



Emergent Strategy with the
AI Macroscope

a project of PolicySolve

CENTER FOR
EVALUATION
INNOVATION

When AI Augments Philanthropy

Disrupting the Typical
Traps That Constrain Our
Systems Change Strategies

Jewlya Lynn & Tanya Beer

October 2026





Table of Contents

Executive Summary	02
Introduction	05
Foundations are optimized for stewardship of the institution, not the systems change strategy	06
Trap #1: Designing for manageability instead of systems change	08
Trap #2: Conflating practice frameworks with systems strategies	11
Trap #3: Moving slowly, safely, and additively	15
Trap #4: Failing to account for how systems defend against change	18
Trap #5: Defining success narrowly, measurably, and in ways that often aren't durable	21
What if AI Doesn't Help Us Transform?	25
What Now?	26
References	27
About the Authors	28

AI Disclosure

All drafts of the article, from initial thinking to final polish, were human-generated. Generative AI (specifically, Claude) was used to edit, remove jargon, and simplify complex concepts and sentences. AI-generated phrases and sentences offered during editing were incorporated when helpful.

Executive Summary

AI presents philanthropy with its first real opportunity in a generation to transform how foundations support systems change. It presents an equally real chance to lock in the patterns that have kept foundations from doing systems change well. **When used poorly, AI will harden the sector's conventions** that work against transformative systems change. A funder that trains AI only on their own thinking gets a more sophisticated way of listening to themselves. **When used well, AI can be a tool for transforming philanthropy and strengthening its systems change impact.** It can surface the assumptions funders don't realize they're making and pressure-test strategies for how systems will defend against those assumptions.

The possibility of using AI to strengthen systems change depends on funders first seeing what drives the current behaviors that prevent success. The traps described below are the ones we encounter most often in foundations' systems change strategies. Each trap can be partially disrupted by strategically using AI to augment human capabilities, enabling us to have greater systemic impact.

Traps	Explanation	AI Disruption Opportunity
1 Designing for manageability instead of systems change	A strategy clean enough to govern and structured enough to evaluate against is one designed for institutional manageability, not for influencing unpredictable and dynamic systems.	AI can push back when strategies drift toward manageability and away from strategic actions; articulate the risks of untimely exits from the system; and protect against the habit of changing priorities.
2 Conflating practice frameworks with systems strategies	One way philanthropy copes with the structural pressures in Trap 1 is by adopting frameworks that seemingly simplify the messy work of systems change (e.g., grantmaking approaches, collaborative process frameworks, leverage point frameworks). These frameworks describe what good <i>philanthropic</i> practices look like, and we design around them instead of around the systemic change we are trying to influence.	AI can expand what we pay attention to and design for systems change, not philanthropic practices, by generating alternative ways of operationalizing values; tracing the causal logic of our strategies to show how far the chain reaches into systems dynamics; and matching the demands of a specific system to the strategy, making visible where mismatches occur.

Traps	Explanation	AI Disruption Opportunity
3 Moving slowly, safely, and additively	Foundations consistently choose strategies that add benefit, or “create more good,” over strategies that disrupt what isn’t working. Adding benefit is easier to fund, justify, measure, and defend; it produces visible results and positive stories and makes few enemies.	AI can help measure decline (which sometimes matters more than growth); surface counterfactuals of inaction; speed up our ability to plan and act; and sort a strategy’s risks into those borne by the foundation versus those borne by people in the system.
4 Failing to account for how systems defend against change	Much of the strategy that emerges from foundations treats the system as something to be acted upon and rearranged, as if a well-funded, well-coordinated intervention will influence it in proportion to the resources behind it. Yet, systems are active structures, organized to protect the status quo.	AI can push back steadily against thinking that forgets to attend to how systems defend; it can draw on bodies of research about how systems defend and apply them to specific systems; it can model systems resistance; and it can listen to how a system is resisting in real-time.
5 Defining success narrowly, measurably, and in ways that often aren’t durable	Most philanthropic measurements are organized around the moment of a win, but systems do not stop reorganizing when a foundation milestone is reached. Either the measurement is narrow enough to be easily tracked regardless of durability or context, or so broad that it fails to learn about how the system is changing in response to strategies.	AI can help us listen to the system first, rather than our own strategies and their outputs; enable cycles of rapid systems sensing to listen for emerging changes; manage the complexity of incoming systems information; and find causality amid complexity.

Bringing Generative AI into a philanthropic organization does not automatically challenge these traps. It could just as easily help us double down on those traps if we focus only on how Generative AI increases efficiency, manages more information, justifies collecting more, and generates additional content. Used that way, it would help us confirm our current beliefs, rather than challenging them, by reinforcing our metrics, deepening our confidence in problematic strategies, and further separating us from the systems we seek to influence.

AI cannot disrupt the core structures of philanthropy that have led to these five traps. However, it can augment our ability to see the traps, see beyond them, and be prepared to change our actions.



It can help us bring deeper systems expertise, move at an otherwise unavailable pace, and listen to systems in ways we have not done before. At the same time, we must protect our uniquely human roles in expanding our relationships, partnering with those most harmed by systems, bringing curiosity and hope to the work, and taking responsibility for the decisions we must make. This combination, where AI augments our thinking while we maintain our commitments and values, offers an opportunity to transform our systems change work and have a greater positive impact on the world.

Introduction

AI presents philanthropy with its first real opportunity in a generation to transform how foundations support systems change. It presents an equally real chance to lock in the patterns that have kept foundations from doing systems change well. This choice arrives at a moment when philanthropy's commitment to systems change has never been more visible, or its results more elusive. Foundations have rewritten strategies, built dedicated portfolios, hired staff to lead the work, adopted shared frameworks, and made systems change a fixture of conferences, learning agendas, and field-wide conversations. **Yet the systems that philanthropy seeks to influence keep behaving in remarkably durable ways.**


AI is beginning to disrupt the philanthropic conventions that produce low impact on systems, exposing their limitations and risks. Whether these disruptions sharpen philanthropy's ability to support transformative change or simply make existing patterns faster and harder to dislodge depends on **whether philanthropic leaders treat this as a moment to reimagine or upgrade their practices.**



When used poorly, AI will harden the sector's conventions that work against transformative systems change. A funder who trains AI only on their own thinking gets a more sophisticated way of listening to themselves. A foundation that uses AI to accelerate its existing strategy process will refresh its strategy faster, carrying forward the same assumptions about how change happens, assumptions that have been formed inside the foundation and never tested against the system. AI in the service of efficiency will entrench what already isn't working. AI treated as a stand-in for community voice, which it cannot be, will substitute the funder's interpretation of the system for the system itself.



