What's your endgame in the age of AI?

Ben Metz, with Claude

April 2026

A response to Gugelev & Stern, "What's Your Endgame?" (Stanford Social Innovation Review, Winter 2015)

A note on this draft

This paper is a work in progress. I'm sharing it now to get feedback before developing it further, not because I think it's finished.

The most significant gap is geographic. The 11 endgames described here are written primarily from a UK and US nonprofit perspective. Several of them, particularly government adoption and commercial adoption, don't translate well to other geographic contexts. This applies to a large proportion of civil society globally, and it means the framework as it stands has a Global North bias I haven't yet resolved.

I'm also exploring whether additional endgames are needed. For instance, one example is and endgame that captures the work of organisations whose primary function is to generate ground truth that would otherwise not exist, rather than to apply or validate expertise that does. This emerged from conversations about civil society operating in contexts where the information AI systems can reach is structurally incomplete. I'm still working out whether it's sufficiently distinct from the endgames already described.

I intend to do more research on both of these gaps before producing a final version. That research is now under way. If you have relevant experience or evidence, I'd welcome it.

In the winter of 2015, Alice Gugelev and Andrew Stern published one of the most useful pieces of thinking the nonprofit sector has seen in recent decades. "What's Your Endgame?", published in the Stanford Social Innovation Review, made a simple but devastating observation: most nonprofits are asking the wrong question. By focusing on "how do we scale up?" they are pursuing a goal, specifically organisational growth, that the structural conditions of the sector make nearly impossible to achieve. The "social capital chasm", the funding gap between early-stage organisations and the scale required to match the problems they address, has remained stubbornly in place. Most nonprofits will never cross it.

Their answer was to reframe the question. Instead of asking how to grow, leaders should ask what specific role their organisation intends to play in the overall solution to a social problem. They called this its endgame. They proposed six: open source, replication, government adoption, commercial adoption, mission achievement, and sustained service. The framework has been widely used in the decade since, and rightly so.

I've been returning to it recently in the context of research I'm conducting on AI and civil society, specifically on how AI-driven disruption threatens the operating models, funding relationships, and workforce pipelines of NGOs and charities. What I've found is that the question Gugelev and Stern identified is now more urgent than ever, but the menu of available endgames has fundamentally changed. Some of the original six have been strengthened by AI. Others have been hollowed out or bifurcated. And five new endgames have emerged that didn't exist, or didn't need to exist, when the original paper was written.

This is my first attempt to update the framework. I want to be clear from the outset: this is a response to Gugelev and Stern's work, not a replacement for it. Their central insight stands. What I'm offering is an update to their menu.

Why the question is more urgent now

Gugelev and Stern's 2015 paper diagnosed the social capital chasm as the central constraint facing nonprofits. The funding architecture of the sector, including grant cycles, overhead caps, and bias toward programmes rather than organisational capacity, made it structurally difficult for organisations to accumulate the capital needed to reach meaningful scale. They were right. A decade later, that architecture has not fundamentally changed.

But AI has begun to change something else: the economics of knowledge work. The research, policy analysis, communications, grant applications, impact reports, and advocacy documents that constitute the primary output of a large proportion of the nonprofit sector can now be produced at dramatically lower cost. Inference costs for AI systems have fallen 280-fold in two years (Epoch AI, 2024). The coding benchmark that stood at 4.4% accuracy eighteen months ago now stands at 71.7% (Stanford HAI, 2025). In the legal profession, a reasonable proxy for knowledge-intensive nonprofit work, AI achieves 91.4% accuracy in contract analysis and 35–70% efficiency gains in document review [source to be confirmed].

This matters for the endgame question in a specific way. Gugelev and Stern were writing in a world where the primary constraint was scale: most nonprofits couldn't do

enough. AI is changing the constraint. For knowledge-producing organisations, the question is no longer primarily "can we do enough?" but "what is the distinctive value of what we do, now that AI can do a version of it at near-zero cost?" That is an endgame question. And unlike the social capital chasm, which was a structural feature of the funding landscape, it requires a strategic answer from each organisation individually.

