TECHTONIC JUSTICE

# Inescapable AI

*The Ways AI Decides How Low-Income People Work, Live, Learn, and Survive*

## EXECUTIVE SUMMARY

**Kevin De Liban**
Founder and President

**November 2024**

TECHTONIC JUSTICE

TechTonic Justice is a multidimensional, community-based effort to strengthen local justice movements and build long-term power among the people and communities that AI is leaving behind.

TechTonic Justice envisions a world rooted in justice for the people AI is leaving behind. We all have the power to decide whether AI is used as a tool to promote – rather than undermine – our civil rights and liberties, access to opportunity, and our fundamental human needs.

When we invest at the ground-level and mobilize those working directly within impacted communities, we can leverage our collective power to resist automated outcomes and create a future where all communities thrive.

**Kevin De Liban**
Founder and President

Kevin De Liban is the Founder of TechTonic Justice, a newly launched nonprofit to fight alongside low-income people harmed by AI. Through multidimensional advocacy, TechTonic Justice supports marginalized communities and their advocates to secure the work, housing, schooling, public benefits, and family stability needed for a thriving life. Kevin previously worked for 12 years at Legal Aid of Arkansas—most recently as its Director of Advocacy--where he represented over 1,800 low-income people in matters involving health care, workers' rights, public benefits, special education, and domestic violence. There, Kevin led campaigns at the cutting edge of anti-poverty advocacy: ending the state's use of algorithms that cut the in-home care of disabled people, stopping Medicaid work reporting requirements that stripped health insurance from 18,000 people, and overcoming qualified immunity to hold state officials personally liable for violating constitutional rights. Kevin regularly presents about imposing accountability on AI and consults with advocates, policymakers, and journalists in the U.S. and abroad.

# Contents

# Executive Summary

The use of artificial intelligence, or AI, by governments, landlords, employers, and other powerful private interests restricts the opportunities of low-income people in every basic aspect of life: at home, at work, in school, at government offices, and within families. While popular discourse has recently centered on the newest versions of AI that generate answers, reports, or images in response to users' questions or prompts, such technologies derive from a lineage of automation and algorithms that have been in use for decades with established patterns of harm to low-income communities. As such, now is a critical moment to take stock and correct course before AI of any level of technical sophistication becomes entrenched as a legitimate way to make key decisions about the people society marginalizes.

Employing a broad definition of AI, this report represents the first known effort to comprehensively explain and quantify the reach of AI-based decision-making among low-income people in the United States. It establishes that essentially all 92 million low-income people in the U.S. states—everyone whose income is less than 200 percent of the federal poverty line—have some basic aspect of their lives decided by AI.

While these decisions may not be universally negative, the use of such technologies over the past two decades demonstrates the capacity for broad, systemic harms with immense suffering at scales and speeds that were impossible with the human-centered methods that precede them. AI is often inaccurate, absurd, biased, and inscrutable. Even when it is not plagued by these problems, it is often carrying out a purpose that is unfavorable to low-income people, such as cutting benefits or making housing, education, or employment opportunities harder to access.

Where humans make decisions without AI, even wrong or biased ones, their scope is limited to their personal reach. Only elected officials, corporate executives, or high-ranking government staff might make decisions for masses of people, and they generally must do it through formal government or corporate policy processes that involve at least minimal levels of public or stakeholder scrutiny.