When used well, AI can be a tool for transforming philanthropy and strengthening its systems change impact. It can help funders engage with the complexity of systems and align strategies to where change can actually happen. It can surface the assumptions funders don't realize they're making and pressure-test strategies for how systems will defend against them. It can challenge conventional thinking about how values translate into practice, loosening the assumption that there is only one way. Used this way, AI acts as a sparring partner with the patience to do the uncomfortable work of examining a foundation's own thinking.



The possibility of using AI to strengthen systems change depends on funders first seeing what drives current behaviors preventing success. The traps described below are those we encounter most often in foundations' systems change strategies. Each is rooted in the structure of the foundations themselves, which is why good intentions and better frameworks have not dissolved them. Each also has a counterpart, something that becomes possible when AI augments human judgment, including partial escape from constraints no foundation can simply wish away. The same technology can just as easily harden each trap, so we close with a clear-eyed look at how AI use could go wrong in philanthropy's hands, and what funders should do to make sure it doesn't.


Foundations are optimized for stewardship of the institution, not the systems change strategy


Foundations are organized around stewardship. The institutional features that shape philanthropic practice, from strategy cycles and measurement architectures to board reporting, learning agendas, and funding boundaries, are stewardship in action. The question is, what is being stewarded?

Today, **most foundations are well organized to steward the institution** itself, meaning its assets, reputation, continuity, and legitimacy within the field. Foundations are much less organized to steward the systems change they aim to support. Stewarding a foundation and stewarding systems change call for different things. Systems change requires attention to time horizons longer than a strategy cycle and signals subtler than a measurement framework can register, and it carries risks larger than most boards will easily absorb.

Boards must be informed, assets must be managed, staff must be supported, reputations must be tended, and legal and fiduciary requirements must be met. These are real responsibilities that are not going away. Yet the question every foundation must ask is which features of its institutional stewardship are required for it to function, and which are **inherited conventions that have outlasted their purpose and now get in the way of effective systems change work.**

[T]he question every foundation must ask is which features of its institutional stewardship...now get in the way of effective systems change work.





In our work supporting foundations' systems change strategies, we routinely observe five additional traps that emerge when institutional stewardship goes unexamined and crowds out systems change practices. Each trap began as a sensible response to a real problem and has, in many foundations, hardened into a default that operates beneath the level of strategic decisions. These traps are entangled and reinforce one another. They concern how we design and manage strategies, use practice frameworks as a shortcut around deeper analysis, weigh risk and pace, and define and measure success, as well as what we account for or ignore in the systems we work in.

Together, these traps produce strategies that are attractive on the page, coherent and easy to explain, measurable and reportable over time, and easy to invite others into, whether funding partners, grantees, or a board that needs to stay committed. However, these same strategies have limited impact on the systems they aim to change, rarely achieve durable results, and are readily co-opted and absorbed by the system.

Generative AI cannot resolve the structural drivers underneath these traps, but it can help us push against them and engage with systems change strategy in new and sometimes transformative ways. Several of AI's capabilities make this possible:

- **AI can change information asymmetries**, making deep systems expertise available to more people. It is needed at the moment, which lets us work inside complexity in ways that were not previously possible.
- **AI can increase the speed at which we make informed decisions**, provided we have the authority to act on them quickly, by helping us analyze how change is happening as it happens.
- **AI can support our commitment to inclusion** by widening access to knowledge about the system, so that those we work with can see into parts of the system they historically could not and help construct strategies that reach upstream drivers.
- **AI can help us push back against the cognitive and social biases** that have built up in philanthropy over time, naming them and making visible very different ways of advancing change.

Understanding Terms

Generative AI, across most platforms, can be trained to think in specific ways, process information through specific workflows, and generate predefined types of outputs. Workflows include such things as project instructions, context documents, skills, and the content of the prompts themselves. In this article, we will refer to these ways of directing AI as training the AI and prompt sequences.

[For more on both, see the resources on AI Macroscope.](#)

Trap #1

Designing for manageability instead of systems change

Foundations frequently mistake the coherence required to *manage a strategy* for the strategic fit required to *change a system*. A strategy clean enough to govern and structured enough to evaluate against is one designed for institutional manageability; whether it is a strategic fit with the system is a separate question. Systems change strategy ideally looks quite different. It is less contained and less pre-planned, and it stays responsive to what the system does. This trap shows up in many of philanthropy's routine patterns:

- Strategy refresh cycles that arrive on the foundation's calendar, rather than when the system is ripe for a different approach.
- Learning agendas shaped by the questions the foundation wants answered about its own role, rather than the questions it needs answered to understand the system it is trying to change.
- Hiring patterns that favor topic expertise, proximate experience, or grantmaking craft over the analytic and strategic skills and knowledge required to understand how a specific system changes.
- Strategies built around what can be planned and projected for a board, rather than what the system might actually require as it shifts.
- Strategic direction that shifts with internal pressures, such as staff turnover or leadership and board transitions that trigger a rethinking of program areas and goals every few years, for reasons unrelated to the system being addressed.
- Boundaries set up front about what the foundation will and won't fund, foreclosing strategies before any analysis of the system has been done.

Systems change strategy...is less contained and less pre-planned, and it stays responsive to what the system does.

Many of these patterns began as reasonable responses to real demands. Yet, **when manageability becomes the implicit test of strategic quality, the foundation becomes the unit of analysis, and the intended systemic changes become secondary.**

Strategies optimized for the foundation tend to underweight the parts of the system the foundation does not already understand and underfund the work that does not fit its existing categories. **The system continues to behave as it always has, while the foundation refreshes its strategy on schedule.**

The system continues to behave as it always has, while the foundation refreshes its strategy on schedule.