> *The social capital chasm was a problem with the funding architecture. The AI disruption is a problem with the value proposition. Every knowledge-producing nonprofit now needs to answer: what do we offer that AI cannot?*

What AI does to the original six

Let me take each of Gugelev and Stern's original endgames in turn, briefly, and assess where AI leaves them.

Open source: amplified but hollowed out

> **GUGELEV & STERN, 2015**
> *A breakthrough idea that is easy for other organisations to adopt and integrate. The organisation conducts R&D, shares knowledge freely, and serves as a knowledge hub for its field.*

The open source endgame involves investing in R&D, developing a breakthrough idea, and sharing it freely through a knowledge hub. AI makes knowledge dissemination near-zero cost, which amplifies the reach of this model. But it simultaneously hollows out the core competency. If anyone can produce plausible-seeming research at scale, an open source CSO's distinctive contribution is no longer the production of knowledge but its validation. The endgame survives but mutates: the core competency shifts from research production to epistemic credibility. Organisations need to invest in the things AI cannot provide: domain expertise built through years of practice, relationships with affected communities, and institutional reputation for rigour.

Replication: bifurcated

> **GUGELEV & STERN, 2015**
> *A breakthrough product or model that is easy for other organisations to adopt and deliver. The organisation defines a replicable model, demonstrates its efficacy, and shares it - becoming a centre of excellence and certification body.*

Replication endgame organisations prove a model works and then help others deliver it. AI disrupts this in two directions simultaneously. It makes codification trivially easy, which is good for the model in the short term. But if the model itself involves knowledge work such as producing assessments, analysis, or information. AI may make the replication unnecessary by delivering the output directly. The endgame bifurcates: replication of human-relational models (community organising, service delivery requiring physical presence) is strengthened. Replication of knowledge-work models is undermined. The key question for any replication

endgame organisation is: are we replicating a human practice or an information process?

Government adoption: strengthened but slower

> **GUGELEV & STERN, 2015**
> *A model with high coverage potential and a capacity for integration into public sector programmes. The organisation delivers results at sufficient scale to make a case for government involvement, then serves as advisor and advocate.*

This endgame, where a CSO proves a concept and advocates for the state to adopt it, is arguably strengthened by AI. Governments are slow technology adopters. CSOs can serve as proving grounds for AI-enhanced public service models, absorbing the implementation risk and ethical complexity that government is structurally ill-equipped to handle quickly. The complication is that the advocacy pathway is harder in an information environment flooded with AI-generated content, and CSOs increasingly lack the technical capacity to make credible cases for AI-enhanced public services to technically sophisticated policymakers.

Commercial adoption: paradoxical

> **GUGELEV & STERN, 2015**
> *A product or service with profit potential that solves a market failure or reduces market risk. The organisation demonstrates viability, then hands off to commercial providers - continuing to target harder-to-reach segments and monitor quality.*

The commercial adoption endgame, addressing a market failure and then handing off to commercial providers, faces a paradox. AI is resolving some of the market failures that originally justified CSO intermediaries (particularly information asymmetries). At the same time, it is creating new ones at speed: algorithmic bias, digital exclusion, data exploitation. The market failures CSOs need to address increasingly are AI-related, and the handoff to commercial providers is structurally complicated by the fact that the companies best positioned to adopt the solutions are often the same companies whose products created the market failure.

Mission achievement: accelerated but harder to exit

> **GUGELEV & STERN, 2015**
> *Defined and achievable outcomes related to solving a discrete problem. The organisation maintains a tight focus on targeted intervention - and when the goal is reached, winds down or redeploys its assets to a new purpose.*

AI could dramatically accelerate mission achievement for organisations with well-defined, bounded goals, particularly where the bottleneck is information processing or coordination rather than political will. But the clean exit that Gugelev and Stern describe, declaring success and closing or redeploying, is messier when AI has been integrated into operations. Who maintains the systems? Who governs the data? The March of Dimes pivot (from polio to birth defects) has a new analogue: the most valuable remaining asset of a mission-achieving CSO may be its AI capability, not its domain expertise.

