Coefficient Giving

2025 Letter from the CEO
March 12, 2026

Coefficient Giving directed over $1 billion in 2025, the most in our history. This is a significant milestone for us as an organization. It's the result of extraordinary dedication from our funders, staff, and grantees, who share the conviction that effective philanthropy can be a powerful lever for making the world a much better place.

In 2025, we set out to lay the groundwork for giving at a greater scale, and — as I told our team at an all-hands recently — I honestly think we crushed it. We increased our total giving by over 50%; onboarded an almost totally new leadership team; launched two new funds and changed our name to Coefficient Giving to reflect our growing multidonor orientation; and more than doubled the funding we directed from non-Good Ventures donors to over $200 million.

Our goals for 2026 are even bigger. We aim to significantly scale our funding again, drive more proactive and ambitious grantmaking to help fill the most important gaps we see, and prepare for even larger growth in 2027 and 2028. We want to scale thoughtfully, growing a lot without losing our ability to make big decisions quickly, becoming too bureaucratic to pivot when new opportunities arise, or losing our willingness to bet on neglected (and sometimes weird) topics.

Two and a half months into the new year, the work still ahead of us feels a little daunting. On the supply side, Good Ventures — our largest funding partner — is eager to accelerate their giving, and the AI boom seems poised to drive a surge of new philanthropy. On the demand side, we see more and more compelling giving opportunities, driven by the advent of new technologies, the retrenchment in government development aid, and the payoffs to our investments in staffing up and exploring new causes over the last decade. And accelerating progress in AI only heightens the stakes.

I've decided to write a more personal reflection this year. Below, I cover what a decade of grantmaking has taught us about the compounding value of early bets, how I'm thinking about the pace of AI progress, our push to ambitiously own big problems, building our team to meet the moment, and becoming more agile as we scale.

In order to deliver on our mission, we need excellent people across the board. If you've ever thought of working with us, now is a great time. Please let us know what you're excited to work on here and get updates on open roles from our monthly newsletter.

Early bets compound

J.C.R Licklider wrote in 1965 that "people tend to overestimate what can be done in one year and to underestimate what can be done in five or ten years." We're now over a decade into Coefficient and this line is more and more front of mind.[1]

[1] Though I wonder how true this will remain as AI speeds up.

For years I had struggled to crisply articulate our track record, but in preparation for our rebrand in November, the team pulled together a short summary of our biggest wins that make me feel enormously proud:

- Our grants to evidence-backed global health interventions have saved over 100,000 lives, and will save tens of thousands more in the years to come.
- We supported the late-stage clinical trials for the R21 malaria vaccine, now being rolled out to protect millions of kids globally.[2]
- We were the earliest major funder of the YIMBY movement to build more housing. Our grantees have led the charge on major wins like City of Yes in New York, and SB 79 in California, which will enable hundreds of thousands of new housing units.
- Beginning in 2015, seven years before the launch of ChatGPT, we jumpstarted the field of AI safety and security, supporting researchers and organizations that have gone on to do field-defining work in alignment, control, and governance.
- Our work so far has improved the lives of over 3 billion farm animals, largely by driving commitments from corporations to improve conditions in their supply chains.

[2] Though our program officer Katharine Collins, who co-invented the R21 malaria vaccine as a grad student, deserves a lot more of the credit than our later funding does.

As I've discussed elsewhere, the compounding benefits of early work feel especially pronounced in AI safety and security. When we began funding over a decade ago, very few philanthropists were thinking seriously about the risks from rapid AI progress. Because it was such a small field, we primarily funded talent pipelines and early-stage research centers, giving grantees a lot of flexibility to explore different approaches to mitigating potential worst-case risks. This funding felt highly speculative at the time, and we were bottlenecked at first by both a lack of talent and strategic clarity on the issue.