AI Augmentation: Designing strategies optimized for systemic change

Of these structural issues, Generative AI is best positioned to address the gap left when staff are hired for reasons other than knowledge of and skill in systems change strategy. Trained well, it can also hold the system in view, pushing back when strategy design drifts toward managerial legibility over the ability to influence the system, and when learning agendas prioritize internal needs over system dynamics. It can articulate the risks of untimely exits from the system and protect against the habit of changing priorities whenever leadership turns over.

The examples below show what this looks like in practice. The first is a hypothetical constructed to illustrate how AI can shift our ongoing learning focus. The second comes from a foundation we work with in the US that ran a new strategy through a series of AI pressure-testing tools (Premortem Analyst, Devil's Advocate, and Systems Defender). Each tool pairs a set of AI instructions with a structured sequence of prompts and is trained on a different set of analytical, evidence-based frameworks, so that each gives a distinct type of feedback.

Structural Issue

What Becomes Possible with AI-Augmented Insight

Learning agendas are shaped around the questions the foundation wants answered, rather than the questions the system needs the foundation to understand.

Pressure Testing a Systems Change Strategy

A hypothetical for illustration: a foundation launching a five-year strategy on youth mental health in a mid-sized state sets its learning agenda before any grants go out. Its four learning questions ask (1) whether school-based counseling grants improve student outcomes, (2) whether the foundation's convening role strengthens the provider network, (3) whether its advocacy investments influence the state Medicaid reimbursement rate, and (4) what the foundation should fund next.

Structural Issue

What Becomes Possible with AI-Augmented Insight

(continued) Learning agendas are shaped around the questions the foundation wants answered, rather than the questions the system needs the foundation to understand.

(continued) An AI trained in systems change frameworks is asked to review the agenda against the strategy's own systems map. It finds that every question tracks a foundation decision rather than a system dynamic; and three of the four questions would be sufficiently answered through grantee reporting, with little other documentation required. This means these questions are about things grantees already know. They don't need this learning.

The AI also named that the map's most consequential features (a workforce shortage driven by licensure rules, a reimbursement structure that rewards crisis care over prevention, and a school funding formula that makes counselors the first cut) appear in none of the questions. Yet, these are issues grantees have named as critical drivers of the systems inequities and progress on changing them is not as visible or easy to measure. The systems-trained AI recommends the foundation use grant reporting to answer the original four internal questions, while working with grantees to identify learning that they collectively need to monitor how the system is changing on the most consequential issues.

What the AI can make visible: Every learning question was a question about whether the foundation's bet was paying off, but none asked what the system was doing when the foundation wasn't looking.

Hiring patterns that favor topic expertise, proximate experience, or grantmaking craft over the analytic and strategic capacities required to understand how a specific system changes.

Pressure Testing a Systems Change Strategy

A foundation designing a multi-state strategy to increase funding for early childhood care drafted geographic criteria to select states to fund, focusing on those with established provider networks, supportive legislatures, and existing advocacy infrastructure, which in practice meant states with left-leaning and centrist electorates.

The AI analysis brought systems knowledge of durable change to those criteria and found that they conflated campaign winnability with the durability of systems change. It pointed instead to states that lack some of that infrastructure but are experiencing disruptions, such as provider closures, Medicaid restructuring, or workforce collapse, that could be leveraged toward durable, non-partisan systemic change.

It also supplied expertise the program team did not have in the intersection of government and market dynamics, explaining that private equity firms were buying and consolidating providers across the states where public funding expanded, and no state's regulatory infrastructure had been built to oversee a shift from early childhood care services as a public-good to services optimized for profit. The foundation team was familiar with a federal-level non-profit working on this issue, but had not included the issue in the state strategy. The AI analysis encouraged working at the state regulatory level to prepare for this dynamic instead of waiting for public dollars to be secured before tackling regulatory reform.

Generative AI Insight: *"The strategy provided a replicable model for appearing progressive on controversial issues while avoiding structural challenges to wealth concentration."*

Trap #2

Conflating practice frameworks with systems strategies

One way philanthropy copes with the structural pressures in Trap 1 is by adopting frameworks that seemingly simplify the messy work of systems change. These frameworks describe what good *philanthropic* practice looks like, and we often design around them rather than around the systemic change we are trying to influence. We may even build a systems map and a theory of change first, naming the system and how it might shift, and then revert to planning a familiar set of practices anyway.

The frameworks themselves are worthwhile. They help us decide how to make grants, to whom, with what requirements, and with what kinds of technical assistance, collaboration, and support. They also give us categories to help us group and label our strategies and outcomes. They fall into four distinct categories, each useful for a different purpose:

1. **Collaborative process frameworks** are intended to guide philanthropy and others on which *practices* are most helpful in systems change work. Collective Impact is a well-known example of this, as are movement-building and field-building frameworks.
2. **Many of the frameworks for systems thinking** offer system-agnostic insights about the nature of systems and how to change them, functioning as cognitive and conceptual tools. They can simplify the entry into systems work, but when used as the primary lens for grantmaking decisions, can also unintentionally limit what we see. The frameworks used most often in philanthropy draw on iceberg and river metaphors, hierarchies of systems behaviors, and distinctions between simple, complicated and complex aspects of systems.
3. **Leverage point frameworks** are even more narrow and focus on one specific aspect of systems change, with relatively little attention to how that leverage point relates to other systemic dynamics. Narrative change, power-building, policy advocacy, and network weaving all represent frameworks in this category.
4. **Values-based practice frameworks** offer concrete philanthropic practices (e.g., general operating grants, multi-year grantmaking) that operationalize critical values to which philanthropy has historically undercommitted. Some of the most often discussed frameworks in this space are trust-based philanthropy and participatory grantmaking.

All of these frameworks offer philanthropy something useful, from clearer vocabulary about collaboration to concrete practices to ways to address real and important problems, often problems in the funder-grantee relationship. However, when used as stand-ins for deeper system-specific analysis, the frameworks overly simplify the complex, bound our thinking, and offer an appealingly straight path forward. **The frameworks are too often used as a replacement for rather than a supplement to the work of designing a systemic change strategy**, which may include elements of the framework practices but should go beyond them to respond to the system's dynamic complexity.