Sustained service: existentially challenged

> **GUGELEV & STERN, 2015**
> *A strong organisation, with a proven ability to sustain funding, that fills a gap the market and state won't provide. Core work: building a cost-effective model and making ongoing efficiency improvements. Budget is the only endgame that continues to grow.*

Sustained service is the most disrupted of the original endgames. It is the only one requiring continuous organisational growth, and it depends on a cost structure AI is compressing, a funding model AI is destabilising, and a workforce pipeline AI is thinning. The endgame bifurcates cleanly. Service-delivery organisations, such as homeless shelters, disability support, and domestic violence refuges, whose core work requires human presence and relationships face indirect pressure but not existential threat. Knowledge-service organisations, such as legal advice centres, policy consultancies, and research services, face a more fundamental challenge: if the service is information provision, and AI can provide the information at near-zero marginal cost, the justification for sustained grant-funded operation erodes.

CASE STUDY: The Knowledge-Service Squeeze

Consider a mid-sized legal advice charity providing welfare benefits guidance. Its operating model is built on knowledge work: specialist caseworkers who understand an arcane and constantly changing benefits system and translate it for people who cannot navigate it alone. AI systems can now answer most welfare benefits questions with high accuracy. The organisation faces a genuine strategic choice: does it continue as a sustained service (arguing that the human relationship, the trust, and the advocacy for individual clients is irreplaceable by AI)? Does it shift toward a government adoption endgame (proving that AI-assisted legal guidance works and advocating for Legal Aid to adopt the model)? Or does it pursue a new endgame entirely, becoming the body that evaluates and validates AI-generated legal guidance for accuracy and safety? Each answer is a different organisation with a different cost structure, a different funding model, and a different theory of change. The question is not which answer is right in the abstract. It is which answer is right for this organisation, with its specific capabilities and relationships. That is an endgame question.

Five New Endgames

The Gugelev and Stern framework was designed for a world where the primary constraint on social impact was organisational scale. In the AI era, that constraint is shifting. Scale is becoming cheaper. The new constraints are trust, legitimacy, domain expertise, democratic accountability, and the governance of powerful systems built by a small number of corporations. This suggests endgames that didn't need to exist in 2015.

7. Algorithmic accountability

A CSO develops the tools, expertise, and institutional credibility to hold AI systems accountable in a specific domain, including criminal justice, housing, employment, healthcare, and environmental compliance. It audits algorithms, brings strategic litigation, publishes findings, and builds public understanding of how automated systems affect people's lives. The core competency is technical and legal expertise combined with the independence and credibility that government bodies and commercial actors lack.

CASE STUDY: Amnesty International's Algorithmic Accountability Toolkit

In December 2025, Amnesty International published its Algorithmic Accountability Toolkit, a practical resource enabling civil society investigators, rights defenders, and activists to hold actors accountable for AI-facilitated harms. It covers audit methodologies, litigation strategies, and advocacy approaches. Amnesty is not pursuing this as a side project: it represents a strategic extension of its accountability mission into a domain where harms are increasingly algorithmically mediated. The toolkit is freely shared (open source endgame logic) but the core capability, the institutional credibility, the legal expertise, and the global network, is Amnesty's. This is a classic algorithmic accountability endgame: prove the methodology, share it widely, maintain the institutional authority to validate results.

8. Domain expertise capture

A CSO positions itself as the authoritative source of domain-specific knowledge that AI systems need but cannot generate: community-level data, practitioner wisdom, contextual judgment, lived experience. It becomes indispensable not by producing reports but by providing the ground truth against which AI outputs are validated. The CSO's value increases as AI improves, because better AI systems need better training data and better evaluation frameworks, and those require human expertise that cannot be automated.

CASE STUDY: NetHope: the 57-NGO AI consortium

NetHope, a consortium of 57 major international NGOs collectively reaching hundreds of millions of beneficiaries, has developed what is arguably the most comprehensive AI governance framework built specifically for the NGO sector. Its AI Lighthouse guide

covers ethical assessment, implementation, accountability, and data stewardship, not adapted from commercial frameworks but built from the ground up for organisations operating in fragile contexts, with vulnerable populations, in environments where commercial AI vendors have no presence and no interest. NetHope is not trying to scale its direct service. It is positioning itself as the domain expert against which AI tools for the humanitarian sector are evaluated. In an AI era where every humanitarian organisation will be evaluating AI systems, the organisation that sets the evaluation standards holds significant structural power.