Many of the grantees that have gone on to be among our most important and impactful didn't start off looking that way at all. For instance, we made our first $250,000 grant to the program that would eventually become ML Alignment & Theory Scholars (MATS) in 2019, when it was a side project by some students affiliated with the Stanford Existential Risks Initiative who thought there should be a summer program to prepare software engineers for careers in AI safety. The MATS 1.0 cohort had 5 fellows and no permanent full-time staff. They have since expanded to run multiple cohorts

a year of around 100 scholars with an admission rate of 4-7%, and report that over 80% of their alumni are now working full-time in AI safety and security (accounting for a meaningful portion of safety staff at some of the biggest companies and government institutes).[3]

Another example here is the YIMBY movement. In the early and mid-2010s, many academics believed that zoning restrictions artificially increased housing costs, but there were almost no organizers and advocates making this a salient issue for policymakers. As then-program officer for our U.S. policy work, I made some small initial grants to groups like CaRLA and CA YIMBY and the first YIMBYtown conference, which was enough to make us the biggest funder of the nascent movement. And as with MATS, the early grants that ended up high-impact didn't always start off seeming especially propitious — I'll always remember writing this slightly tongue-in-cheek footnote in the public writeup on our original grant to CaRLA. Most of the work remains ahead, but the YIMBY movement now feels ascendant.

[3] They now count a few dozen full time staff, including my wife, who joined in December as a research manager after a decade as a data scientist at Uber.

YIMBYtown attendance and pro-housing policy adoption over time

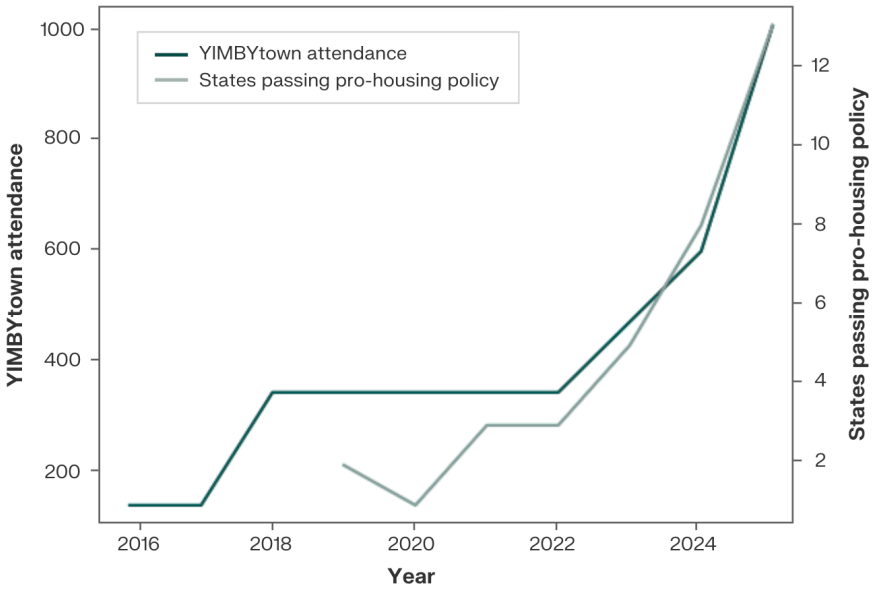

A highly scientific analysis that a grantee shared with us

The same compounding growth also shows up in the trajectories of our teams. We hired Lewis Bollard to lead our farm animal welfare giving in 2015, and wrote publicly at the time about our worry that, just a couple years out of law school, he might not have enough experience to succeed in the role. It's honestly hard to comprehend the impact he and his team have had since then. Their grantmaking has already improved the lives of billions of animals, and if the reforms they've helped secure stay in place (even putting aside future wins), we project that they will improve the lives of another 10 billion animals over the next 10 years.

Confronting rapid AI progress

Over the past year and especially the last couple of months, I have been struck by how much more accurate the predictions of people with short AGI timelines have been than mine. There was maybe a lull in mid-2025,[4] but advancements in AI capabilities seem to have accelerated in recent months with Opus 4.6 and Gemini 3. Epoch's Capabilities Index shows no sign of slowing, and METR's time horizon measurement appears to be speeding up.[5] We are starting to see evidence of AIs doing significant chunks of engineering work themselves, like Claude Code helping build Claude Cowork and a C compiler. New models are reaching benchmarks much faster than most superforecasters or domain experts predicted.[6]

[4] Sometimes it feels like timelines get longer in the summer and shorter in the winter.