To be clear, the values shaping contemporary progressive philanthropic practice (e.g., equity, agency, self-determination, the redistribution of power, attention to who holds voice and who has been silenced) are not in question. They are essential, hard-won commitments, and deserve the central place they have come to occupy in the field. The trap lies in how the frameworks meant to operationalize these values are used.

This shows up in recognizable ways:

- Strategies are described in the language of collaboration, coalition, partnership, and field, with little articulation of how the system being addressed will actually change via these activities.
- Theories of change describe who is at the table and how decisions get made, without describing how any of that will disrupt the dynamics holding the system in place.
- Portfolios are organized around funding systems change *actors* and fail to also fund systems change *moments*. These windows, when a system is genuinely open to shift, often arrive on timelines and through actors that the established collaboration cannot capitalize on.
- Strategic responsibility is increasingly shifted to grantees most proximate to harm, with foundations stepping back from the analytic and political work of deepening their own understanding of how the broader system and power-holders within it operate.¹

¹ One particularly difficult tension in this trap involves the role of proximate knowledge. The principle that those most affected by a system must have a voice and power in changing it is a hard-won and necessary commitment, and proximate knowledge is indispensable to any serious systems change work. At the same time, the dynamics that produce and protect a system's harms (e.g., market forces, political institutions, capital flows, regulatory regimes) are often distant from and intentionally opaque to those experiencing the harms.

Consequently, when philanthropy treats centering proximate experience as a *sufficient* theory of how change happens, the parts of the system hidden from proximate view go unanalyzed. A commitment to shifting power to those harmed can be bolstered by a commitment to also learning with and from other systems actors who can see into the more opaque power-centers of the system. Without this additional knowledge, a foundation can fund a strategy that identifies systemic racism or economic exploitation as the problem, feel the moral weight of having taken it on, and yet (unintentionally) design every intervention to avoid threatening the interests of the very actors who sustain those conditions.

AI Augmentation: Putting practice frameworks *in service of* systems change strategy

One reason we default to practice-based frameworks is that they are easier to use than the sprawling, jargon-heavy academic literature on how systems change. A simplified version of that literature does not solve the problem, because no single short framework can make sense of a given system, and not all frameworks are equally relevant to all systems.

A Generative AI partner trained in theoretical and research-based systems frameworks can make systems thinking as accessible as practice-based frameworks. It can also quickly identify and say where practice-based thinking has eclipsed strategic, systems-level thinking.

Over time, the practice-based frameworks also let us settle into mindsets that focus attention on how we act and with whom, at the expense of how the system changes and why. AI can help loosen those mindsets in several ways:

- Making our assumptions and biases visible and how we've locked them into our strategies.
- Generating alternative ways of operationalizing the same values, this time aimed at systemic change.
- Tracing the causal logic of a strategy to find where the chain runs only through our own collaboration and practice, without touching the larger dynamics of the system.
- Testing whether a strategy that puts those closest to the harm in the lead has left anyone working on the parts of the system they cannot reach.
- Matching the demands of a specific systems strategy to the frameworks best suited to meet them, so that collaboration, leverage, and values-based practices are deployed where they will make a difference instead of being applied by default.

Structural Issue

What Becomes Possible with AI-Augmented Insight

Strategies are described in the language of collaboration, coalition, partnership, and field, with little articulation of how the system being addressed will actually change as a result.

Pressure Testing a Strategy

A field-building strategy designed to bring coherence and manage complex power dynamics in a field struggling to achieve any meaningful impact was soundly rejected by an AI trained in systems thinking and systems defense. The AI noted that the diagnosis of the field's issues was robust, but the strategy defaulted to improving the field without ever asking whether the field's decline was necessarily a bad thing. It made visible several alternative fields already emerging in the system that the funder could choose to support, rather than protecting the actors it had always funded, who represented an older model of change that had lost relevance.

Generative AI insight: *"The strategy correctly diagnoses the field's internal pathologies. But a perfectly healthy network that lacks the power to alter anything outside itself is, from the system's protection perspective, an ideal outcome."*

Strategic responsibility increasingly shifting to grantees most proximate to harm, with foundations stepping back from the analytic and political work of understanding how the system and its powerholders operate.

Pressure Testing a Strategy

A foundation's strategy is committed to following the lead of community-based organizations closest to the harm, with the foundation's own role limited to resourcing and convening. Asked to identify where the strategy located the system's power, a well-trained AI found that the strategy described the affected population in detail, but barely described the actors sustaining the power imbalance in the system. It named the asymmetry directly, noting that the foundation had devolved strategy to the actors with the least leverage over the system, while retaining for itself relationships and standing with regulators, industry, and other funders that no grantee could access, and then declining to use them for fear of abusing its own power. The finding sparked considerable debate within the foundation and prompted a conversation about how and when the foundation should use the standing it holds in the system to create pressure on powerholders that the people most harmed by the system cannot generate alone.

Generative AI insight: *"This strategy asks those with the least structural power to design the change, and reserves for the foundation the one form of power the strategy never deploys. Deference is being performed where it costs the least."*

Trap #3

Moving slowly, safely, and additively


Foundations consistently choose strategies that add benefit, or “create more good,” over strategies that disrupt what isn’t working. Adding benefit is easier to fund, justify, measure, and defend; it produces visible results and positive stories and makes few enemies. Disrupting, in the sense of actually loosening the dynamics that hold the system in its current shape, is harder to attribute, more likely to fail, and tends to provoke pushback from the institutions and interests protected by the system. Faced with this asymmetry, philanthropy defaults to the additive. The system retains its shape while accumulating more programs and supports designed to mitigate its effects.

The slowness with which philanthropy designs and deploys these strategies is another part of the problem. The pace is due in part to the manageability issues noted above and in part to strong collaborative values. However, systems do not move slowly, particularly in disruptive times. By the time a thoughtful, additive strategy has been designed and implemented, the window in which the system was vulnerable to change has often closed.

