Staying Ambitious: Harnessing AI for Social Impact

Introduction

Social impact organisations are operating in a period of overlapping crises, rising needs and constrained resources. Across sectors and geographies, organisations are being asked to do more, for more people, with less. The challenge is not only how to respond to this pressure, but how to stay ambitious within it.

Artificial intelligence offers one route. Used well, AI can help organisations unlock knowledge, improve decision-making, reduce administrative burden, extend frontline capacity and reach people with more personalised support. It can help teams analyse large volumes of data, summarise evidence, translate content, identify patterns, support learning, improve services and make better use of scarce resources.

At the same time, AI is not a shortcut to impact, nor a substitute for trust, relationships, good governance, strong institutions or community knowledge. If used poorly, AI can amplify bias, deepen exclusion, undermine accountability, spread misinformation and create new risks for the very people it is intended to support.

The social impact of AI can be understood through three connected areas: lives, livelihoods and learning. AI can have positive and negative effects on people's health and wellbeing; jobs, incomes and economic opportunity; and access to education, skills and knowledge. These areas remind us that AI is not only a technology issue. It is a social, economic and institutional issue that will shape who benefits, and who does not. They also provide a practical framework for business action: through core business operations, philanthropy and advocacy.

The future of AI for social impact will not be determined by technology alone. It will be shaped by the choices made by governments, businesses, civil society organisations and communities. It will depend on whether AI is deployed in ways that are human-centred, inclusive, transparent, accountable and grounded in real needs.

For Business Fights Poverty, the aim is a human-first, AI-enabled and social impact-focused future. This means using AI as a complement to human expertise, not a replacement for it. It means recognising the limits of AI, and ensuring that community knowledge and diverse perspectives shape the outputs and solutions that emerge.

This insight summary draws on several years of Business Fights Poverty's work on AI and social impact, including online and in-person convenings with business leaders, civil society organisations, academics, funders and practitioners across our global community. It builds on the 2023 framework paper, "Generative AI and Social Impact: The Role of Business", published by Zahid Torres-Rahman, Co-Founder and CEO of Business Fights Poverty, and Jane Nelson, Founding Director of the Harvard Kennedy School Corporate Responsibility Initiative. It also draws on discussions from the University of Cambridge Centre for Human-Inspired Artificial Intelligence and Business Fights Poverty Institute event in May 2026, "From Promise to Practice: How Partnerships Can Deliver AI for Global Development", alongside earlier Business Fights Poverty convenings on generative AI, equity and the future of work.

The Five Key Insights

1. AI's promise depends on the system around it

AI's impact depends less on the model than on the system it enters. A technically impressive tool can fail if the surrounding conditions are weak. A relatively simple AI application can be powerful if it is embedded in the right trusted relationships, workflows, institutions and infrastructure.

This means paying attention to the "infrastructure layer" of AI for social impact. This includes electricity, connectivity, access to devices, data governance, institutional ownership, compute capacity, digital literacy, user trust, policy, regulation, financing, operational workflows and long-term maintenance. These conditions are not separate from AI adoption. They are part of what determines whether AI can be useful, safe and sustainable.

The example of agricultural advisory services makes this practical. On the surface, the idea is simple: smallholder farmers receive timely, locally relevant advice on what to plant, when to plant, how to manage pests, and how to respond to weather variability. The technology is increasingly feasible. But the hard questions sit around the model. Which farmer registry does it draw on? How is that registry connected to soil, weather and market data? In which language does the farmer receive advice? Who owns the relationship with the farmer? Who pays for compute as usage

grows? Who is accountable if the advice is wrong? Who updates the system when local conditions change?

None of these questions is purely technical, yet each determines whether the tool works in practice. This matters because much of the current AI conversation is still driven by excitement about tools, pilots and use cases. These are important, but they are not enough. Social impact organisations operate in complex systems. A health AI tool depends on clinical validation, referral pathways, trained workers, patient trust and regulatory oversight. A financial inclusion tool depends on consumer protection, data rights, fraud prevention and responsible lending. An education tool depends on teachers, curricula, language, access to devices, and the confidence of students and parents.