[5] METR is struggling to continue their time horizon measurement because the AIs are saturating their work tasks.

[6] Over the last couple of years, domain experts have been more accurate than superforecasters in predicting the rate of progress.

Time horizon of software tasks different LLMs can complete 50% of the time

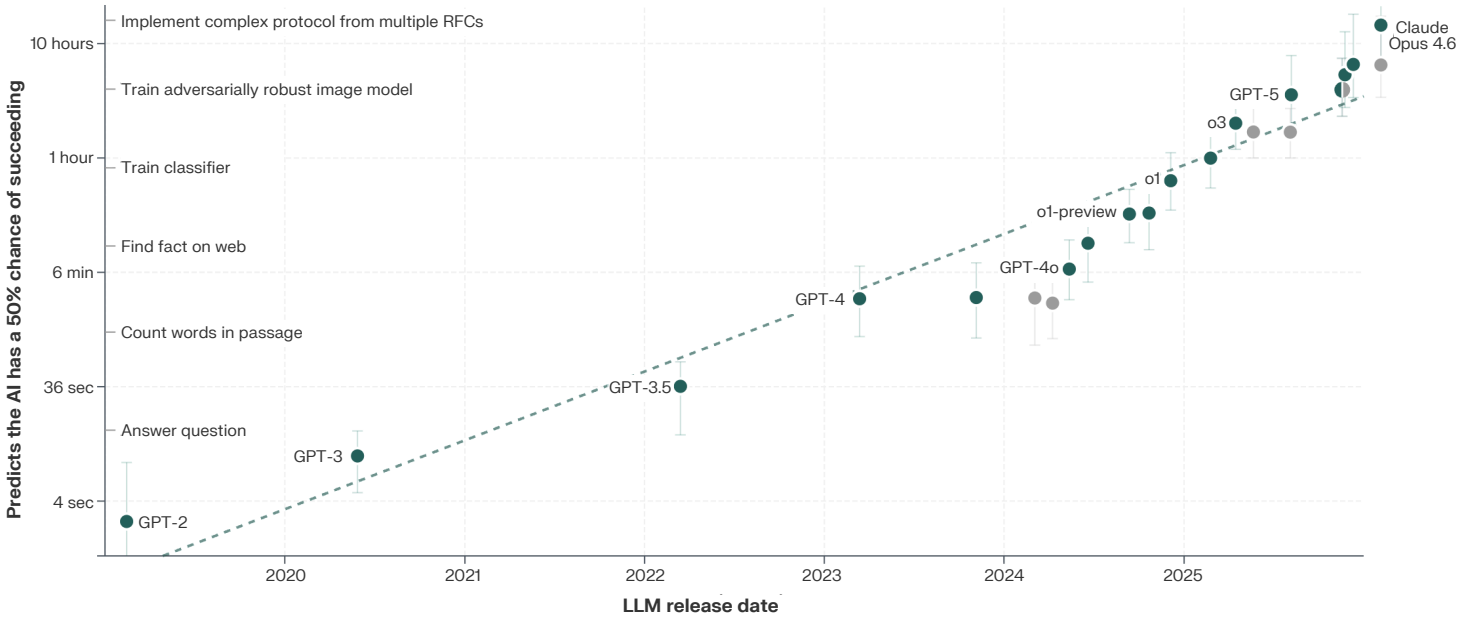

Source: METR (2026)

For most of my time at Coefficient, I worked on our global health and wellbeing grantmaking, first on researching new causes, then U.S. policy, and then overseeing the entire portfolio. I was always glad that my colleagues were focusing on risks from AI, but I was initially skeptical that this work was as important or tractable as they thought. I've started out each of the last few years with a view that rapid AI progress is not necessarily going to slow down, but on some level expecting that it probably would. With each passing year of that implicit expectation being dashed and capability measurements continuing to move up and to the right, continued fast progress increasingly looks to me like the default path rather than an aberration.[7]

For now, Coefficient is pursuing a diversified approach to AI giving, spreading our bets and supporting work that seems robustly good across many different futures. We don't have a single theory of change for how to

[7] Of course, this is exactly the sort of extrapolative expectations dynamic that can drive bubbles.

make AI go well. But our leadership and AI team leads have definitely been paying more focused attention to the possibility of short timelines in recent months. One of our goals early in 2025 was to prepare to respond more quickly and effectively if we do find ourselves in short timelines worlds — Claire Zabel is now leading that work, and is currently building a team and triaging projects that could be useful. And we're looking out for changes across our portfolios to act with greater urgency given the possibility of short timelines.