9. Transition stewardship

A CSO exists specifically to manage the social transition caused by AI disruption in a specific sector or community, including retraining displaced workers, supporting affected communities, advocating for policy responses, and documenting what is happening. It is time-limited by design: when the transition is complete, or the acute phase has passed, the organisation winds down or pivots. This is analogous to disaster response organisations, but for economic and institutional disruption rather than natural catastrophe. Previous technology transitions happened slowly enough that existing institutions could absorb the adjustment. The evidence suggests AI may move faster.

10. Democratic AI governance

A CSO builds the infrastructure for meaningful public participation in AI governance decisions, not just consultation but genuine democratic input into how AI systems are designed, deployed, and held accountable. The model exists: Taiwan's vTaiwan platform used structured deliberation processes to produce government action on 80% of technology policy issues it examined. The endgame is that participatory AI governance becomes embedded in institutional practice. The CSO then either evolves into permanent democratic infrastructure or hands off to government. The concentration of AI capability in a small number of firms, combined with the documented lag of regulatory responses, creates a democratic deficit that existing governance structures are not equipped to fill.

11. Epistemic commons maintenance

A CSO maintains the shared knowledge infrastructure, the "epistemic commons", in a specific domain, ensuring that reliable, verified, publicly accessible information continues to exist even as AI-generated content floods the information environment. This is a public library function for the age of synthetic media: not producing knowledge, but guaranteeing its reliability.

CASE STUDY: Wikimedia Foundation: commons in the age of AI

Wikipedia is the world's largest repository of freely available verified knowledge, maintained by a community of volunteer editors applying consistent standards of sourcing and neutrality. The Wikimedia Foundation has spent the last two years wrestling with what AI means for this model. AI can generate plausible-sounding

encyclopaedic content at enormous scale and zero cost, which looks like a threat to an organisation built around human-curated, sourced, verifiable content. But the inverse is also true: in an information environment saturated with AI-generated text of uncertain provenance, the Wikimedia model, built on transparent sourcing, community verification, and open editing history, becomes more valuable, not less. The Foundation is not trying to out-produce AI. It is maintaining the epistemic commons: the shared infrastructure of verified knowledge that the information ecosystem depends on. That is an endgame.

The Revised Framework

Eleven endgames, then. Six original, reinterpreted. Five new. The table below maps them against the questions that matter most for a CSO navigating AI disruption.

Endgame	Core question	AI changes...	Budget shape
Open source (evolved)	Can we be the trusted authority, not just the producer?	Amplifies reach; demands higher credibility investment	Stable to growing
Replication	Is our model a human practice or an information process?	Codification easier; knowledge-work replication may become unnecessary	Peaks then declines (human) or collapses (knowledge)
Government adoption	Can we prove an AI-enhanced public service model?	CSO as AI proving ground; advocacy is harder	Grows during advocacy, then declines
Commercial adoption (evolved)	Are we fixing a market failure that AI created?	New market failures emerge; handoff conflicts of interest	Variable
Mission achievement	Can AI accelerate our defined goal?	Faster to achieve; harder to exit (data legacy)	Peaks then declines; exit costs higher
Sustained human service	Is our service inherently relational?	AI handles admin; humans handle relationships	Grows steadily
NEW ENDGAMES FOR THE AI ERA			
Algorithmic accountability	Who holds AI systems accountable in our domain?	Creates the need; CSO is the answer	Grows then stabilises
Domain expertise capture	Do we hold the ground truth AI systems need?	Symbiotic: CSO validates AI, AI amplifies CSO's value	Stable; potentially self-sustaining
Transition stewardship	Who manages the human cost of this transition?	AI is the cause; CSO manages the effect	Rises then declines
Democratic AI governance	Who ensures public voice in AI decisions?	Creates the democratic deficit; CSO fills it	Grows, then hands off or sustains
Epistemic commons	Who guarantees reliable information in our domain?	AI degrades the commons; CSO maintains it	Stable and ongoing

What this means for your organisation

Gugelev and Stern's practical advice was: define your endgame early. A nonprofit that knows its endgame makes better decisions about what to invest in, what to stop doing, and how to communicate its theory of change to funders. That advice holds, and is now more urgent.