AI must therefore be designed around the problem, the domain and the institutional conditions that can absorb it. A use case is not promising simply because the technology exists. It is promising when it addresses a clear gap, fits a real workflow, can be maintained over time, and improves outcomes for people and communities.

This calls for a purpose-driven approach to AI. Staying ambitious does not mean adopting AI wherever it is available. It means starting with the social problem, defining the intended benefit, understanding who could be harmed or excluded, and being willing to adapt, pause or stop when the evidence shows that AI is not the right tool.

AI's opportunities and risks will be shaped by deep-seated inequities such as income, gender, race and geography. A tool that works for well-connected, English-speaking, digitally confident users may fail for those with limited connectivity, low literacy, low trust in institutions, or limited time and money to experiment with new technologies. The system around AI must therefore be judged by whether it expands or narrows access for those who are already underserved.

For businesses and funders, this points to a shift in investment. Funding pilots is not enough. The sector also needs investment in data readiness, organisational capacity, local research, digital infrastructure, governance systems, user training and long-term adoption. These elements may be less visible than new products, but they are what determine whether AI becomes useful at scale.

For civil society organisations, the implication is equally clear. AI strategies should start with organisational purpose and social outcomes, not with available tools. Where are the biggest constraints to impact? What knowledge is locked away in reports, datasets, field notes or community feedback? Which tasks consume time without adding value? Which decisions could be improved with better evidence? Which relationships must remain human? These questions should guide whether, where and how AI is used. AI's promise is real, but it is conditional. It depends on the strength of the system around it.

2. Responsible scale requires partnerships that clarify ownership, risk and accountability

Partnership is not a warm aspiration. It is the operating model for responsible AI in social impact. No single actor can provide everything needed to deliver AI responsibly at scale. Technology companies may bring engineering expertise, infrastructure and product development capacity. Governments bring legitimacy, public systems, policy authority and responsibility for public services. Civil society organisations bring trust, community relationships, practical knowledge and delivery experience. Academia brings evidence, evaluation and critical scrutiny. Funders and multilateral institutions bring convening power, patient capital and cross-country learning. Frontline workers bring knowledge of real workflows and practical constraints. Communities bring lived experience, local legitimacy and the right to shape what is built for them.

The challenge is not simply to bring these actors together. It is to clarify what each is responsible for. Effective partnerships need to answer difficult questions early. Who defines the problem? Who pays for customisation? Who owns the data? Who maintains the system? Who monitors performance after deployment? Who carries the risk if the tool causes harm? Who is accountable when outcomes are not as expected? Who decides whether to scale, adapt or stop? How will communities benefit from the value their data or knowledge helps create?

Without this clarity, partnerships can become fragmented, duplicative or extractive. Too many disconnected tools, overlapping pilots and unshared datasets waste resources and weaken learning. In a resource-constrained environment, duplication is not just inefficient. It diverts scarce resources from urgent needs.

Responsible scale also requires a mature view of the private sector. Businesses should not be treated only as vendors selling tools into development contexts. Equally, businesses should not treat AI for social impact simply as a market opportunity. The strongest partnerships are problem-led, not product-led. They require co-design, adaptation, transparency, local capacity building and a willingness to share responsibility for outcomes.

Businesses can act through three connected levers: core business, philanthropy and policy engagement. Through core business, companies can shape how AI is designed, deployed and governed in their operations, value chains, products and services. This includes conducting social impact and human rights assessments, addressing risks to workers and communities, investing in diverse datasets, reducing bias, supporting employees, and sharing productivity gains more fairly.

Through philanthropy and community investment, companies can support civil society organisations to use AI responsibly, fund independent audits and research, strengthen local datasets, support training and retraining, and back projects that improve health, livelihoods and learning for vulnerable communities.

