn't count as a readmission.

HOST: That's just a definitional trick.

DR. PATEL: Exactly. It games the metric without improving care. Worse—observation status means patients aren't covered the same way by Medicare. They end up with huge bills they weren't expecting. And if they actually need to be admitted after 24 hours of observation, that doesn't count as a readmission because they were never technically discharged. It's a shell game.

HOST: Did readmission rates go down?

DR. PATEL: Measured readmission rates went down. Actual patients being readmitted? Harder to say because of all the observation status games. And other outcomes got worse—patients being discharged too early because we're scared of readmissions, leading to worse health outcomes at home.[12]

HOST: So you optimized the metric...

DR. PATEL: ...and destroyed the system. Again. We made care worse while making the numbers look better.

[COFFEE SHOP AMBIENCE - 29:00]

HOST: What's the alternative?

DR. PATEL: Qualitative assessment by skilled clinicians who know the patient. I can tell you if someone's ready for discharge better than any metric. But that requires trusting my judgment. And judgment is messy. It can't be standardized. It can't be automated. So instead we use metrics that can be gamed and call it quality improvement.

HOST: Why do administrators prefer metrics to clinical judgment?

DR. PATEL: Because metrics are objective. Quantifiable. Defensible. If something goes wrong and you followed the metric, you're protected. If something goes wrong and you used judgment, you're vulnerable. So the system incentivizes cowardice. Follow the metric even when you know it's wrong, because at least you can't be blamed.

[MUSIC TRANSITION - 30:30]

[ACT SEVEN: WHAT TRADITIONAL KNOWLEDGE KNEW - 31:00]

HOST: So I've been thinking about this through the lens of what we talked about in episode five—pratyaksha pramana, direct perception, embodied knowledge.

Traditional knowledge systems—indigenous knowledge, Ayurveda, traditional ecological knowledge—didn't rely on metrics. They relied on long-term observation of whole systems.

I called Dr. Lakshmi Bharadwaj—the scholar of Indian philosophy from episode five.

DR. BHARADWAJ: Traditional systems understood something we've forgotten: you cannot understand a system by measuring its parts. You have to observe the whole, over long periods, through direct perception.

HOST: Give me an example.

DR. BHARADWAJ: Okay, so traditional farming. A farmer doesn't measure soil nitrogen content, pH levels, moisture percentages. They look at the soil. They smell it. They feel it. They observe what grows, how it grows, when it struggles. They watch for decades. They develop intuition about the whole system—soil, water, sun, seeds, pests, weather patterns.

HOST: But modern agriculture uses metrics.

DR. BHARADWAJ: Right. Measure NPK ratios, optimize fertilizer input, maximize yield. And it works! For a while. Yields go up. But soil health degrades. Topsoil erodes. Biodiversity collapses. Water tables drop. You've optimized yield and destroyed the agricultural system.

Traditional farmers know: you can't optimize for yield alone. You have to maintain the whole system. And maintaining the whole requires seeing the whole. Which requires time, observation, presence. Not metrics.

HOST: Why don't we do that anymore?

DR. BHARADWAJ: Because it doesn't scale. Traditional farming requires intimate knowledge of specific land. You can't transfer that knowledge to someone else's land. You can't automate it. You can't teach it in a classroom. You have to develop it over decades of direct observation.

Modern agriculture can be standardized. One formula for nitrogen. One formula for irrigation. Same everywhere. Measurable. Optimizable. Scalable. And it destroys the land within 50 years. But in the short term, the metrics look great.

[MUSIC TRANSITION - 34:30]

[ACT EIGHT: THE COBRA EFFECT - 35:00]

HOST: I want to come back to the cobra story because it's so perfectly absurd.

British colonial India, early 1900s. Too many cobras in Delhi. Government offers bounty for dead cobras.[13]

Initially, it works. People kill cobras, bring in bodies, collect bounties. Cobra population drops.

Then people realize: why hunt wild cobras when you can breed them? Set up cobra farms. Breed cobras. Kill them. Collect bounties. Easy money.

Government discovers this. Ends the bounty program.

Now cobra breeders have worthless cobras. What do they do? Release them. Cobra population explodes—higher than before the program started.

This is the ultimate Goodhart's Law example. The metric was "dead cobras brought to government." The goal was "reduce cobra population." And optimizing the metric made the goal worse.

[COFFEE SHOP AMBIENCE - 36:30]

I asked people: have you ever gamed a metric and made things worse?

PERSON 1: I work in customer service. We're measured on calls per hour. So I get people off the phone as fast as possible. Doesn't matter if I solved their problem. I hit the metric. And then they call back, and someone else has to deal with it, and the total time spent on the customer goes up. But my individual metric looks good.

PERSON 2: My company measures lines of code written. So I write verbose, inefficient code. More lines = better metric. But the codebase becomes unmaintainable. We're optimizing for quantity and destroying quality.

PERSON 3: My kid's school measures attendance. So sick kids come to school because parents don't want the attendance ding. Then those sick kids infect other kids. Total sick days probably go up. But attendance metrics look good.

PERSON 4: My gym measures membership numbers. So they make it really hard to cancel. People keep paying even though they don't come. Membership numbers look great. Actual gym usage is terrible. People hate the gym. But the metric is optimized.