The commitment to carefully crafted additive strategies shows up in several patterns:

- Treating growth in positive activity (e.g., more collaboration, engagement, programs, and/or participation) as equivalent to systems change, even when that activity runs alongside the system without challenging its underlying dynamics.
- Seeking more models and pilots to layer on top of existing system dynamics, rather than tackling head-on the systemic barriers that prevent successful models and pilots from becoming the system’s default behaviors.
- Portfolios weighted toward supporting those who are marginalized (e.g., leadership development, movement building), without *also* seeking to disrupt the systemic dynamics that have prevented past leaders and movements from having a significant impact.
- Harm reduction as the default response to disruptive moments, such as emergency support to keep nonprofits stable or direct services for those in crisis. This orientation is values-aligned, and it privileges the most immediate and visible harms over the larger, longer-term harms that philanthropy could choose to prioritize.

By the time a thoughtful, additive strategy has been designed and implemented, the window in which the system was vulnerable to change has often closed.



The avoidance of disruptive strategies has its own recognizable patterns:

- A pattern of avoiding contested or politically uncomfortable work, even when that work is what systems change would require.
- Risk assessments organized around the foundation's reputational and legal exposure, rather than around the system's likely response to the strategy.
- Action withheld in uncertain moments, rather than treating the choice *not to act* as the riskier failure.
- Taking a long time to plan the "right" strategy, without recognizing that slow action allows current harm to continue and preserves a harmful future.

Two dynamics keep this trap in place. The first dynamic is institutional and related to Trap 1: Foundations are organizations with reputations to protect and boards to satisfy, and they reasonably consider the safety of their grantees. These pressures push toward strategies that the foundation can defend if questioned and that are planned and implemented slowly to minimize mistakes, even at the cost of missing windows of opportunity.


The second dynamic is psychological. Additive work carries fewer of the tensions that come with disrupting what isn't working, including fewer strained relationships and uncomfortable questions about the foundation's own role in what the system protects. It can also feel more immediately rewarding.

Together, these dynamics **produce a sector that excels at making bad systems more tolerable** and struggles to actually change them.

None of this is a failure of individual judgment, and the instinct to add good is not one philanthropy should be asked to give up. Making people's lives better is real work and worth doing. The trap lies in letting that instinct displace the harder work of changing the systems that make those lives harder in the first place. **The foundation gets to claim impact while the system that produces the harm remains intact, ready to generate the next cohort of people the foundation will need to help.**

AI Augmentation: Acting rapidly, boldly, and disruptively

In many of its uses, AI makes us uncomfortable by revealing what we did not want to see or could not see before.



It can articulate when a strategy avoids the underlying issues; show why additive strategies will fail to matter for the larger systems change aspiration, even when we can measure them as successes; and offer information we can act on proactively, shifting our focus toward disrupting problematic systemic dynamics as well. AI can also help overcome this trap in other ways:

- Helping us measure decline, since sometimes less of something matters more than growth in something else.
- Surfacing the counterfactuals of inaction and the asymmetries in who bears harm.
- Speeding up our ability to plan and act, reducing the likelihood that action comes too late.
- Making expertise rapidly available to those we hope will design with us so that we can act together more quickly.
- Sorting a strategy's risks into those borne by the foundation and those borne by people in the system, so we can see whose risk we are actually managing.

In terms of speed, AI enables the sourcing of insights at a previously unavailable pace. Where we once had to limit what we paid attention to and took far too long to learn and act, AI can rapidly make systemic issues visible, even as they are changing, and help us work through them with partners. **Real-time use of AI during collaborative systems sensing** is one of the most powerful ways to engage it, disrupting our thinking and expanding our knowledge base in the very moment we decide how to act.

Structural Issue

Treating growth in positive activity (e.g., more collaboration, engagement, programs, and/or participation) as equivalent to systems change, even when that activity runs alongside the system without challenging its underlying dynamics.

What Becomes Possible with AI-Augmented Insight

Iterative Process of Learning Sprints and AI Feedback Loops

Through its engagement with AI, a philanthropic team focused on transforming youth education and pathways to economic opportunity is learning that its strategies are more additive than disruptive. The team discussed whether additive strategies would be enough and why they resonated so strongly. As the team began a six-month series of deep listening conversations, it built in an iterative AI review that analyzed each day-long conversation with systems change frameworks, explored the implications, and shaped the next conversation to go deeper. The combination supports ongoing strategy refinement and surfaces experiments to act on along the way.

Structural Issue

What Becomes Possible with AI-Augmented Insight

(continued) Treating growth in positive activity (e.g., more collaboration, engagement, programs, and/or participation) as equivalent to systems change, even when that activity runs alongside the system without challenging its underlying dynamics.

(continued) Encouraged by the AI findings to look beyond additive strategies, the listening conversations explored systems dynamics in ways that made barriers more visible. Partners in the dialogues still offered familiar truisms (“we need more collaboration,” for instance), and the facilitator used the AI analysis and the team’s reflections to push further, surfacing the underlying dynamics that explain why those truisms persist and what has to be solved structurally before any additive collaboration strategy will have an impact.

Taking a long time to plan the “right” strategy, without recognizing that slow action allows current harm to continue and preserves a harmful future.

Real-Time Use of AI in Systems Sensing Sessions

A team with a strong systems change practice has already adapted to funding systems change moments and actors. Its AI use takes this further, augmenting systems sensing sessions with a deeper look at current system dynamics and planning strategy in response. The team uses a “real-time AI partner” trained to listen during meetings, surface what sits underneath what people are saying and the assumptions they are making, and investigate systemic dynamics where the team lacks current or historical knowledge.

Trap #4

Failing to account for how systems defend against change

Systems are active structures organized to reproduce themselves with their own logics of resistance, absorption, and accommodation. Yet much of the strategy that emerges from foundations **treats the system as something to be acted upon and rearranged, as if a well-funded, well-coordinated intervention will influence it in proportion to the resources behind it.** The assumption is rarely stated, and it shows up in observable patterns:

- Theories of change that describe what the foundation will do without describing what the system will do in response.
- Evaluation frameworks that measure outputs and outcomes but rarely track resistance, co-optation, or absorption.


- Strategy documents that treat the system as a problem to be solved rather than as an “actor” with its own interests.
- Funding that goes to the organizations most successful at playing by the existing system’s rules, rather than those working on transformative alternatives.
- Portfolios that fund the work without funding the analysis of how the system has previously responded to similar work.
- Risk assessments organized around grantee performance and the foundation’s reputational exposure, rather than around the system’s likely defensive moves.

