# Tokenmaxxing for Good

## A Framework for Philanthropy's Role in AI

*Inspired by Gary Community Ventures*


WRITTEN BY

**Santhosh Ramdoss**

President & CEO, Gary Community Ventures



PUBLISHED

**September 2026**




`tokenmaxxing` *n.*

Leaning into AI and using it extensively. *Tokens* are the units of text an AI model reads and writes; the more you use AI, the more tokens you burn.

**good** *n.*

Whatever betters humanity and our communities: what makes life better for everyone.

The role of philanthropy in this current moment is rather profound. Philanthropy literally means the love of humanity. And given the exceptional scale and diversity of American philanthropy, and its presence across communities far and wide, this institution is tailor-made for defending and co-creating what a human-centered AI future could look like.

But philanthropy has often been at the fringes of other major technological shifts that have had a significant impact on humanity, often sitting them out or arriving a little too late to the party. Given how polarizing the conversations around AI have become in the last few months, the natural predisposition of philanthropic leaders and donors is to avoid AI and try to stop it, rather than work deeply with it to help chart a positive coexistence while minimizing its negative impacts.

> **Philanthropy literally means the love of humanity.**

This working paper offers a playbook, primarily as a resource for philanthropy that cares about place and communities. It was originally written as an internal document to provide direction for our AI strategy at Gary Community Ventures, and it still serves that purpose: clarifying how we at Gary are navigating and leveraging this moment. Our proximity to community, our unique philanthropic model (including an in-house engineering and product team), and our sunset in 2035, with the urgency that comes with it, all led us to grapple with questions of AI use, impact, and coexistence far earlier than most institutions would. This paper is a reflection of those learnings, which I hope will be of value to other philanthropic institutions and leaders.

PART ONE

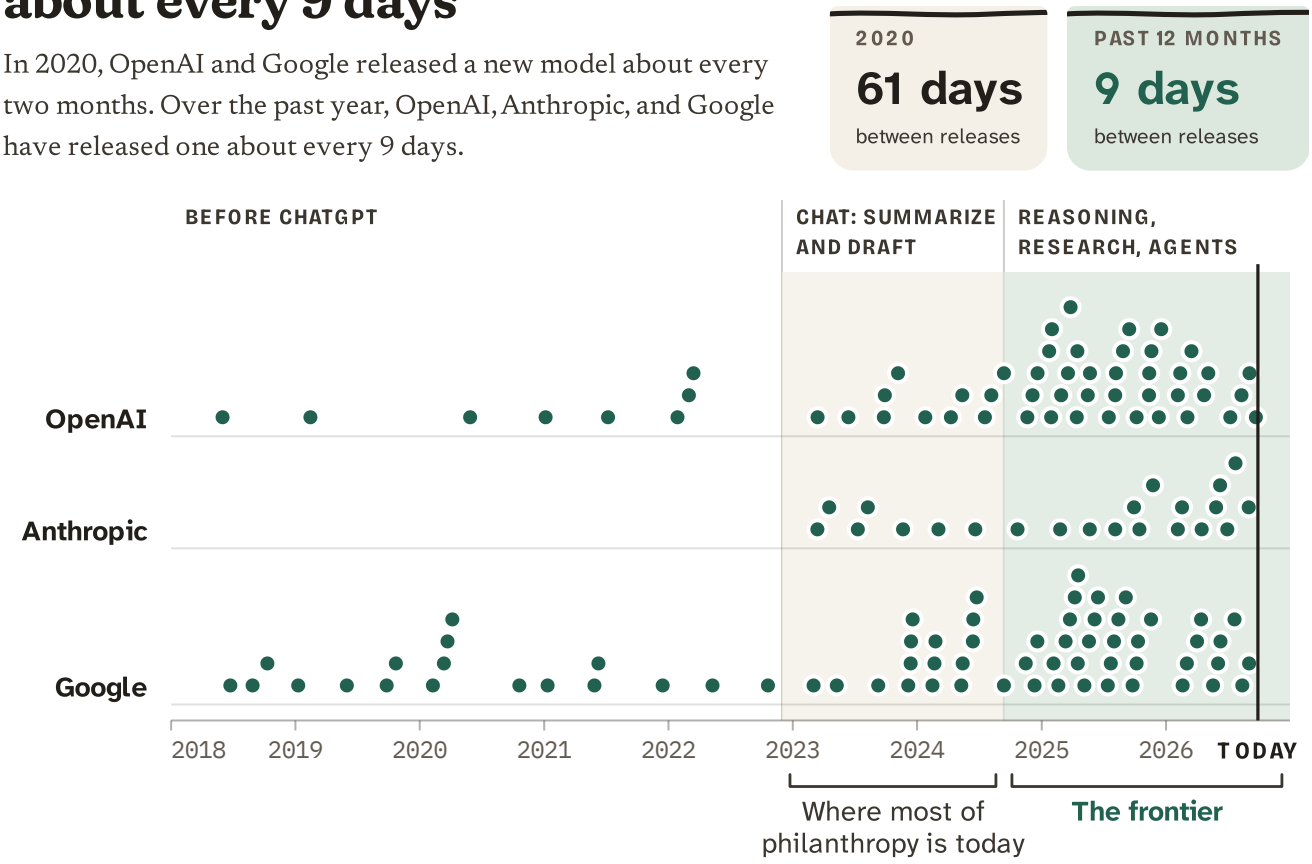
# Taking Stock of the Current Moment:

*Utopia, Dystopia, and Wonder*

HOW FAST THIS IS MOVING

## A new model about every 9 days

In 2020, OpenAI and Google released a new model about every two months. Over the past year, OpenAI, Anthropic, and Google have released one about every 9 days.