AI, though, enhances the ability of both high-ranking and lower-level decision-makers to act at scale, often through informal or surreptitious means. They can cut the in-home care of 4,000 disabled people in Arkansas despite them having medical conditions that have not gotten better, falsely accuse 40,000 people in Michigan of Unemployment Insurance fraud, or subject 4,000,000 people in Texas to the potential loss of health insurance through a labyrinthine Medicaid enrollment system that even agency staff cannot navigate—all real examples in which the bulk of harm happened within months of the AI system going live. And that is just in the area of public benefits. AI-based decision-making also happens on grand scales with high stakes in other areas, such as housing, employment, K–12 education, domestic violence, and child welfare.[1]

1 AI also makes decisions about low-income people's lives with respect to the criminal legal system, immigration, credit, tax enforcement, child support enforcement, health care treatment (independent of insurance), and voting, all of which fall outside the scope of this report. Additionally, AI harms low-income people independently of the decisions it makes about them. The technology demands huge reserves of energy, raw materials, and processing power, creating environmental harms that tend to fall hardest on low-income people. And AI requires labor from low-paid and otherwise-exploited workers to gather or refine raw materials, label data used to train the systems, and operate certain aspects of the systems, such as managing violent or hateful outputs. These nondecisional harms are critical to understand but outside this report's scope.

Of course, the root problem is not the technology, but rather the social forces that make life difficult for low-income people and the actors empowered to perpetuate such forces. Hostility toward low-income people, racism, sexism, ableism, homophobia, and expanding corporate power all shape a world in which it makes sense to use technology to impose austere and precarious living conditions on disfavored groups of people. From this warped vantage point, low-income people "merit" ongoing suspicion and must be prevented from "cheating" the system, "slacking" at work, or "tricking" their landlords—all beliefs that have long been embedded in laws and are now enforced through AI. Moreover, the decades-long hollowing out of government and unyielding desire to "streamline" both government and industry provide a neat justification for the use of technology like AI that promises (though may not deliver) huge gains in administrative efficiency.

While AI is not the root problem of societal inequities, it is a particularly dangerous manifestation of them. Acting at scales previously unfeasible, this new enabler of harms challenges existing political, market, and legal mechanisms for accountability. Its technological complexities, the lack of a constituency with sufficient political power at hand to oppose it, the unavailability of affordable attorneys, insufficient community organizing resources, limitations of existing laws, and general precarity of poverty frustrate the ability of low-income people and their allies to fight back against harmful AI or the forces enabling it. Even where there have been successful challenges, the resulting changes offer incomplete relief for people's suffering and do not alter the perverse incentive structures that cause, or allow for, AI to be used harmfully in the first place.

Ultimately, empowered constituencies of low-income people and their allies will need to win and enforce additional societal protections to meaningfully stave off AI-enabled harms. TechTonic Justice exists to support this fight, strengthening local justice movements by training, learning from, litigating alongside, organizing with, offering strategic assistance to, and building infrastructure that supports low-income people and their allies to resist immediate harms and win long-term gains.

The way ahead is as daunting as it is necessary. Only by reimagining the status quo and reconfiguring the forces that sustain it will harmful AI and the poverty it reinforces become escapable. This—not overhyped technologies—will involve truly revolutionary innovation.

**Figure 1. Low-income people in the United States and the ways in which they are exposed to AI decision-making.**

The 92 million people in the US who presently live below 200 percent of the federal poverty line are considered low income. Additionally, people flow into and out of low-income status based on changes in the economy and their personal circumstances. Such churn means that the number of people ever exposed to AI decision-making while they were low income is much higher than 92 million.