The defenses that systems mount are well documented. Institutions decouple what they say from what they do. When they act, they do so slowly, with many pauses for seemingly justifiable reasons. They engage in cosmetic compliance, making changes that look good but leave the system undisturbed in any significant or durable way. This combination lets them outlast many philanthropic strategies (Meyer and Rowan, 1977; Marquis and Qian, 2014).

Interest convergence occurs in many systems, as reforms that begin as challenges to dominant interests end up serving them, often because the policy or program was designed in ways the system was already prepared to absorb (Crenshaw, 1988; Bell, 1980). **Strategies drift** toward what existing coalitions and institutions are willing to do, and dominant discourses define the problem in ways that constrain possible solutions (Hajer, 1995; Stone, 2002). **Movements get domesticated** as their demands are reframed into demands the system can tolerate (Piven and Cloward, 1977; INCITE!, 2007).

Beyond these overt actions to prevent change, systems also have structural patterns that make it very difficult to alter the status quo, even in the absence of active resistance. **Institutional lock-in** (or “path dependence”) occurs when earlier decisions become embedded in structures, relationships, and investments that raise the cost of switching paths over time. The longer a system runs along a particular trajectory, the harder it becomes to choose alternatives (Pierson, 2000). **Feedback loops** pull conditions back toward established patterns or amplify what is already dominant. These loops do quiet, continuous work, rewarding expected behavior and absorbing pressure for change before it accumulates (Meadows, 1999).

Systems also stay put because of **aligned incentives**. Actors across the system respond to the same rewards, so their behavior converges on what looks like a consensus without any explicit agreement. Breaking from that convergence is costly, so the existing path becomes the rational choice and the status quo holds without anyone having to defend it (DiMaggio and Powell, 1983).



A final way that systems resist change without active defense is through **cognitive and narrative** constraints. The system's stories about what works, what counts as a real problem, who the legitimate actors are, and which solutions are viable get quietly infused throughout it, shaping hiring, training, and professional norms, until the boundaries those stories create feel like reality rather than constructions (Hajer, 1995).

Foundations that imagine themselves as outside agents directing change toward an inert system have already misjudged where they stand and what they are working against.

These dynamics and many other ways systems defend themselves have been documented for decades. Yet the **foundation strategies that fund systems change work rarely anticipate systemic defense, and even more rarely build in mechanisms to address it.**

The trap deepens when foundations forget that they are inside the system. Philanthropic capital is part of the systems philanthropy seeks to change, accumulated through them, governed by them, and granted legitimacy by them. Foundations that imagine themselves as outside agents directing change toward an inert system have already misjudged where they stand and what they are working against. As actors in the current system, foundations are often complicit in the defenses described

above, unintentionally driving interest convergence and rewarding cosmetic compliance. The strategy becomes a way of demonstrating commitment to change while structurally ensuring that the most uncomfortable changes never have to happen. Strategies get absorbed faster than they can shift anything, and an entire field is left surprised, year after year, by how durable the systems within which it works have turned out to be.

AI Augmentation: Anticipating and responding rapidly when systems defend

AI's specific contribution here is to push back steadily against thinking that forgets to attend to how systems defend against change. It removes the risk of never considering the defense at all and brings the accumulated observation of dozens of theorists and researchers who have watched systems defend themselves, historically and in the current environment, and can anticipate common defense pathways as a result. In practice, this means giving the AI a set of analytical frameworks for how systems have defended themselves in similar contexts and directing it to look for evidence that a system is defending itself as a strategy deploys.

AI can help in several further ways:

- Modeling system resistance and pressure-testing strategies for co-optation, cosmetic compliance, decoupling, and interest convergence.
- Listening to the system in real time to identify weak and strong signals of resistance as they emerge.
- Finding examples of how the system has resisted similar attempts in the past, and which kinds of resistance succeeded.

Structural Issue

What Becomes Possible with AI-Augmented Insight

Failure to account for system resistance and defense of the status quo.

Real-Time Use of AI in Theory of Change Sessions


A hypothetical for illustration: a foundation supporting strategies to shift district hiring and retention practices toward a teacher workforce that reflects its students is preparing to bring its theory of change into a set of AI-supported sessions with its strategic learning partners. The AI will be trained on frameworks for how systems defend themselves, with particular attention to the education system's capacity for cosmetic compliance and decoupling, in which districts adopt recruitment targets and affinity programs while certification pathways, placement decisions, and the seniority rules that govern layoffs stay as they were. The intent is to identify, before the strategy deploys further, which of its assumptions about district and union behavior the system is most likely to absorb rather than change, and build tracking for those defensive moves into the team's learning practice.

Generative AI Insight: *"The strategy's assumptions all sit upstream of the decision that determines who is actually in a classroom in September, and none of them survive a layoff cycle."*

Trap #5

Defining success narrowly, measurably, and in ways that often aren't durable

Trap 5 undergirds the four traps that define how strategy is typically designed (traps 1-4). Inattention to durable systemic change is a deep structural problem, shaped by board and staff tenures, the limits of a foundation's strategy cycles, and the norms of the monitoring and evaluation field. Most philanthropic measurements are organized around the moment of a win, but systems do not stop reorganizing when a foundation milestone is reached. Policies narrow in implementation, and coalitions disperse as a crisis recedes, so the visible win at Year 3 may be largely absorbed by Year Seven.



Short-cycle measurement has a real purpose, surfacing near-term signals supporting learning, adaptation, and developmental work. The trap is short-cycle measurements focused on narrow outcomes with little attention to durability or whether the system has actually shifted. Without that connection, measurement generates incentives that pull foundations toward visible, attributable wins and away from durable systemic change.

Foundations often define success in one of two ways:

Narrow enough to be easily tracked and attributed regardless of durability or context, and short-cycle enough to be reported within a grant cycle or strategy refresh.


OR **So broad** that the focus falls on the ultimate population outcomes the system produces (e.g., improved educational outcomes), even when the core systems arrangements that drive the problem are unchanged, and the drivers of those outcomes are far too many and variable to parse out where the foundation's work contributed.