Source: Epoch AI, AI Models database (CC BY). Publicly available language and multimodal models, one dot per lab per release day, through September 22, 2026.

## Philanthropy is not used to adapting to change this fast.

Plenty has been said and written about how fast AI capabilities have accelerated over the last three years, and I won't belabor the point here (see: pelicans; or, more rigorously, METR).

Our reactions to this foundational technological shift have mostly leaned toward one of two poles: utopia or dystopia. The utopians believe AI could make life dramatically better for every human being, in ways we could never have imagined (see: Machines of Loving Grace). The very people building these systems

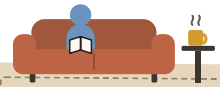
have been among the loudest in describing the dystopian possibilities: a future in which these technologies could, quite literally, kill us all (see: the Statement on AI Risk, signed by the heads of OpenAI, Anthropic, and Google DeepMind).

One of the more powerful voices, if a less heard one, belongs to researchers like Princeton's Arvind Narayanan and his co-author Sayash Kapoor, who rightly characterize AI as transformative and significant, yet still as "normal technology." That framing lets us pause, breathe, and notice what is actually happening in front of us, rather than swinging on a pendulum between exceptional hope and exceptional fear. This is also my call to philanthropy, and to my own team at Gary.

I will say this much: working with these models, from the earliest chat interfaces to the steady swarm of agents I rely on today, has been, to me at least, one of the most wondrous things I have experienced. It is a testament to what humanity is capable of. Our capacity for ingenuity, innovation, and creation is astounding and beautiful, even when the creation takes the form of machines.

The beauty, for me, is not in the suboptimal art or writing that reads like the relentless average of human expression. It is in these models' capacity to navigate human complexity, especially in the machine-like worlds we have built for ourselves.

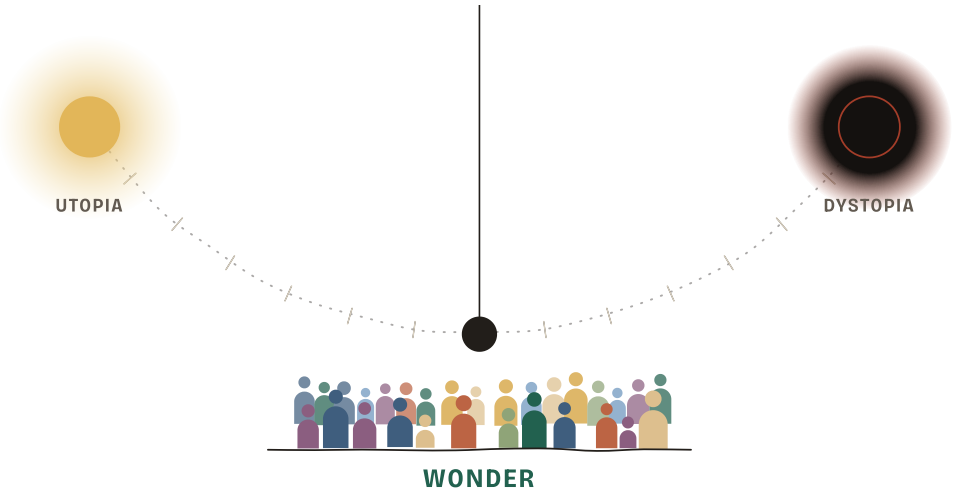
Lately, the wonder is in being freed from the slog of clicking through hundreds of websites and pop-up ads to understand a topic, or of fighting my way through my daughter's pediatrician's patient portal to book her vaccination. The wonder is sitting on my couch with a good book and a warm cup of coffee, knowing that an agent is handling the needless complexity of the world we have built, and letting me focus on what I am good at: being human.



> **The wonder is sitting on my couch with a good book and a warm cup of coffee, knowing that an agent is handling the needless complexity of the world we have built, and letting me focus on what I am good at: being human.**

The wonder is also watching a well-informed organizational agent slip in and out of our team's Slack threads, contributing analysis that would otherwise have taken hours. That gives us time for the judgment, conversation, and decision-making we so often skip because we are too busy wading through PDFs and other complexity of our own making.

Most of all, the wonder is in what all of this could mean for the communities we care about most: families facing mind-bending complexity while they figure out how to pay for the next week's groceries, or worrying about a child's doctor's appointment while afraid of losing their Medicaid coverage.




*AI as transformative, yet "normal technology." A place to pause, breathe, and notice what is actually happening in front of us.*

Select a pole to read it. The pendulum always settles in the middle, where our communities are.

---

This call to curiosity and wonder shapes our approach, and it asks for balance. At one end, we notice and pay close attention to the potential harms of this technology. At the other, we understand it, use it, and see its possibilities. And in the middle sit our communities and our humanity.