Through policy engagement, companies can support safe, ethical and inclusive regulatory frameworks, advocate for digital inclusion, contribute to public debate, and help build the skills and institutions needed for an AI-enabled future.

This three-part role matters because AI's social impact will not be shaped only by the companies that build AI systems. It will also be shaped by companies across all sectors that adopt AI into their workflows, value chains, customer services, recruitment systems, risk assessments and community programmes.

Government ownership is also essential, particularly where AI is being used in public systems such as health, education, agriculture, financial services or social protection. But government ownership does not mean government-only delivery. Public systems often depend on partnerships with civil society, business, academia and communities. The key is to ensure that AI strengthens public accountability rather than bypassing it.

Partnerships also need to build trust by being clear about data sources, oversight and redress. Trust is not only a question of whether people like or accept AI. It depends on whether they can understand where the data comes from, how outputs are generated, who has checked them, and how decisions can be challenged. For high-stakes uses, especially in health, finance, employment or public services, people need routes for contestability and redress.

The human side of adoption also matters. Many workers understandably fear that AI is being introduced to replace them, expose underperformance or reduce headcount. Adoption is more likely to succeed when AI is framed and implemented as a capability-building tool, not simply a cost-cutting tool. This requires honest communication, worker involvement, training and clear boundaries around how AI will be used.

Partnerships must include the people closest to the work. In healthcare, clinicians need to understand and trust AI tools before they use them. In poverty reduction, social workers need to know how AI can help them prioritise time without weakening relationships. In education, teachers need to see AI as a support, not a threat. In agriculture, farmers need advice that is understandable, locally relevant and delivered through channels they trust.

The most effective partnerships are therefore not only multi-stakeholder. They are system-level. They are designed to change how knowledge,

resources, decisions and accountability flow through a system. They are specific enough to solve real problems, but broad enough to address the conditions that make solutions sustainable.

For leaders, this means moving beyond general declarations about collaboration. AI partnerships should begin with a shared problem statement, a governance model, a data framework, a financing plan, an evaluation approach and an accountability mechanism. These may sound procedural, but they are what turn promise into practice.

3. The most useful AI may be small, specific and human-augmenting

Bigger AI is not always better. In many social impact contexts, the most useful AI may not be the largest, most advanced or most general model. It may be a smaller, task-specific tool that works reliably within real-world constraints.

This is especially important in low-resource environments. Many communities face unreliable connectivity, expensive data, limited access to smartphones, unstable electricity and constrained compute capacity. AI tools designed for high-bandwidth, cloud-dependent environments may fail to reach the people they are supposed to serve. In these contexts, “small AI” can be more relevant: models that are lightweight, robust, locally deployable, able to run offline or at the edge, and designed for specific tasks.

For a frontline health worker, the best AI tool may not need to perform well on global benchmarks. It needs to help triage urgent cases, support diagnosis, flag risks, translate information or reduce paperwork within an existing workflow. For an agricultural extension worker, it needs to provide useful advice in the right language, based on local crops, soil, weather and market realities. For a social worker, it needs to help identify which families need human support most urgently, without undermining trust or agency.

Across poverty reduction, healthcare, financial inclusion, education and humanitarian response, the strongest use cases are about augmenting people, not replacing them.

In poverty reduction, technology can help families understand their own multidimensional poverty, identify priorities and build plans for change. AI can help social workers reach more families and tailor support more effectively. But the goal is not to replace the social worker. It is to save time on tasks where AI can help, so that human workers can focus on the harder aspects of poverty that require trust, judgement, empathy and persistence.

In healthcare, AI can extend scarce specialist capacity by helping to triage urgent conditions, support screening or accelerate diagnosis. This can be especially valuable in settings where the disease burden is high and specialist availability is limited. But AI is not a magic wand. Tools must be validated in the specific populations where they are deployed, monitored over time, and integrated into clinical workflows. Patient demographics, equipment, protocols and local conditions can all affect performance.