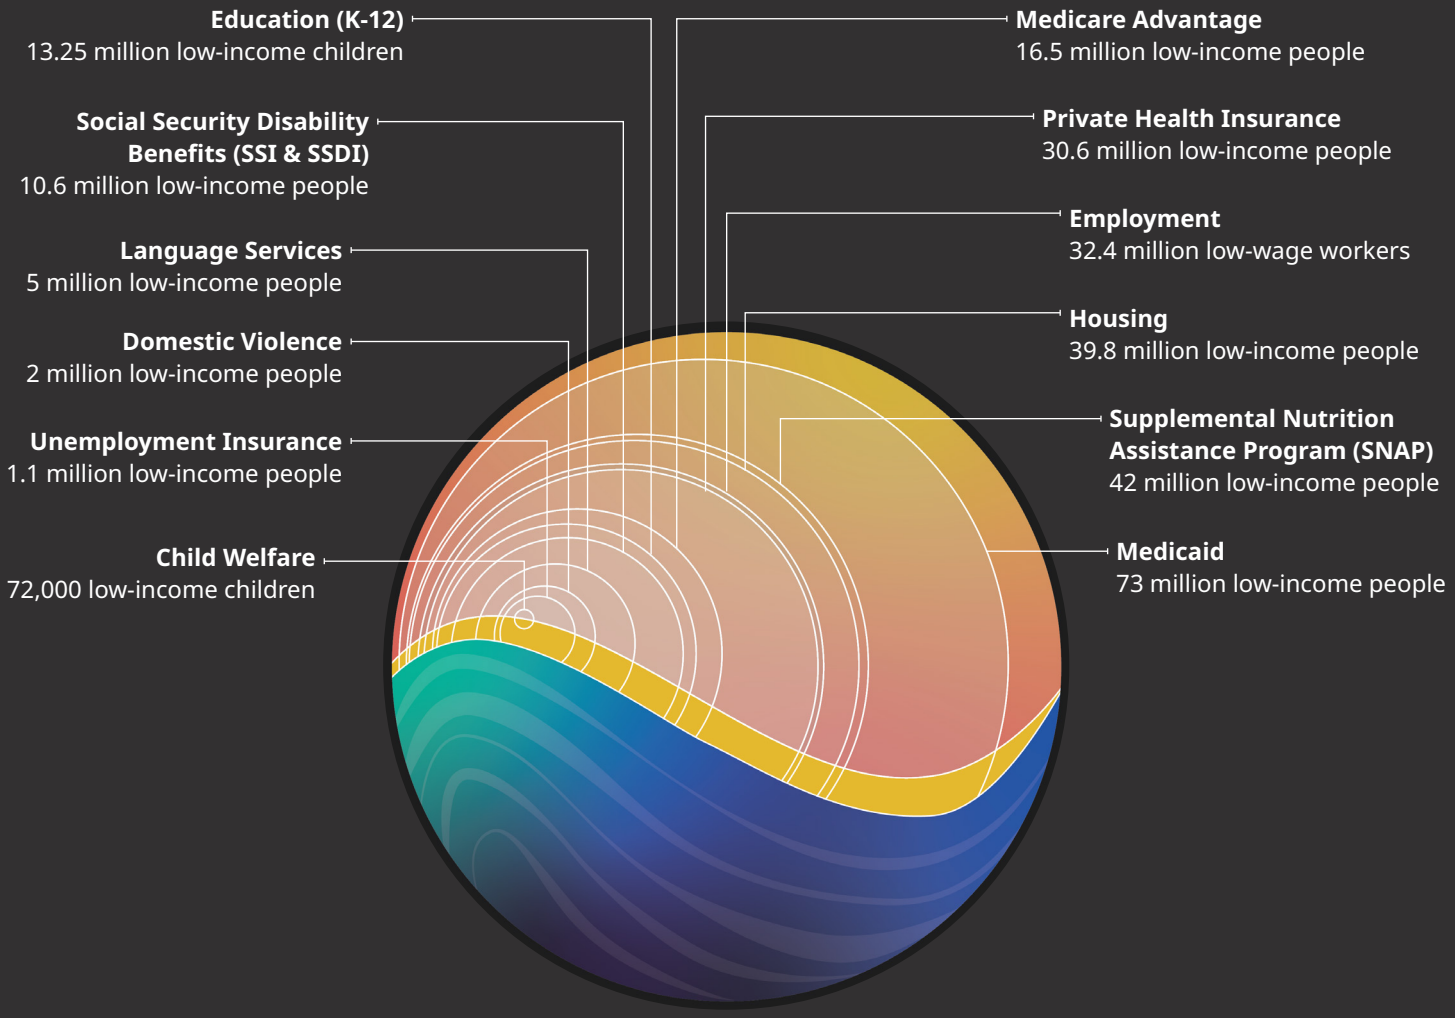


In this chart, the various types of AI decision-making that low-income people are exposed to should be read as concentric arcs, with the number of low-income people exposed to each type of AI decision-making visualized proportionally. The area under the arcs represents the proportion of low-income people exposed to AI decision-making in each of the specified categories. The area of interest is the full area of each arc.