PART TWO

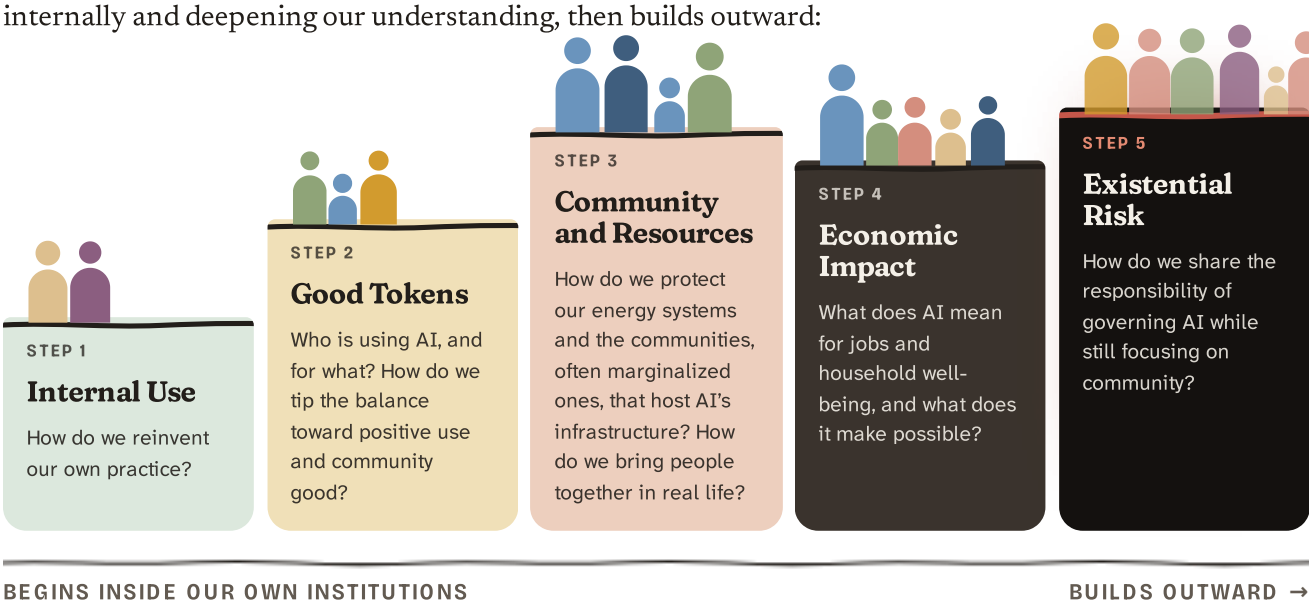
# AI in Philanthropy:

## *Tokenmaxxing for Good*

We now turn to practice: approaches and insights that I hope will hold up even as the AI ecosystem keeps shifting beneath us, and that help us manifest the middle and embrace our curiosity. The question throughout is how philanthropy can lean in to understand, use, and see the possibilities of this technology, while also noticing its harms and helping to mitigate them.

The cheeky title, Tokenmaxxing for Good, is intentional. My argument all along has been that to shape this technology, we must first understand it, and to defend our communities, we must first integrate it into our work, build with it, and learn to coexist with it. I would argue that for philanthropy, more than for any other sector, this is a prerequisite for our broader goals: creating good for humanity and keeping humans at the center.

It is in that spirit that I have organized what follows as a spectrum of AI use and strategic engagement. Think of it as a step function for philanthropic organizations. It begins with adopting these tools internally and deepening our understanding, then builds outward:




A step function for philanthropic organizations: Internal Use → Good Tokens → Community and Resources → Economic Impact → Existential Risk. Each step stands on the one before it.

In each area, I look at the possibilities philanthropy should lean into and the probable threats we need to navigate and protect our communities against. Sometimes the two converge. And because this playbook comes from practice, each section closes with the approach we have taken at Gary. A caveat: the existential risk section is deliberately light, because this guide is designed primarily for place-focused, community-oriented institutions, which make up the vast majority of philanthropy.

STEP 1 OF 5

# Internal Use

**LEAN INTO**

Reinventing our practice through faster understanding, more building, and more time in community.

**GUARD AGAINST**

Dehumanizing the sector and displacing its workforce.

The practice of philanthropy is in desperate need of reinvention. As institutions, we are increasingly out of sync with the moment. Resources keep accumulating in the sector, and AI wealth is about to multiply them further. Yet our practices were designed for the slower pace of academic learning and research, with little bias toward experimentation, building, and moving big ideas forward. This is not universal, but I would argue that AI can be an accelerant for philanthropies big and small to reimagine the practice: not by stepping away from human talent, but by concentrating our incredible density of talent on the hardest problems, the ones machines cannot solve.

Three possibilities stand out. The first is speed to understanding. Today you can reach a fair grasp of the current state of knowledge on almost any topic quickly and effectively, compressing weeks, if not months, of work that once went to consultants assembling a research report or white paper. That frees our teams to focus on solving the problem itself, alongside the people in the room who are closest to it.