AI can also unlock organisational capacity behind the scenes. Practical uses include drafting reports, cleaning survey data, summarising evidence, translating documents, mapping systems, managing communications, supporting proposals, improving knowledge management and matching organisations to donor opportunities. These uses may not sound as exciting as frontline applications, but they can free up scarce human capacity for higher-value work.

This is particularly important because many organisations already hold large amounts of underused knowledge. Programme reports, evaluations, community feedback, meeting notes, case studies, survey data and lived experience often sit across fragmented systems. AI can help organisations make better use of this organisational memory. It can identify patterns, surface lessons, connect people to relevant knowledge and reduce the time spent searching for information.

Youth Business International's work on AI and young entrepreneurs highlights how AI can be a practical leveller for young founders and small firms. Generative AI, low-code platforms and other digital tools can help entrepreneurs design products, automate services, reach new markets and make better decisions without large teams or extensive capital. But this potential will only be realised if young entrepreneurs have affordable access to tools, reliable connectivity, relevant-language support, AI literacy, mentorship and practical business training. Otherwise, AI could widen the gap between those with the resources and confidence to adopt it, and those who remain excluded from the next wave of digital opportunity.

Productivity gains can also be unevenly distributed. AI may complement some jobs and displace others. It may create new opportunities for some workers while increasing insecurity for others. It may enable lower-skilled workers to improve their performance, but only if they have access to the tools, training and confidence to use them.

A particular risk is that AI removes the first rungs of opportunity. If entry-level tasks are automated without redesigning pathways into work, young people and under-represented groups may find it harder to build experience. Businesses and training providers therefore need to co-design new pathways, such as apprenticeships, fellowships and earn-and-learn models, so that AI augments entry-level opportunity rather than narrowing it.

AI is also changing what it means to manage work. The future will require people who can lead blended teams of human colleagues and AI tools, knowing when to delegate to AI, when to question it, when to override it and when not to use it at all. This makes critical judgement, curiosity, data literacy and responsible AI practice core leadership skills, not specialist technical skills.

This is why the word “augmenting” matters. A human-centred approach asks how AI can strengthen human capability, dignity and agency. It does not assume that every efficiency gain is automatically positive. It asks who benefits from the time saved, who carries the risk of automation, and whether workers and communities share in the gains.

Productivity is not the same as impact. Faster reporting does not automatically improve outcomes. More efficient grant writing does not necessarily lead to better programmes. Women entrepreneurs may save time using AI without necessarily growing their businesses. NGOs may become more efficient without improving people’s lives.

The test is whether AI helps organisations make better decisions, strengthen relationships, reach more people, improve services, reduce harm and deliver more meaningful results. Efficiency matters, but it should be in service of impact.

For leaders, the practical implication is to look for high-friction, high-value tasks where AI can extend human capability. Start with clear tasks, clear users and clear safeguards. Avoid designing humans out of systems that depend on trust, judgement and care. The goal is not AI-first. The goal is human-first and AI-enabled.

4. Impact depends on integration, outcomes and safeguards, not adoption alone

AI success should not be measured through usage, satisfaction or technical performance alone. A tool may be popular, efficient or impressive without improving social outcomes. The real test is whether AI improves health, livelihoods, education, safety, resilience, decision quality, access to services or community agency.

The lives, livelihoods and learning framework is useful here. For lives, we should ask whether AI improves health and wellbeing, strengthens access to care, reduces risks, and protects people from harm. For livelihoods, we should ask whether AI improves incomes, job quality, financial resilience, productivity and access to opportunity, while managing the risks of job displacement and exclusion. For learning, we should ask whether AI improves access to education, skills, knowledge and confidence,

particularly for people who have historically been excluded from good-quality learning opportunities.