# Key Findings

Part one of this report quantifies the number of low-income people in the United States and the ways in which they are exposed to AI decision-making.[2]

**Medicaid: 73 million low-income people** are exposed to AI-related decision-making in Medicaid through the eligibility and enrollment process, the determination of home- and community-based services, or, where a state uses private companies to manage the Medicaid program, the prior authorization process for medically necessary services. As a result, people are denied health insurance, home-based care needed to avoid nursing facilities, and medically necessary treatments and medicines.

**Medicare Advantage: About 16.5 million low-income people** are exposed to AI-related decision-making through the prior authorization processes used in Medicare Advantage programs. As a result, people are denied medically necessary treatments and medicines.

**Private Health Insurance (through employers or federal subsidies): About 30.6 million low-income people** are exposed to AI-related decision-making through the prior authorization processes in private health insurance. As a result, people are denied medically necessary treatments and medicines.

**Supplemental Nutrition Assistance Program, or SNAP: 42 million low-income people** are exposed to AI-related decision-making in SNAP through the eligibility and enrollment process or detection of alleged fraud. As a result, people are denied vital benefits needed to buy food, wrongly disqualified from the program, and falsely accused of wrongdoing.

2 Because low-income people are exposed in multiple ways, the following numbers are not cumulative. Even within a single category, like health insurance, the totals will exceed the number of low-income people because people can sometimes be eligible for more than one program or because available data sources cannot easily be reconciled. In addition, these are point-in-time numbers. Because people's incomes vary and they cycle into and out of exposure to AI in the ways described below, the actual number of people who ever have been or could soon be exposed is much higher.

**Medicaid**

73 million low-income people are exposed to AI-related decision-making in Medicaid through the eligibility and enrollment process, the determination of home- and community-based services, or the prior authorization process.

**Medicare Advantage**

16.5 million low-income people are exposed to AI-related decision-making through the prior authorization processes used in Medicare Advantage programs.

**Private Health Insurance**

30.6 million low-income people are exposed to AI-related decision-making through the prior authorization processes in private health insurance.

**Supplemental Nutrition Assistance Program (SNAP)**

42 million low-income people are exposed to AI-related decision-making in SNAP through the eligibility and enrollment process or detection of alleged fraud.

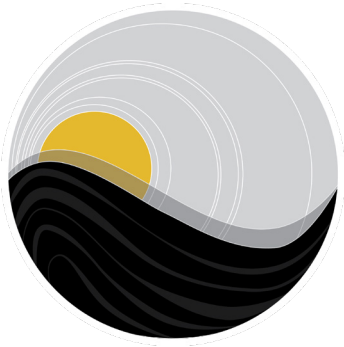

**Social Security Disability Benefits (SSI & SSDI)**

10.6 million low-income people, are exposed to AI-related decision-making in the Social Security Administration's disability benefits programs through current or planned uses of AI technologies in various parts of benefit administration.

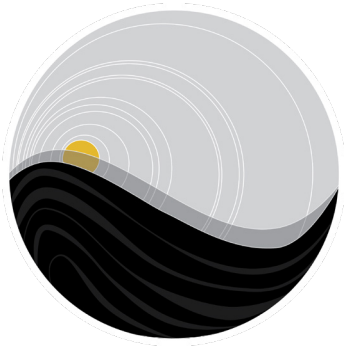

**Unemployment Insurance**

1.1 million low-income people are exposed to AI-related decision-making in the Unemployment Insurance program through the eligibility and enrollment process, identity verification practices, and the detection of alleged fraud.