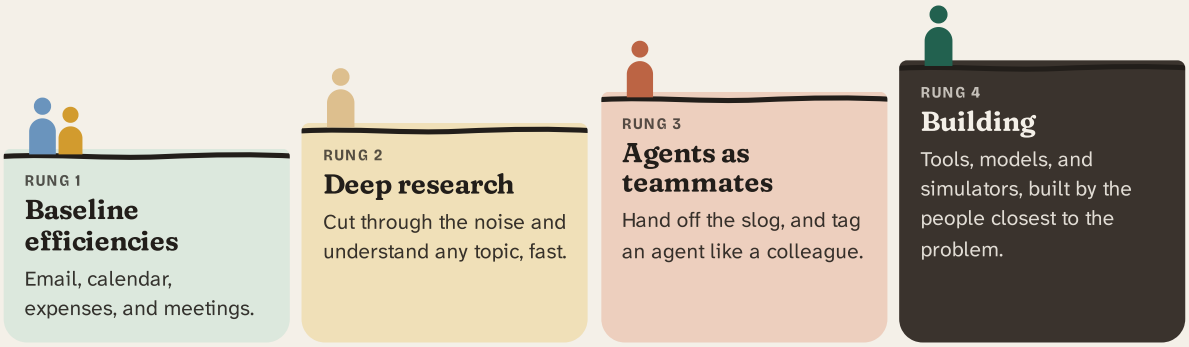
The second is building. Twenty-first-century philanthropy has shied away from it, focusing instead on moving small pools of capital out of the endowment. Philanthropy in the early twentieth century had a builder's bias (think of Carnegie's 2,500 libraries or the nearly 5,000 Rosenwald schools), and AI can help us recover it by dramatically accelerating what we can build.

The third is proximity. These tools can bring us closer to the work and to community, so that program officers spend less time in front of screens and more time in front of the communities they serve.

The risks are real: dehumanizing our sector and displacing its workforce. But we have a rare opportunity to put philanthropic teams at the forefront of learning to coexist with AI and using it to solve hard problems, and to say clearly that this is not about displacement and dislocation. It is about reimagining the way we work as a sector.

## THE APPROACH AT GARY

It starts with reimagining how our team works, one rung at a time. Every colleague climbs at their own pace, and the point of each rung is the same: hand the slog to the agents, and spend the time on people and hard problems.

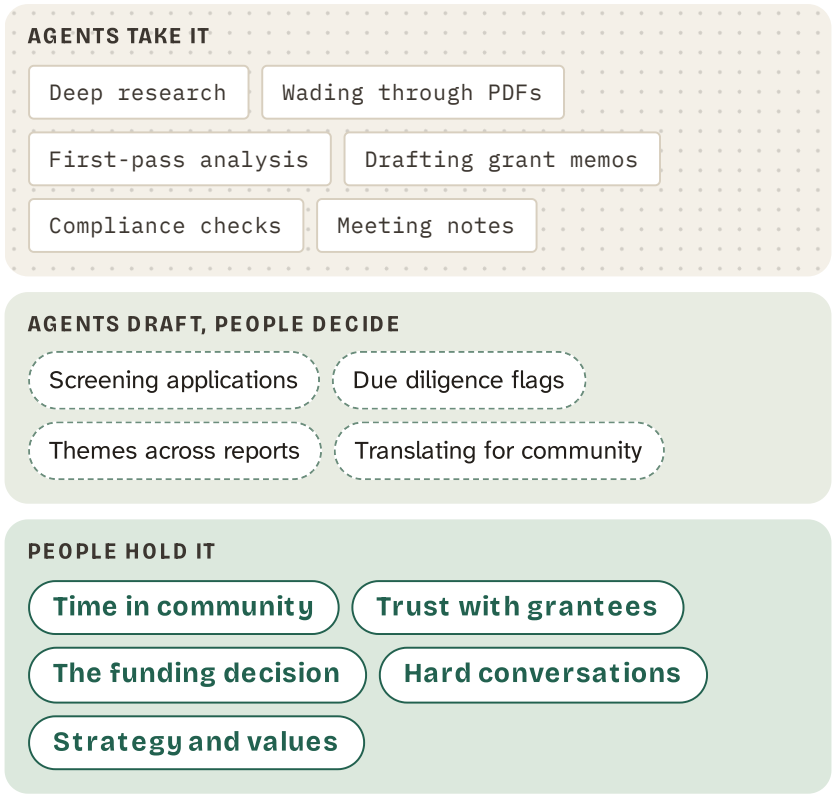



Tools are the easy part. Adoption is the work. We make a range of models, workspaces, and harnesses (the setups of tools, instructions, and context that let AI agents do real tasks) widely available, and we encourage everyone to use them, because this is how knowledge work will be done. More important, we have an in-house coach who specializes in exactly this transition. The coach sits side by side with each colleague to understand their workflows, build harnesses with them one workflow at a time, and develop the skills for a reality in which each of them has an army of agents, freeing them for the hardest problems and for what gives them the most purpose and joy. We use enterprise accounts with privacy controls, and we keep sensitive private financial data out of the models. Other institutions can do the same with internal talent or external partners.

Building is part of our job description now, across all functions. Our team has built a portfolio model of our path to sunset in 2035 and an AI navigator prototype that walks first-time homebuyers through the programs they may qualify for. In our childcare affordability pilot, OCR and AI cut document review time by about 40 percent.

At the same time, we try to embody a human-centered workplace. Well-integrated agentic workflows can pull people away from their computers and toward one another. Only a group of humans, together, can solve the problems we are uniquely positioned to take on.

## HAND THE SLOG TO THE AGENTS




Illustrative, not measured.

**COMMUNITY VOICE, BUILT IN**

Building faster only helps if the people it serves shape what gets built. Through the Colorado Design Insight Group (CDIG), a Gary venture, Colorado families are paid to share their experiences and co-design the policies, programs, and services that affect their lives.

STEP 2 OF 5

# Good Tokens

**LEAN INTO**

Backing the good uses of AI aggressively, and a shared social impact eval.