Evaluation must be built in from the beginning. Before AI is introduced, organisations need a clear understanding of the baseline: current workflows, costs, waiting times, quality, accuracy, user experience, outcomes and risks. After deployment, evaluation should examine whether the AI tool improves the system in which it is used. Does it reduce time to diagnosis? Does it improve patient outcomes? Does it increase incomes? Does it strengthen learning? Does it improve the quality of decisions? Does it reduce administrative burden without weakening accountability? Does it increase access for those who were previously excluded?

Research from the Cherie Blair Foundation for Women on women entrepreneurs in low- and middle-income countries underlines the importance of distinguishing between AI adoption and AI integration. Many women entrepreneurs are already using AI, but often for lower-risk, customer-facing tasks where experimentation is easier and the benefits are immediate. The deeper opportunity lies in embedding AI into routine business workflows, especially finance, operations, planning, stock management and decision-making. Without this shift, AI may save time and reduce pressure without necessarily changing business trajectories or supporting sustained growth.

This is a useful lesson for the wider social impact sector. The question is not only whether organisations have started using AI. It is whether AI is embedded in ways that improve decisions, strengthen delivery, reduce risks and create measurable value. Light experimentation may be useful, but deeper integration requires skills, trust, governance and workflow redesign.

Evaluation should also be tiered across different levels. At the individual level, AI might improve access, experience, income, health, learning or productivity. At the community level, it might strengthen voice, trust, inclusion or resilience. At the system level, it might improve institutions, workflows, governance, service delivery or accountability. A tool can perform well at one level while creating problems at another. For example, it may make a provider more efficient while reducing user trust, or improve data collection while weakening community control over information.

The full cost must also be counted. AI can appear cheap if only software costs are considered. But responsible deployment requires compute, data governance, training, maintenance, monitoring, infrastructure, cyber security, organisational change, environmental management and long-term affordability. In low-resource settings, these costs matter. A solution that depends on unaffordable cloud infrastructure, external expertise or donor funding that soon ends is not sustainable.

Environmental costs also need greater attention. AI depends on energy, water, compute and digital infrastructure. At the same time, AI systems intended to serve vulnerable communities must be resilient to climate shocks. Floods, storms, heat, conflict or power failures can disrupt the infrastructure on which digital tools depend. AI for development should therefore be designed for low energy use, offline functionality, local resilience and emergency conditions.

The risk side of the equation is equally important. AI can amplify bias, misinformation, fraud, unsafe decision-making and exclusion. In financial services, AI can expand access to credit, but it can also enable scams, discriminatory lending or opaque decisions. In healthcare, AI can support diagnosis, but false positives, false negatives or model drift can cause harm. In humanitarian settings, AI can help identify needs, but it can also raise surveillance and data protection risks. In recruitment, AI can speed up screening, but it can reproduce structural bias.

Misinformation and hallucination are particular risks. As AI-generated text, audio, images and video become easier to produce, trust in evidence-based knowledge can be weakened. For social impact organisations, whose work depends on credibility and trust, this is a serious concern. AI-generated outputs need expert verification, clear sourcing, transparency and careful communication. Speed should not come at the expense of truth.

Post-deployment monitoring is essential because AI systems can drift, degrade or behave differently over time. A model trained on one population may perform differently in another. A tool that worked during a pilot may become less accurate as equipment, behaviours, demographics or workflows change. Responsible AI is therefore not a one-off approval decision. It requires ongoing evaluation, traceability and accountability.

Organisations also need to decide where not to use AI. In some cases, the priority may be more teachers, better clinics, stronger governance, safer roads, reliable electricity, affordable connectivity or direct cash support. AI should be considered one tool within a wider system, not the default answer to every development challenge.

For organisations, this requires explicit AI policies, board-level engagement, clear rules on what data can and cannot be uploaded, and processes for assessing risk before tools are adopted. It also requires critical AI literacy, not just technical training. Staff and partners need to understand not only how to use AI, but how to question it: What data was it trained on? Whose knowledge is missing? What assumptions does it make? What are the risks? Who benefits? Who may be harmed? When should a human override the output?