**Housing**

39.8 million low-income people, are exposed to AI-related decision-making through landlords' use of background screening, and more are exposed through the use of rent-setting algorithms and ubiquitous surveillance technologies.

**Employment**

32.4 million low-wage workers are exposed to AI in the context of work, such as AI-related decision-making through employers' use of AI to determine who gets hired, to set workers' pay, and to surveil and manage them.

**Education (K-12)**

13.25 million low-income children are exposed to AI-related decision-making through school districts' use of AI to determine if they are likely to drop out or engage in criminal activity, and more are exposed through ubiquitous AI surveillance technologies.

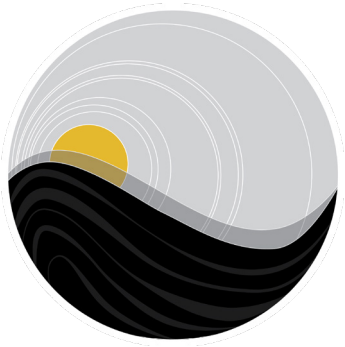

**Language Services**

5 million low-income people with limited English proficiency are exposed to some form of AI-based translation in government benefit offices, schools, hospitals, medical clinics, law enforcement agencies, or courts.

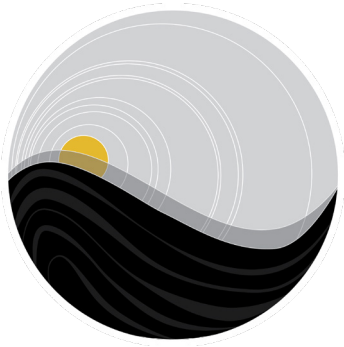

**Domestic Violence**

2 million low-income people are exposed to AI-related decision-making through police departments' use of AI to assess the risk of further violence to survivors.

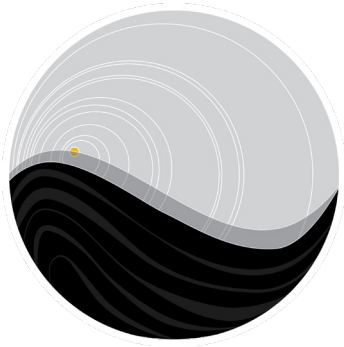

**Child Welfare**

72,000 low-income children are exposed to AI-related decision-making through government child welfare agencies' use of AI to determine if they are likely to be neglected.

**Social Security disability benefits (SSI and Social Security Disability Insurance): About 13.8 million people, including 10.6 million low-income people**, are exposed to AI-related decision-making in the Social Security Administration's disability benefits programs through current or planned uses of AI technologies in various parts of benefit administration, including the eligibility determination process and enforcement of asset limits. As a result, people experience temporary or permanent losses of income and are wrongly accused of being overpaid benefits.

**Unemployment Insurance: About 1.1 million low-income people** are exposed to AI-related decision-making in the Unemployment Insurance program through the eligibility and enrollment process, identity verification practices, and the detection of alleged fraud. As a result, people are denied critical income, experience grave delays in receiving benefits, and are falsely accused of fraud.

**Housing: About 16.3 million low-income households, or 39.8 million low-income people**, are exposed to AI-related decision-making through landlords' use of background screening systems, and more are exposed through rent-setting algorithms and surveillance technologies. As a result, people are denied housing, must pay higher rents than they otherwise would, and experience the pressures of being constantly watched while at home.

**Employment: At least 32.4 million low-wage workers** are exposed to AI in the context of work. Of these, at least 24.4 million low-wage workers are exposed to AI-related decision-making through employers' use of AI to determine who gets hired and to surveil, manage, and evaluate them. And an additional **8 million low-wage "gig" workers who work full-time** have their wages set by AI. As a result, people are denied job opportunities, fair pay, and fair working conditions.