**GUARD AGAINST**

Extractive uses that arrive first and take the most from those already struggling.

The question of good tokens is one of the foundational challenges of AI today. Honestly, it is something the labs have been reluctant to talk about, and its effects are already visible. It plays out mostly in the application layer, the products and services built on top of the models. The common fear is that these applications will replace humans. The deeper issue is a split inherent in the application layer itself: there are already good applications of AI, and bad ones.

When we view AI as normal technology, albeit a transformational one, a familiar pattern emerges: throughout history, technology has been used for good and for terrible ends. The same ships that carried explorers to new lands, fostered trade, and connected civilizations also carried colonizers and enslaved people. The email and messaging that let us support someone we love who lives far away, and make human connections that would otherwise be impossible, have also been used to coordinate heinous terrorist attacks.

> “Beneath it all lies humanity’s capacity for joy and goodness, and its capacity for evil. **AI is no different.**

Such is the nature of technology, and of human progress. Beneath it all lies humanity’s capacity for joy and goodness, and its capacity for evil. AI is no different.

**IT CAN BE USED TO**

- ✓ coordinate life-saving care for a loved one who needs support;
- ✓ make public benefit programs accessible, so families don’t struggle to put food on the table;
- ✓ act as your own agent, negotiating better prices on essential services like insurance, filing your

**IT CAN ALSO BE USED TO**

- ✗ manipulate pricing algorithms to make essentials more expensive;
- ✗ aggressively deny insurance claims for people who desperately need them;
- ✗ accelerate gentrification and capture wealth by helping corporate buyers acquire homes from owners on the verge of foreclosure.

taxes on time, and claiming the credits you are owed.

These examples illustrate why philanthropy has a foundational role in making sure good tokens outnumber bad ones. It starts, obviously, with the philanthropic and nonprofit sectors actively using these tools ourselves, but it also means backing the good uses of AI far more aggressively. Otherwise, the actors working on behalf of communities will be the last to arrive, long after the extraction has happened, while market actors, with far more to gain, use these tools to take a greater share from those who are struggling. In fairness, the labs themselves are keen to ensure beneficial deployments and to accelerate positive uses of AI in areas like scientific research. Nevertheless, this fight is real, just as it was for the railroads, for telecommunications, for the internet, and for every technology that came before.

One significant opportunity philanthropy has barely touched is evaluations, or evals: the standardized tests that measure what AI models can do. There is not yet an industry-standard eval for social impact. Given how hard frontier labs push to optimize their models against evals, a joint investment by philanthropy in a strong social impact eval could reward the models that perform best on what matters to communities: helping people access public benefits, showing less racial and gender bias, and protecting kids and communities, to name only a few dimensions.

Promising early efforts exist. MIT's Media Lab is building open benchmarks for AI's impact on human well-being, and a new Public Benefits Bench from Vals AI, the Center for Civic Futures, and Code for America now tests how well models answer real SNAP questions. But none has reached the scale that makes the labs take it seriously. Philanthropy can change that by backing a social impact eval with the rigor and reach to sit alongside the benchmarks labs already report with every new model, like ARC-AGI for reasoning or SWE-bench for coding.

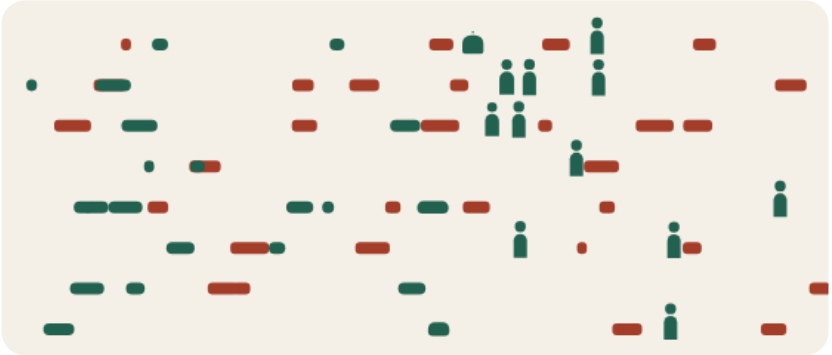
| | | |
|---|---|---|
| **ARC-AGI** | reasoning | REPORTED WITH NEW MODELS |
| **SWE-bench** | coding | REPORTED WITH NEW MODELS |
| **A social impact eval** | benefits access, bias, protecting kids and communities | THE MISSING ROW |

The benchmark table labs publish with every new model has no row for what matters to communities.

## THE APPROACH AT GARY

Our work on good tokens centers on the experience of struggling households: making sure they can reach the services they depend on, and building systems that are accessible and easy for them to use. We helped pioneer benefits-navigation technology through MyFriendBen, which has helped more than 55,000 Colorado households apply for $33 million in benefits. We are now exploring similar models in housing and homeownership and in supporting families with young children. We are investing in helping public schools adopt AI, while also pushing for more opportunities for young people to step away from technology and build skills in the real world. And through deep partnerships, especially with the State of Colorado, we are working to help public systems leapfrog into an AI-first world, where the experience of residents and safety-net participants is seamless and dignified.

### TIP THE BALANCE OF TOKENS



Good 65% Extractive 35%

**55,000+**
Colorado households helped to apply through MyFriendBen

**$33M**
in benefits applied for

The stream is illustrative. The figures are from the MyFriendBen case study.