Foundations sometimes do both, envisioning population-level change while measuring grant outputs, and conflating progress on the latter with progress on the former or connecting the two with little causal interrogation.

[A] foundation that defines change as population-level outcomes can remain confident in a strategy that is failing to change the system, on the grounds that systems change takes a long time.

These measurement choices are consequential. They **reward the behaviors that produce Traps 1-4 and discourage systems-aware work**. A foundation that defines success as policy passage (narrow and trackable) gets a cleaner contribution story, while a foundation that defines change as population-level outcomes can remain confident in a strategy that is failing to change the system, on the grounds that systems change takes a long time.

A foundation that defines success as durable systemic change inherits the much harder problem of observing change in systems arrangements (e.g., practices, norms, resource and information flows, power dynamics, decision-making protocols, and policy). It must also track how the system absorbs, redirects, or co-opts a policy over years and decades.






Measurement systems tend to reward narrow, tidy wins and quietly close off the work of durable change. In practice, this takes familiar forms:

- **Outcomes scoped to fit within grant cycles** and strategy refresh windows, rather than within the less predictable and highly variable timelines on which systems actually shift.
- **Attribution-focused metrics** that require the foundation to show its resources produced or significantly contributed to the outcome, incentivizing strategies with traceable, narrow theories of change.
- **Policy wins are counted as durable systemic change** without tracking what happens after passage, as the system absorbs the policy or narrows its scope.
- Measurement frameworks that **stop at tracking how well grantees work together** without understanding whether their work has actually influenced the system.
- **Measurement of engagement and participation** of those most proximate to harm, without also measuring whether those most proximate are actually less harmed as the system changes, a pattern that extends the proximity trap from Trap 3 (conflating practice frameworks with systems strategies) into the measurement architecture itself.
- A consistent **gravitational pull toward downstream indicators** where harm is visible and trackable and away from the upstream drivers (e.g., wealth concentration, market dynamics, institutional design) where the system's architecture and outcomes originate.
- Comprehensive studies that explain more significant change stories, but **completed after a strategy is itself completed**, or after a strategy refresh has come and gone.
- **A focus on measuring outcomes**, with little attempt to interrogate why, how, and under what conditions the outcomes (and other changes) emerged, inclusive of drivers well beyond the philanthropic strategy.

A related dimension of this trap is that foundations are rarely solo actors in systems change, yet they almost always measure as if they were. Each foundation tracks its portfolio against its own indicators and inside its own contribution space. **Signals that would reveal whether the broader system is moving are produced by the combined work of many actors and are only visible when measurement is shared and interpreted collectively. They go largely unmeasured because no single foundation owns them**, and the infrastructure that would support cross-foundation, cross-sector learning has not been built. Foundations learn about their own grants, but they rarely learn, together with allies, about the system.



Foundations do commission reports from the nonprofits they fund, designed to support advocacy. These documents, often a mix of prevalence studies, documentation of policy wins and losses, and descriptions of the current context, serve as the ongoing evidence of the system's problems and needs. They pay little attention to how the system behaves as a whole, instead attending to discrete parts of the system or sets of outcomes that persuade funders, policymakers, and others to act. Such analyses serve their purpose and are a critical part of systems change work. As typically constructed, though, they differ from sensemaking about the system and how it is changing, which is what a foundation needs to adapt its strategy and respond when the system defends itself.

AI Augmentation: Finding success in the system's story

Short-cycle, narrowly defined success measurement is driven by many structural issues. It is also a product of the knowledge, capacity, and scope of monitoring, learning, and evaluation teams and their consultants, as well as the readiness of program teams to guide their learning partners. AI can bring new expertise into our learning spaces. It can also enable a more complex kind of listening to how change emerges, drawing on multiple signals and patterns instead of single metrics.

So far, AI's early uses in evaluation have focused on efficiency (e.g., coding qualitative data, cleaning quantitative data sets). An augmentation lens suggests that AI can support learning practices embedded in the complexity of how systems change. We become more able to study the system itself, looking well beyond the strategy for evidence of how the system is changing overall and where our strategy's influence shows up within it. We gain the capacity to investigate causal pathways by scaling the data we collect and analyze to more easily seek both confirming and disconfirming evidence, pushing against our own thinking and building a stronger case for when our work is truly influencing change. We can also generate analyses of complex systems on timelines that make them actionable within a strategy, instead of spending a year or more on a study whose findings are out of date by the time they arrive.

Structural Issue

What Becomes Possible with AI-Augmented Insight

Comprehensive studies that explain more significant change stories, but completed after a strategy is itself completed, or after a strategy refresh has come and gone.

Rapid Systems Analysis with AI

A foundation with a strong systems change and emergent strategy approach has used Generative AI from 2024 through 2026 to produce rapid systems sensing analyses. Its teams engage in deep and wide-ranging conversations with grantees, funding partners, and others to make sense of how the system is changing. Using AI, hundreds of pages of transcripts are analyzed with systems frameworks. Within a week, an analysis is ready for use in systems sensing and goal-setting conversations inside the foundation and with external partners (Measuring what matters, 2025).

Structural Issue

What Becomes Possible with AI-Augmented Insight

A focus on measuring outcomes, with little attempt to interrogate why, how, and under what conditions the outcomes (and other changes) emerged, inclusive of drivers well beyond the philanthropic strategy.

Managing Complexity in Systems Analysis with AI


A foundation with more than a decade of deep involvement in a specific global and in-country system commissioned a causal pathways analysis of how the system changed over 10 years, reaching well beyond the scope of its own strategy and of philanthropic influence generally. AI supplemented the analysis in two ways: finding confirming and disconfirming evidence for more than 350 causal relationships in the 10-year story and challenging the early anchoring on key findings that naturally emerged as analysts and systems actors explored the data together. The result was a set of conditions for systemic change that made visible how and why the system actually changed, including the finding that the foundation's funding of systems change "moments," alongside systems change actors, was one of many critical drivers (Lynn and Phootong, 2025).

What if AI Doesn't Help Us Transform?

Generative AI has created an unusual moment for many systems and sectors, philanthropy among them. As the examples above show, it can push against constraints we have largely accepted as givens and let us operate in new ways, including stepping outside the traps described here.