I would add two questions to the ones they proposed. First: does your current endgame depend on the cost of knowledge work remaining roughly what it is today? If a significant part of your value proposition rests on the labour-intensive production of research, analysis, documentation, or communications, and if the answer to that question is "yes", then you need to interrogate the endgame. Not because AI will necessarily replace what you do, but because the implicit pricing of that work in the grant economy is likely to shift, and organisations that have thought through their response will be better placed than those that have not.

Second: is the thing you are protecting, the public good you exist to provide, one that AI actively degrades? Reliable information. Democratic accountability. Human-centred public services. Community voice in technology decisions. If so, one of the five new endgames may be more honest about what your organisation is actually for than the endgame you currently claim.

> *The question is not whether AI will affect your organisation. It will. The question is whether you choose your response or have it chosen for you.*

A note on my own context

I work at The Chancery Lane Project, which develops climate-aligned legal clauses and shares them freely with lawyers, businesses, and governments working on decarbonisation. In Gugelev and Stern's original framework, we are a classic open source endgame organisation: we produce knowledge, share it freely, and function as a knowledge hub.

The analysis in this paper confronts me directly. AI is already capable of drafting climate-aligned legal clauses. The technology is not perfect: hallucination rates on specific legal queries remain high, and domain expertise in climate law is exactly the kind of nuanced, contextual knowledge that current AI systems handle poorly. But the trajectory is clear. If TCLP's endgame is the production of clauses, that endgame is under pressure.

What I think is actually true, and what I am exploring in conversations with potential partners, is that TCLP's value may lie less in clause production and more in the combination of domain expertise capture and epistemic commons maintenance. We hold the ground truth of what a climate-aligned legal obligation actually means, tested against legal practice, scientific evidence, and real-world implementation. In an AI era where climate law will be drafted at scale by AI systems, the organisation that maintains the benchmark for what "correct" looks like may be more valuable than the one that produced the original clauses.

That is a different endgame. Naming it clearly would change what we invest in, how we communicate our value to funders, and what partnerships we prioritise. Which is exactly Gugelev and Stern's point.

Conclusion

Gugelev and Stern ended their paper with a Steve Jobs quotation: "Death is the destination we all share." Their point was that a nonprofit, like a person, should live purposefully toward an inevitable end. Scaling up is not a purpose. An endgame is.

Ten years on, the urgency of that argument has increased. The AI disruption is not a distant threat to civil society. It is already reshaping the economics of knowledge work, the expectations of funders, the pipeline of the sector's future workforce, and the information environment in which CSOs make their case. Organisations that define their endgame clearly, choosing one that is honest about what they offer that AI cannot, will be better placed to pursue it. Those that do not will find the endgame being chosen for them.

The question is the same as it was in 2015. The menu has changed.

Acknowledgements

This piece builds directly on Alice Gugelev and Andrew Stern's "What's Your Endgame?" (Stanford Social Innovation Review, Winter 2015). The framework they developed is theirs; what follows is my attempt to update it for a changed landscape. The evidence base draws on research conducted for a broader project on AI and civil society, details of which are available on request. I am grateful to colleagues at The Chancery Lane Project for reactions to earlier drafts of the ideas here.

Key sources cited

Gugelev, A. and Stern, A. (2015) 'What's your endgame?', Stanford Social Innovation Review, Winter 2015. Epoch AI (2024) inference cost analysis, November 2022 to October 2024. Stanford HAI (2025) AI Index Report 2025 [SWE-Bench benchmark data]. The 91.4% contract analysis accuracy figure cited in this paper requires source confirmation before final publication.

About the author

Ben Metz is Executive Director of The Chancery Lane Project, a legal charity that develops climate-aligned contract clauses. He writes about AI, civil society, and the future of law at benmetz.org.