STEP 3 OF 5

# Community and Resources

**LEAN INTO**

Balance sheets that speed the renewable transition, communities that share in the upside of AI infrastructure, and funding IRL: in-person spaces for human connection and belonging.

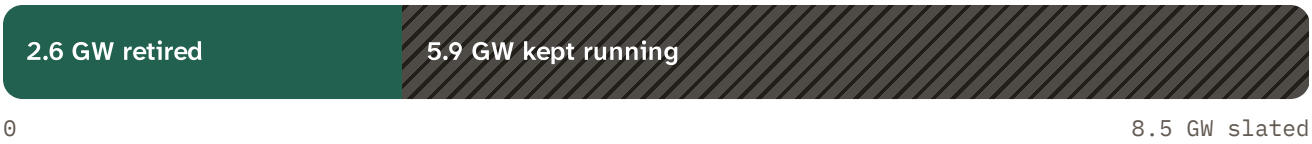
**GUARD AGAINST**

Higher energy bills, delayed decarbonization, host communities that bear the costs without the benefits, and deepening isolation as more of life moves onto screens.

The critics are right that some of AI's most pernicious effects are being felt by the communities hosting massive new data centers and superclusters. Data centers already use nearly 5 percent of the country's electricity, a share Lawrence Berkeley National Laboratory projects could more than double by 2030, and in some regions the costs are landing on household bills. On PJM, the grid serving much of the Mid-Atlantic and Midwest, the independent market monitor attributes $29.4 billion in capacity costs, nearly half the total across the last four auctions, to data center demand.

The industry's thirst for power is also slowing the retirement of carbon-intensive plants. Only 2.6 of the 8.5 gigawatts of coal capacity slated to close in 2025 actually retired, the smallest amount in fifteen years, and Georgia Power has cited data center growth in keeping coal plants running into the late 2030s. Every delay slows the shift to renewable energy, with significant long-term consequences for the climate.

**COAL CAPACITY SLATED TO CLOSE IN 2025**




Source: U.S. Energy Information Administration. The smallest retirement total in fifteen years.

On energy, philanthropy's role is to keep our investments centered on the renewable transition, and to accelerate the role our balance sheets can play in it. It is shocking that even today, among foundations that practice impact investing, the median puts only about 5 percent of its endowment to work that way, and across the sector the share is likely lower still. There is no better time than now to invest in renewables: doing so helps meet AI's surging demand while accelerating the broader energy transition we have all been hoping for. Many foundations have already made this a core priority.

On the community side, the relationship between the AI industry and the places that host it is deeply broken (local opposition blocked or delayed nearly $200 billion in data center projects in the first half of 2026 alone), and philanthropy can play a profound role in re-brokering the social contract. Part of that means reimagining community benefit, shifting from tokenism to a real share of the upside. Imagine true community ownership of data centers, where residents decide together, knowing there is shared sacrifice but also shared economic benefit. Philanthropy has pioneered models like this in land use and large-scale urban development, yet we have largely sat out this moment, missing the chance to bring those innovations to a better social contract.

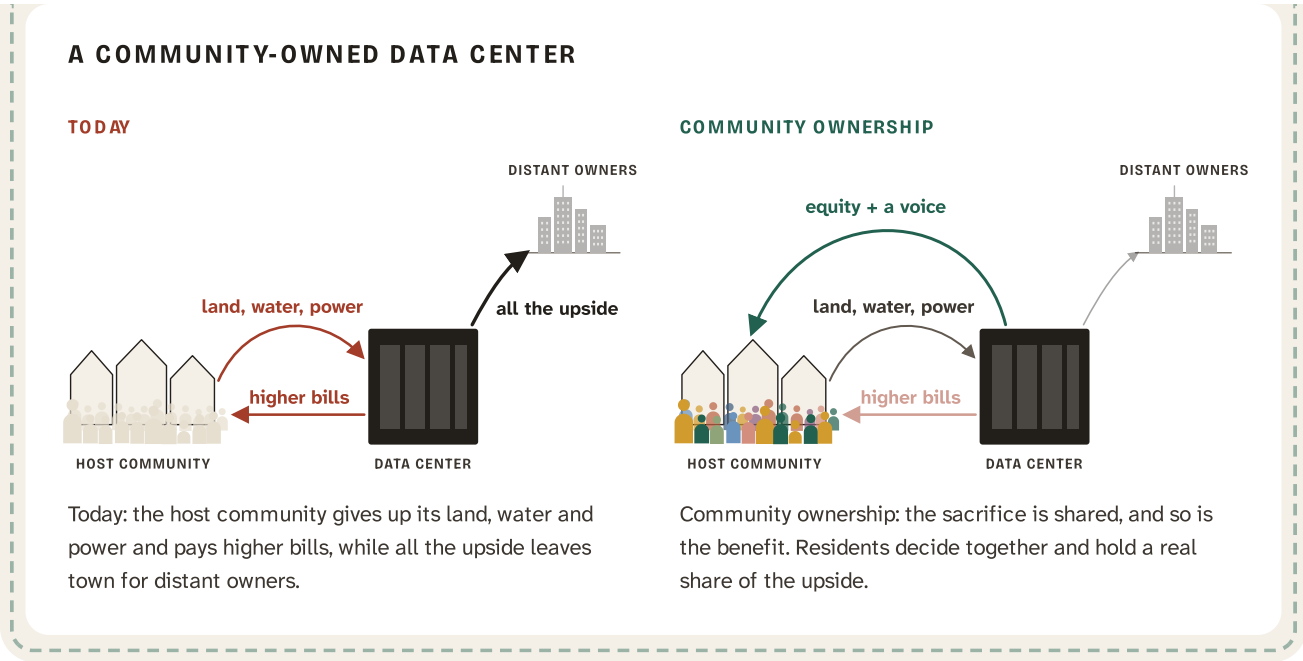
There is another dimension of community that this technology makes more urgent: our connection to one another. As AI takes on more of our work, our learning, and even our companionship, time together in person grows more precious, and easier to lose. The U.S. Surgeon General was already warning of an epidemic of loneliness and isolation in 2023, with about half of American adults saying they feel lonely. In this moment, more than ever, philanthropy should do more to bring people together: to convene, to create community, and to create spaces where we can foster human connection and belonging. I believe this is a foundational opportunity, so here is my ask of the field: fund IRL. Fund the gatherings, places, and programs where people meet in real life.