Our Technical AI Safety (TAIS) team rose to meet the urgency of the moment after the slow start I discussed last year. In early 2025, we saw that better AI capabilities were making empirical safety work more tractable, and as a result that the main bottlenecks to technical AI safety progress were no longer useful research paradigms, but the availability of funding and talent. The TAIS team responded by releasing a detailed RFP on several priority research directions, which helped more than triple our total TAIS giving from $46 million in 2024 to $157 million in 2025, across more than 200 grants. To enable the team to move faster, we changed our grant approval process for this RFP, and other teams are now considering something similar.

As we look ahead, there's a tension between increasing funding for existing grantees and creating new highly absorptive organizations that can scale over time. Both compete in the near term for scarce grantmaker time, and our experience is that helping create new grantees, while far more labor intensive per dollar moved at the beginning, can pay bigger dividends in terms of subsequent capacity and impact later.[8] We expect different teams to tackle this tension differently this year: some will focus more on steadily expanding support for existing grantees and responding to inbound opportunities, while others will work to more proactively seed high-capacity organizations that could grow substantially in 2027 and 2028.

This could mean that our AI giving doesn't actually increase dramatically in 2026. While we hope it does, we'd prefer to prioritize longer-term growth if the two come apart.

Grappling with scale

Our mission is to help others as much as we can with the resources available to us, and historically, that's meant operating in a "marginalist" mode, thinking hard about how to equalize marginal returns within and sometimes across cause areas. (The idea is that each additional dollar we spend should produce roughly the same amount of good if allocated to any of our funds, because if it doesn't we should be shifting resources toward the higher-return area until it does.)

[8] There are also several examples we've come across in the history of philanthropy of funder-intiated start-ups — like Ed Scott helping start the Center for Global Development — that grew to fill important ecosystem gaps.

The trade-offs that drive a focus on equalizing marginal returns remain. Funding still falls well short of the opportunity set across our focus areas, and we still face some zero-sum budget decisions across worldviews and funds.

But with nearly every fund working with a higher budget in 2026, and more growth expected in the future, we need to shift more of our attention away from marginal trade-offs, and toward more ambitious goals. Many of our best grants have been deeply inframarginal: we had to recruit a founder or — effectively — incubate an organization, but once we had and there was finally something fundable, the eventual grant was far above our bar. The kind of dedicated and proactive ownership from program staff required to enable these grants can pay off in a big way, especially over time, but it trades off with a focus on optimizing the marginal dollar in the near term.

I thought the importance of taking full responsibility for a problem was captured well by Nan Ransohoff's piece "There should be 'general managers' for more of the world's important problems." Lewis, who Nan also points to in her post, exemplifies this in his work on farm animal welfare. Another colleague who really embodies this spirit is Andrew Snyder-Beattie, who runs our biosecurity and pandemic preparedness work. Andrew joined Coefficient in 2019 and has taken it upon himself to reduce risks from worst-case pandemics with unusual and impressive dedication and purpose. Recently, this has included work on reducing risks from mirror bacteria and a four-pillared plan to avoid engineered biological catastrophes.[9]

I'm inspired by these examples and want Coefficient to continue being a place where outstanding people's responsibilities and ambitions can grow to match the scale of the world's most important problems. However, I also want us to avoid the common "strategic funder" trap of thinking we have all the answers and just need to slot grantees into our vision, almost like subcontractors. An important virtue of the marginalist approach to funding is that it's relatively strategy-agnostic: if an opportunity clears the cost-effectiveness bar, you should fund it. That tends to facilitate a healthy openness to grantees' stronger local context and knowledge.