Critical AI literacy requires both curiosity and caution: curiosity to explore how AI can improve work and expand opportunity, and caution to question its outputs, assumptions, data sources and unintended

consequences. Impact measurement must therefore move beyond adoption. The question is not simply: are people using AI? The question is: is AI helping people live safer, healthier, more dignified and resilient lives?

5. Inclusion requires community ownership, critical literacy and the right to shape AI

Inclusion is not simply about “adding more data”. It requires asking deeper questions: who defines the problem, whose knowledge counts, who owns the data, who shapes the design, who benefits from the value created, and who has the right to say no?

AI systems are often built from data that reflects wealthy, connected, English-speaking and institutionally visible populations. This creates the risk that the knowledge, languages and lived realities of marginalised communities are excluded or misrepresented. Women, rural communities, low-income households, oral cultures, offline populations and speakers of under-represented languages may be particularly affected.

Language is one of the clearest inclusion challenges. Many AI systems work best in English and a relatively small number of high-resource languages. But language is not only a technical issue. It is a matter of knowledge equity. If people cannot ask questions, access services or express their needs in their own language, they are excluded from shaping the digital future.

There is also a risk of extractive AI. Communities may contribute language, cultural or local knowledge data that external organisations later use to create commercial products. If those communities then have to pay to access tools built partly from their own knowledge, AI becomes extractive. Responsible AI therefore requires fair data governance, community benefit, transparency, consent, local stewardship and clarity over commercial use.

Community ownership matters because communities often understand local constraints better than external actors. They know which channels people trust, which languages and dialects matter, which risks are most sensitive, which forms of support are acceptable, and which solutions already exist locally. AI initiatives should therefore begin with community-defined problems and community participation in design, testing and governance.

This also means respecting non-digital approaches. Digital inclusion is not achieved simply by assuming that everyone has a smartphone, uses WhatsApp or can navigate an app. In some contexts, printed materials, in-person support, voice interfaces, community intermediaries or basic

phones may be more appropriate. Inclusive AI must meet people where they are, not where designers assume them to be.

Digital exclusion remains a major barrier to AI-enabled opportunity. If people lack internet access, devices, digital skills, language access or confidence, they may be excluded from the benefits of AI even as AI becomes more central to healthcare, work, education and public services. This makes digital inclusion a social impact priority, not just a technology priority.

Inclusion is not only about access to tools or training. For many women entrepreneurs and frontline workers, the binding constraint may be time, confidence, cognitive load or trust. Training needs to be practical, flexible and task-based, fitting the realities of people managing businesses, households, care responsibilities and limited resources. Otherwise, AI support may be available in theory but unusable in practice.

This point is especially important for women entrepreneurs. The Cherie Blair Foundation for Women's research shows that many women entrepreneurs are interested in AI and are already experimenting with it, but the gains are often concentrated in time savings rather than business growth. To move from light use to meaningful impact, support needs to help women integrate AI into higher-value business functions such as finance, operations, planning and stock management. It also needs to recognise time poverty, care responsibilities, confidence, trust and cognitive load as real barriers to deeper use. Training and support should therefore be practical, flexible, task-based and designed around the realities of women running businesses while managing wider responsibilities.

Critical AI literacy is central to inclusion. Communities, frontline workers and organisations need the ability not only to use AI, but to interrogate it. This includes understanding bias, data rights, environmental impact, misinformation, privacy, consent, power and the limits of AI outputs. Literacy should build agency, not dependency. People should be able to decide whether and how AI is used in their lives.

Business Fights Poverty's own AI principles are especially relevant here. A human-first approach uses AI as a complement to human expertise. Inclusivity and diversity require active work to eliminate structural biases and represent diverse views. Rigour and quality mean using AI to enhance work rather than compromise it. Trust and credibility require transparency and expert validation. Conscious and mindful usage means recognising that AI is not always the right tool. Accessible and engaging content means using AI to reach people in new languages, formats and mediums. Community-driven innovation means recognising and integrating the valuable insights generated by communities themselves.