[PAUSE]

HOST: Every single person I asked had an example. Everyone's participating in systems they know are broken. Gaming metrics they know are destroying the actual goals. And we're all trapped in it.

[MUSIC TRANSITION - 38:30]

[ACT NINE: CAN WE REBUILD SYSTEMS THINKING? - 39:00]

HOST: So can we fix this? Can we learn to think in systems again?

Dr. Thompson:

DR. THOMPSON: It would require a complete shift in how we approach management, education, governance. We'd need to:

One, teach systems thinking from early education. Show kids how things connect. How feedback loops work. How optimizing one variable affects everything else.

Two, accept that not everything can be measured. Some goals are qualitative. You have to use judgment. And judgment requires expertise and trust.

Three, extend time horizons. Stop optimizing for quarterly results. Look at five-year, ten-year, generational consequences.

Four, value maintenance of systems over growth of metrics. A system that's stable and sustainable at lower output is better than a system that's growing but degrading.

Five, listen to people embedded in systems. The teacher knows if education is working better than test scores do. The doctor knows if care is good better than readmission rates do. Trust expertise over metrics.

HOST: That sounds impossible.

DR. THOMPSON: It is impossible within current structures. Because all of those things—qualitative judgment, long time horizons, trusting expertise—they're unmeasurable. They can't be quantified. They can't be standardized. So our systems can't value them.

HOST: So we're stuck?

DR. THOMPSON: We're stuck until the systems collapse from their own internal contradictions. Which they will. You can only optimize metrics while destroying systems for so long before the systems stop functioning entirely.

[COFFEE SHOP AMBIENCE - 41:30]

HOST: Sandra, the teacher, told me something that stuck with me.

SANDRA: I've stopped trying to fight the system. I teach to the test because I have to. But I also teach real things in the cracks. Five minutes here, ten minutes there. I answer the curious questions even though I shouldn't. I slip in books that aren't on the curriculum. I create little pockets of actual education inside the testing regime.

HOST: Does that work?

SANDRA: It's not enough. But it's something. And something is better than nothing. I can't fix the system. But I can subvert it in small ways. Help a few kids actually learn instead of just perform. That's what I've got.

HOST: That sounds exhausting.

SANDRA: It is. But the alternative is participating fully in something I know is destructive. And I can't do that. I'd rather be exhausted and subversive than complicit and comfortable.

[MUSIC TRANSITION - 43:00]

[ACT TEN: CLOSING - 43:30]

HOST: So here's what we've learned about Goodhart's Law, about metrics, about systems:

When you optimize the metric, you destroy the system.

Every. Single. Time.

Recycling robots → people stop sorting → less recycling Less-lethal weapons → more force used → more deaths Test scores → teaching to test → education destroyed Readmission rates → gaming definitions → worse care Cobra bounties → cobra breeding → more cobras

The pattern is universal. The mechanism is predictable. And we keep doing it anyway.

Why?

Because we've lost the capacity to see systems. We can only see metrics. We can only manage what we measure. We can only value what we quantify.

And the things that matter most—ecosystem health, educational quality, care, meaning, wisdom, resilience—can't be quantified.

So we measure proxies. And we optimize the proxies. And we destroy the things we actually care about.

[PAUSE]

The traditional knowledge Dr. Bharadwaj described—observing whole systems over long periods through direct perception—that's the capacity we need to rebuild.

Not rejecting measurement entirely. But recognizing measurement's limits. Knowing when metrics help and when they harm. Developing the judgment to see whole systems even when we can't measure them.

Marcus, from waste management, said it best:

MARCUS C: The robots are amazing at what they do. But they can't see the whole system. Only humans can do that. And only if we're paying attention. Only if we're present. Only if we resist the urge to outsource our judgment to the machines.

[THEME MUSIC - 46:00]

[OUTRO - 46:30]

HOST: Things Overheard at the Coffee Bar is produced by Greenheart Media. Our theme music is by Lauren Pastrana.

Next week: The final episode. "Irreversible Windows." The myopia epidemic. The decline of free play. What happens when children miss critical developmental periods. And why some losses can never be recovered.

If you've watched metrics destroy systems, if you've gamed numbers while knowing you were making things worse, if you've tried to see wholes in a world that only values parts—send us a voice memo.

Thingsoverheardpod@gmail.com

Special thanks to Dana Mitchell, Marcus Chen, Dr. Jennifer Hartwell, Sandra Williams, Dr. Lisa Patel, Dr. Amara Thompson, and Dr. Lakshmi Bharadwaj.

And to everyone trying to preserve something real inside systems designed to destroy it.

Close your laptop. Put your phone down. Look at something whole. See if you can resist breaking it into measurable parts.

[END - 48:00]

[PRODUCTION NOTES: This episode should feel increasingly frustrated and urgent. Dana should sound resigned but angry. Marcus C should sound like someone who's learned a hard lesson. Dr. Hartwell should sound like she's revealing an uncomfortable truth about policing. Sandra should sound exhausted but defiant. Dr. Patel should sound like she's trapped in an

absurd system. The cobra story should be told with dark humor—it's so absurd it's almost funny except it's real. Music should build tension throughout.]

---

2 episodes remaining:

- Episode 9: "Irreversible Windows" (FINAL EPISODE)

Should I complete the season with Episode 9?