I want us to navigate this tension as thoughtfully as possible, bringing the ambition and ownership of general managers without losing the curiosity and humility of marginal funders. The right balance will vary across focus areas and individuals, but we should be intentional about trying to strike it.

[9] The four pillars are (1) personal protective equipment like elastomeric respirators, (2) biohardening of buildings and infrastructure (3) pathogen-agnostic detection systems, and (4) rapid medical countermeasures. For more, listen to a recent interview with Andrew on the 80,000 Hours podcast.

Growing our team

The rest of the Coefficient leadership team were less than 6 months into their roles at the beginning of 2025, and I had only become sole CEO a little bit earlier. We're now a year into this new incarnation of the leadership team, and I feel incredibly lucky to work with such an outstanding group of people. They each bring a deep commitment to our mission, excellence in their respective domains, and genuine good humor.

Our Leadership Team at our 2025 org-wide retreat

We also welcomed a wave of talented new colleagues across the organization last year. Our Communications team has doubled in size since the end of 2024, and our Partnerships team grew from 2 to 10, which will be increasingly important as we grow our work with a wide range of donors. Matt Clancy brought on what feels like my entire Twitter feed to the Abundance & Growth Fund with Dylan Matthews, Alex Armlovich, Willow Latham-Proença, and Saloni Dattani joining (Saloni in a part-time volunteer capacity). Justin Sandefur also joined us from the Center for Global Development to lead our new Global Growth Fund. I'm excited that we're (finally) doing more to accelerate economic growth, and was excited to see both funds hit the ground running last year. We also hired Stephanie Hill as our new Head of People — she had previously helped GiveDirectly scale from 250 to 950 people, which bodes well as we aim to bring on ~70 new staff this year.

On the flip side, some of our longest-tenured staff focused on big-picture AI strategy departed last year, including Tom Davidson, Ajeya Cotra, and Joe Carlsmith. Coefficient has benefited a huge amount from their work, as have I personally, and I'm very sad to lose them. But it's a good sign that there are more places where people can do high-level work preparing for a

world with transformative AI, and that research organizations like Forethought and METR can provide an excellent home for it. While I'm glad to see this work thrive elsewhere in the ecosystem (and that we can support it in some cases), I also feel a growing internal need for researchers who can help us translate that big-picture external work into grantmaking strategy. For instance, the gap between a high-quality theoretical paper on "how can AI labor be used for AI safety and security" and actually developing a funding plan is large.[10]

[10] Rebuilding an AI research team to more effectively inform our grantmaking is a potential priority for this year.

One unusual challenge we face, relative to most other philanthropies, is more directly competing with companies for talent: some of the people we're recruiting see working on safety at big AI companies, where they could earn millions of dollars a year, as comparably impactful to our grantmaking roles. We're not going to match the companies' total compensation, but we're considering some changes to make our pay more competitive, and expect to support others in the AI safety ecosystem to do the same.

Becoming more agile

As we grow, direct more funding, and work more with other donors, there will be a natural and somewhat appropriate temptation toward increased risk aversion and caution. I want to make sure we strike the right balance: appropriately weighing the value of being able to work more with others against the ongoing need to make occasionally weird-seeming bets on underappreciated topics and respond with agility to rapid changes.

Staying nimble partly comes down to process improvements like pushing more autonomy down in our org chart and having shorter investigations and approval processes for lower-risk grants. But I think most of what success looks like on this front has to do with culture. We've taken inspiration from Netflix's mantra about minimizing rules and bureaucracy with scale by continuing to invest in ever more high-performance people, and we're trying to follow the same path.

We also want to continue to invest in efficiency as we scale. Compared to other philanthropies of roughly our size where we've been able to source data, we are somewhat unusual in terms of both moving more money and making a large number of grants per staff member, and I'm proud of that combination. (The couple big funders that move a lot more per staff member tend to do a lot more regranting. We agree regranting can be very impactful and we do it at scale when we see a distinctive case for it, e.g. with GiveWell or our Regranting Challenge. But I don't think that moving grantmaking staff from our payroll to a grantee's is necessarily a big efficiency win.)