> **“ I believe this is a foundational opportunity, so here is my ask of the field: fund IRL.**

### THE APPROACH AT GARY

Part of Gary’s portfolio has strong exposure to renewables, and to the extent we are burning tokens, I believe those investments offset the carbon we currently use as an organization. Our other bet is on the community side. We are exploring early pilots that reimagine community benefits agreements as true equity ownership, giving communities a share of the upside in large-scale developments. We are looking to translate those models to data centers, with a dual purpose: a meaningful voice for communities, and a share of the upside to soften the economic impact.

We fund IRL ourselves, too. Our belief in human connection is why we have anchored on a convenings practice and brought people together through gatherings like ASSEMBLE100. It is also why we piloted My Spark, which put $1,000 in families’ hands for after-school enrichment, so young people could get off their phones and into community, in real life. More than 4,000 Denver students took part, and six more Colorado communities are now exploring programs like it.




*A note on water.* I have deliberately left it out. Nationally, data centers' direct water use, about 17 billion gallons in 2023, is well under 1 percent of what US crop irrigation consumes, and much of that irrigation grows feed for livestock. The largest share of the country's water still goes to pre-AI practices in industrial agriculture and water-intensive animal protein. Local impacts are a different matter: roughly two-thirds of new data centers built since 2022 sit in water-stressed areas, and they deserve scrutiny, even though the stresses may have existed before the data centers were set up.

STEP 4 OF 5

# Economic Impact

**LEAN INTO**

Leapfrogging a safety net and tax system we wanted to fix anyway, and putting asset ownership at the heart of economic mobility.

**GUARD AGAINST**

Job disruption that deepens inequality and the affordability crisis.

AI's economic impact is, I believe, an immediate threat to our communities, one that could accelerate the inequality and affordability challenges they already face. This is an area where the labs have been leading. Anthropic and OpenAI have each published economic policy proposals, and economists at Google DeepMind have explored similar ideas. Roughly summarized, they converge on three:

1. helping people navigate the shift in the job market;
2. building a better safety net, especially for people who lose work;
3. giving people a financial stake in the upside, if most future economic output comes from models and machines.

I believe the latter two represent the real unlock: a chance to leapfrog systems we wanted to fix long before AI began to accelerate. Returning to access and equity, that means transforming the safety net so people can actually reach it, and reimagining the tax system so there is a greater cushion for those who are hurt, delivered effectively through tax credits and safety-net programs. The broader tax base created by the AI boom could accelerate both.

> **"The real divide is between households that own assets and those that don't, and AI could make it dramatic, with enormous consequences for our most vulnerable community members.**

The shift toward capital was under way before the AI boom. Labor's share of income has been falling for two decades, and in mid-2026 it hit 52.8 percent, the lowest since records began in 1947. Meanwhile, the top 1 percent hold a record 32.5 percent of household wealth. The real divide is between households that own assets and those that don't, and AI could make it dramatic, with enormous consequences for our most vulnerable community members. As the Harvard economist Richard Freeman put it, "who owns the robots rules the world."

52.8%

Labor's share of income, Q2 2026. The lowest since records began in 1947.

Bureau of Labor Statistics

32.5%

Share of U.S. household wealth held by the top 1 percent, a record.