For businesses, this has practical implications. Inclusive AI should involve investing in diverse datasets, low-resource languages, accessible design,

local research capacity and participatory governance. It should also involve listening to communities as co-creators, not treating them merely as users, beneficiaries or sources of data.

For funders, inclusion means supporting the slower, relational work that responsible AI requires: community engagement, translation, local capacity building, ethical review, participatory design and long-term stewardship. These costs should not be treated as extras. They are core to responsible implementation.

For civil society organisations, inclusion means protecting the trust they hold with communities. AI should not be introduced in ways that undermine that trust. Sensitive beneficiary data must be handled with care. Tools should be tested for bias and harm. Communities should understand how their data is used. Human judgement should remain central in decisions that affect people's lives.

The goal is not simply to extend AI to more people. It is to ensure that people and communities have the power to shape AI in ways that reflect their own priorities, knowledge and aspirations. AI should be built with communities, not merely deployed to them.

Conclusion

AI offers a genuine opportunity for social impact organisations to stay ambitious in a time of crisis and constraint. It can help unlock organisational memory, reduce administrative burden, strengthen decision-making, extend frontline capacity, improve access to information and support more personalised services. In a world of rising needs and shrinking resources, these opportunities matter.

But AI will not deliver social impact by itself. Its value depends on the systems, partnerships and choices around it. The first choice is to start with the problem, not the tool. AI should be applied where it can address a clear need, fit existing workflows, strengthen human capability and improve measurable outcomes.

The second choice is to invest in the foundations. Connectivity, compute, data governance, digital literacy, local capacity, institutional ownership, trust and regulation are not peripheral. They determine whether AI can work in practice.

The third choice is to build partnerships with accountability. Responsible AI for social impact requires governments, businesses, civil society, academia, funders, frontline workers and communities to work together with clear roles, shared risks and transparent governance.

The fourth choice is to keep humans at the centre. AI should augment judgement, empathy, relationships and accountability, not replace them. The most powerful uses may be those that free people to do the work that only people can do.

The fifth choice is to ensure that inclusion means agency. Communities should not only be represented in datasets. They should help define problems, shape design, govern data, share in benefits and decide whether and how AI is used.

For businesses, this means using core business, philanthropy and policy engagement to ensure that AI strengthens lives, livelihoods and learning, while managing risks to rights, safety and trust. For civil society, it means adopting AI carefully and confidently, with clear policies, critical literacy and a focus on outcomes. For funders, it means supporting not only pilots, but the infrastructure, capacity and evidence needed for responsible scale. For governments, it means creating enabling environments that protect people while encouraging innovation that serves public priorities.

The future of AI for social impact is not predetermined. It will be shaped by the choices we make now. A human-centred, AI-enabled and social impact-focused future is possible, but only if ambition is matched by responsibility, inclusion and partnership.

Acknowledgements

This note was prepared by Zahid Torres-Rahman, Co-Founder and CEO of Business Fights Poverty, with the generous input of many people across our community.

Resources

Business Fights Poverty (2023): Generative AI and Social Impact: The Role of Business. Available here.

Business Fights Poverty (2023): Collaborating on Generative AI For Social Impact. Available here.

Business Fights Poverty (2024). How will AI shape equity across societies? Available here.

Business Fights Poverty (2025): How Can We Prepare the Workforce for an AI-Driven Future? Available here.

Business Fights Poverty (2026): Ambition at a Time of Crisis and Constraint: Harnessing Partnerships and AI for Global Development. Available here.

Cherie Blair Foundation for Women (2026): Adopted not embedded: AI, productivity and uneven gains for women entrepreneurs. Available here.

Youth Business International (2025): Policy Paper - Harnessing AI and Digital Solutions to Empower Young Entrepreneurs. Available here.