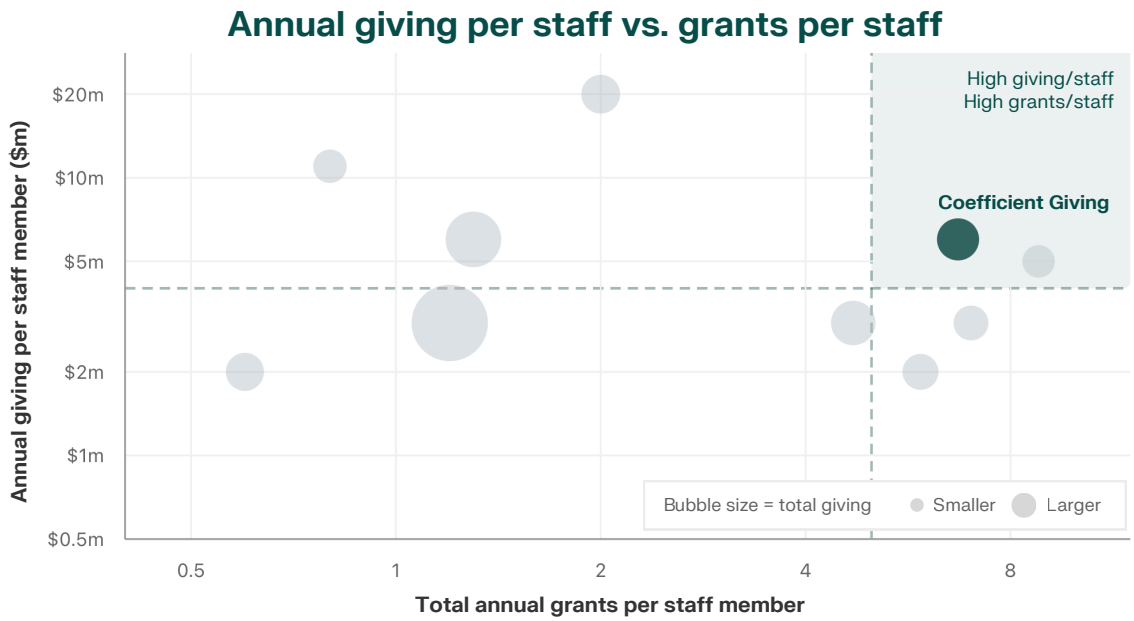

We'd like to maintain our relatively light staffing ratio, even as we take on the complex task of working with more donors, which will require some additional headcount. One concrete way of doing that is to accelerate our internal AI adoption. Chris Webster is leading the charge on AI enablement across the org, and it is early days, but I've been heartened by the initial take-up and sharing of new tools (e.g. Claude skills for grant writeups or cost-effectiveness calculations).

Closing thoughts

Most of this letter is about trying to navigate some real tensions with good judgment, rather than resolving them once and for all. But based on my last few years running CG, I think recruiting and empowering exceptional people is the closest thing we have to an unalloyed good. Our President Emily Oehlsen,[11] who oversees our work on global catastrophic risks, and Otis Reid, who leads our work on global health and wellbeing, both joined Coefficient in 2021 as individual contributors and quickly took on much larger roles, transforming their teams' work and enabling us operate at a much greater scale. Naina Bajekal began running our Communications team in July 2024, and in 2025 led a successful rebrand that has already helped us expand our work with other donors. I'm very proud of what they (and many others at Coefficient) have built in the past few years, and feel deeply that we'll need many more people willing to take on similar kinds of responsibility if we're going to execute on the opportunities ahead of us.

If you've ever thought about working at Coefficient, now is the time to apply. I've focused heavily on AI in this piece, where hiring is urgent given the rapid pace of progress. But our needs extend across everything we do: from biosecurity and science to partnerships and operations. We face a unique opportunity to deliver enormous positive impact in the years ahead, and I hope you'll consider joining us.

[11] If you want to get a sense of Emily's enormous virtues as a leader; I recommend listening to this podcast with our former colleague Ajeya Cotra. I teared up listening to her description of working with Emily.