Federal Reserve Distributional Financial Accounts

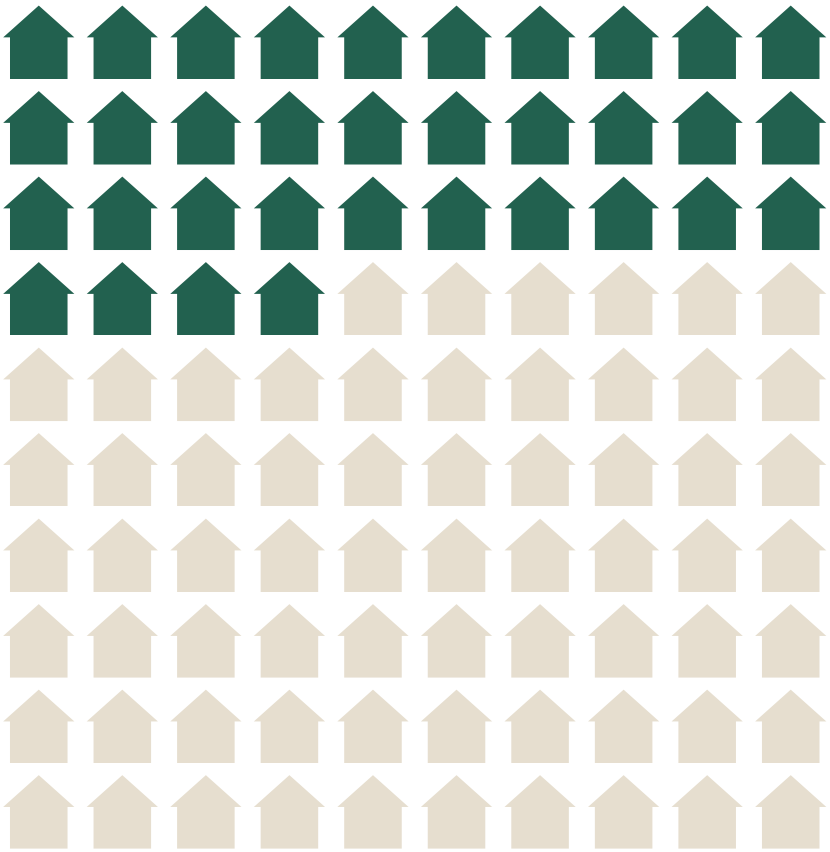
For philanthropy as a whole, this should be a wake-up call: a call to reorient our work toward assets, not just income, and to use this moment to renew our focus on universal access to wealth, through policy innovations that fundamentally strengthen the balance sheets of America's most vulnerable households.

### THE APPROACH AT GARY

Our significant bets are on a better safety net and a share of the upside. We are invested in a partnership with the State of Colorado to help reimagine our safety net and how benefits are delivered. We are also helping to shape the state's tax structure to protect those who may fall behind. Colorado families claimed more than $810 million in the first tax season of the Family Affordability Tax Credit, which we helped pass.

For the last nine years, our biggest bet is on policies and structural solutions that make asset and capital ownership universal in Colorado. It is a core part of our wealth and ownership strategy, which aims to bring hundreds of thousands of Colorado households into essential wealth, where they own at least a reasonable, diversified pool of assets, whether a home, a 401(k), an early wealth account, or a stake in a small business. By the measure we built with the Aspen Institute, roughly two-thirds of Colorado households fall short of it today. I want to help move the rest of the field in this direction, too.

**ESSENTIAL WEALTH IN COLORADO**



**34 of every 100** Colorado households have essential wealth today.

- A home
- A 401(k)
- An early wealth account
- A small-business stake

Source: Colorado Essential Wealth Metric, Aspen Institute Financial Security Program with Gary Community Ventures.

STEP 5 OF 5

# Existential Risk

### LEAN INTO

Governance and safety, led by the national and global funders best placed to do it.

### GUARD AGAINST

Models that turn on us or are weaponized to take down critical systems, and, for place-based funders, distraction from the immediate harms at home.

The AI doomsday scenario has drawn a great deal of recent media attention. These risks go beyond those discussed in the previous sections: the chance that models gain agency, become sentient, and turn against humanity. There is also the risk that, without the right protections and guardrails, models are weaponized by state or non-state actors to shut down critical systems and cripple our infrastructure.

This is the area that has drawn the most philanthropic interest, rightly so, especially from larger funders with national and global reach. Philanthropy can play a foundational governance role here: by sitting on the cap tables of the leading labs, and by advocating strongly for pacing the frontier and for strong regulatory oversight of frontier labs. As a sector, we can also convene coalitions and fund the leading institutions anchoring the work on AI safety.

### THE APPROACH AT GARY

I would argue that for a place-focused, community-obsessed institution like Gary, focusing on the existential, apocalyptic threats of AI (which are real) can be a distraction from AI's immediate community-level and societal impacts on our economy, jobs, energy prices, and affordability, impacts that undermine our transformational work with the communities and people we care about most, in our immediate vicinity. This is not to shrug off responsibility but to share it: to know where we are well equipped to be proximate and make a meaningful difference, and to inform the sector as a whole, while letting our larger national and global peers take on the governance and safety work that is so desperately needed.

## SHARING THE RESPONSIBILITY




Select a ring. Each kind of funder takes the ring where it can be most proximate, and informs the other.

CODA

# We Are Made for This

I would argue that the tectonic shift AI is bringing is, in some ways, the best thing that has happened to philanthropy.

Don Chen, president of the Surdna Foundation and one of my favorite fellow philanthropic leaders, likes to say that philanthropy has no natural predator. Without one, our institutions are prone to complacency, slowness, bureaucracy, and a lack of perspective, even as we have grown unimaginably large, and are expected to grow larger still, in the scale of our assets and the privilege we hold.

AI has the potential to reorient us toward our core values: centering communities and people. As an institution, we are tailor-made for this moment: a bridge between the machine builders and the community organizers (the tech bros and the teachers). We can bring everything we are, as the deepest lovers of humanity, as problem solvers, and as builders, to shaping a better AI future.




**We are made for this.**

---

## Tokenmaxxing for Good

A Framework for Philanthropy's Role in AI, Inspired by Gary Community Ventures.
Santhosh Ramdoss, President & CEO, Gary Community Ventures. September 2026.

Cite as: Ramdoss, S. (2026). *Tokenmaxxing for Good: A Framework for Philanthropy's Role in AI.* Gary Community Ventures.