Generative AI can just as easily help us double down on those traps if we focus only on how it increases efficiency, manages more information, justifies collecting even more, and generates additional content. Used that way, the harm would be specific:

- It would confirm our current beliefs, since AI trained only on a funder's own thinking returns that thinking in a more sophisticated form.
- It would separate us further from the systems we are part of and seek to influence by rewarding continued attention to relationships and funding practices.
- It would reinforce our focus on reportable metrics by making the collection and analysis of grantee-reported outputs and outcomes more efficient.
- It would deepen our confidence in coherent strategies that predict tidy pathways to change at the expense of the healthy uncertainty and steady listening that reveal how change really happens.



Philanthropy's traps have structural roots that AI will not touch, whether the framing of board fiduciary duty, legal architecture of perpetuity, proximity to donor wealth, career incentives within foundations, or the political economy of tax-exempt status. Yet AI can disrupt the other ways we are trapped, including inherited and unquestioned norms, practices that have hardened into defaults, beliefs that restrict more than they enable, and limits on capacity and knowledge. How we use AI will determine whether philanthropy doubles down on its least effective habits for influencing complex systems or steps into roles and forms of influence it has historically been unable to reach, advancing equity even as change in the wider society accelerates.

What Now?

The traps described here are held in place by structure. No single tool can disassemble philanthropy's core structure. What a well-trained AI changes is the cost of seeing. Deep systems expertise, the record of how systems have defended themselves before, the discipline of asking whose questions a learning agenda serves, and evidence of what a strategy is avoiding all become available at the moment a decision is being made to people who did not previously have them. That is enough to loosen the default traps we have described, provided that funders are clear about what they are asking AI to do and what they see as the distinct role of humans.

As we augment our systems change work with AI, what do we need to be sure we protect as our uniquely human role? Leading from within a system, with an honest account of where the foundation sits in it, is human work. So is making sense of what an analysis means for this foundation, in this moment, and making the choices that follow. Facilitating dialogue across groups that do not trust each other and building the relational structures that a strategy depends on are human work. So too are holding the values that decide which harms count and bringing curiosity and hope to work that offers little of either. AI cannot make these contributions and should not be asked to imitate them. A foundation that hands any of them to a tool has repeated Trap 2 in a new form, letting a practice stand in for judgment.

The urgency extends beyond philanthropy's own practice. AI is disrupting the systems that foundations work in, at a pace that will not wait for the next strategy refresh, from labor markets and information ecosystems to public services and the economics of the nonprofit sector itself. Systems in flux are systems open to change, and the actors best positioned to shape what comes next are already moving. If philanthropy approaches this moment as an efficiency upgrade, it will arrive at the systems change moments of the next decade with faster versions of the strategies that missed the last windows of opportunity.

References

Bell, D. A. (1980). *Brown v. Board of Education and the interest-convergence dilemma*. Harvard Law Review, 93(3), 518–533.

Crenshaw, K. W. (1988). *Race, reform, and retrenchment: Transformation and legitimation in antidiscrimination law*. Harvard Law Review, 101(7), 1331–1387.

DiMaggio, P. J., and Powell, W. W. (1983). *The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields*. American Sociological Review, 48(2), 147–160.

Hajer, M. A. (1995). *The politics of environmental discourse: Ecological modernization and the policy process*. Oxford University Press.

INCITE! Women of Color Against Violence (Ed.). (2007). *The revolution will not be funded: Beyond the non-profit industrial complex*. South End Press.

Lynn, J. and Phootong, P. (2025). *Understanding Systemic Change: A 10-Year Retrospective: Addressing Forced Labor and Human Trafficking in the Thai Seafood and Fishing Industries*. Humanity United.

Marquis, C., and Qian, C. (2014). *Corporate social responsibility reporting in China: Symbol or substance?* Organization Science, 25(1), 127–148.

Meadows, D. H. (1999). *Leverage points: Places to intervene in a system*. The Sustainability Institute.

Measuring what matters: From conversations to collective impact. (2025, August 25). Imaginable Futures Lessons and Learnings.

Meyer, J. W., and Rowan, B. (1977). *Institutionalized organizations: Formal structure as myth and ceremony*. American Journal of Sociology, 83(2), 340–363.

Pierson, P. (2000). *Increasing returns, path dependence, and the study of politics*. American Political Science Review, 94(2), 251–267.

Piven, F. F., and Cloward, R. A. (1977). *Poor people's movements: Why they succeed, how they fail*. Pantheon Books.

Stone, D. (2002). *Policy paradox: The art of political decision making (Rev. ed.)*. W. W. Norton.

About the Authors



Dr. Jewlya Lynn


Jewlya Lynn is a facilitator, advisor, and researcher who helps leaders navigate complex systems change. Over more than 20 years, she has worked across philanthropy, government, and social change initiatives to understand how change happens in dynamic systems and how strategy can adapt accordingly.

Jewlya's current work explores how AI is disrupting systems, how it can strengthen systems change practice, and how leaders can make better use of evidence and learning in complex environments. Jewlya has led or advised major systems initiatives including the Causal Pathways Initiative, the U.S. National Science Foundation's (NSF) EarthCube Initiative, Horizon 2045, N Square, and efforts by Humanity United and the Freedom Fund to address forced labor.



Tanya Beer

Tanya Beer is a strategist, facilitator, and learning partner who has spent nearly 25 years helping foundations and nonprofits navigate complexity, uncertainty, and systems change. Her work focuses on how organizations learn well in complex environments, challenge their own assumptions, and use evidence to strengthen strategy rather than simply confirm it. She has worked extensively with foundation learning and evaluation leaders, including as the former Associate Director of the Center for Evaluation Innovation and co-director of the Evaluation Roundtable, a network of leaders from more than 100 foundations. Today, Tanya advises teams and collaboratives working on issues ranging from food systems transformation and climate philanthropy to gender equity and social justice, and works with Jewlya Lynn on improving the use of AI for systems change strategy and learning.



a project of PolicySolve

Emergent Strategy with the **AI Macroscope**

CENTER FOR
EVALUATION
INNOVATION