**Education (K-12): About 13.25 million low-income children** are exposed to AI-related decision-making through school districts' use of AI to determine if they are likely to drop out or engage in criminal activity, with many more exposed through the ubiquity of AI surveillance technologies. As a result, children are labeled as failures, harassed by law enforcement, and experience the pressures of being constantly watched.

**Language services: About 5 million low-income people with limited English proficiency** are exposed to some form of AI-based translation in government benefit offices, schools, hospitals, medical clinics, law enforcement agencies, or courts. As a result, they experience delays, denied benefits, misunderstandings about vital information, or inability to access services.

**Domestic violence: At least 2 million low-income people** are exposed to AI-related decision-making through police departments' use of AI to assess the risk of further violence to survivors. **In addition, countless more survivors** are subjected to abusers' use of AI to create deepfakes to coerce them and to falsify evidence in court. As a result, people experience violence not being taken sufficiently seriously by authorities, diminished ability to escape violent situations, and greater obstacles to obtaining any relief in court.

**Child welfare (sometimes called "family policing"): At least 72,000 low-income children** are exposed to AI-related decision-making through government child welfare agencies' use of AI to determine if they are likely to be neglected. As a result, these children experience heightened risk of being separated from their parents and placed in foster care.

There are additional areas that feature extensive AI-related harms that are mentioned but not analyzed in this report, including the criminal legal system, the immigration system, access to credit, IRS tax auditing, health care treatment (independent of insurance), and voting.

# Understanding How AI is Different

Part two of this report identifies critical ways AI-based decision-making differs from human-centered methods that also may be inaccurate, irrational, or biased.

**AI provides a cloak of unwarranted rationality** to actions that intensify or perpetuate injustices. Some AI is based on supposedly valid statistical science, but much of it is not. In practice, though, the difference is largely meaningless. The alleged rationality and objectivity of the system allow the

users of AI to justify harmful actions, like benefit cuts or law enforcement harassment of students, or reinforce status quo power imbalances, such as that between employers and employees. Revealingly, the AI systems applied to low-income people almost never improve access to benefits or other opportunities.

**AI expands the scale of risk** beyond the reach of individual decision-makers. Without AI, the risk of harm from a benefits caseworker, employer, or landlord who makes mistakes, intends harm, or discriminates is limited to their personal reach. AI, though, centralizes and expands the risks of misapplied policies, coding errors, intentional cruelty, or bias to masses of people while erecting novel barriers to accountability that are more difficult to surmount than with humans.

**AI enables harmful functions** that were previously infeasible. Without AI, government workers simply lacked the capacity to perform the necessary administrative actions to accuse 40,000 people of fraud in the receipt of public benefits, as Michigan's Unemployment Insurance agency did. Or landlords would not have been able to easily collude at the scale needed to meaningfully control rental prices in medium or large markets. AI allows this to happen.

**AI enables policymakers to weaponize** insufficient government capacity. With AI's augmentation of government agencies' ability to act on masses of people, government officials can put administrative burdens on members of the public that the government agency cannot then handle. For example, when state officials implemented AI to determine the eligibility of people for Medicaid, agencies sent out so many requests for information that their staff could not timely process all of the information that people submitted in response. This resulted in improper denials and months-long delays of necessary medical treatments.

**AI makes the lives of low-income people fundamentally less predictable**. Part of AI's promise is that it discerns patterns that humans cannot. Low-income people subject to its decisions are not meant to be able to predict what will happen or why. In addition, the many design and implementation flaws in AI undermine the ability to know when something harmful might happen.

**AI feeds an unyielding desire** to collect more data about individuals at scales unfathomable before. This, in turn, threatens privacy and associated concepts of autonomy, expression, and intimacy and entrenches power imbalances between corporations and consumers who are unable to opt out of the regime of mass data collection and aggregation.

**AI extends the negative consequences** of past actions while limiting the ability of individuals to provide context. Using AI, unfortunate past events, including evictions, unpaid bills, arrests, and convictions are serially gathered, aggregated, and used to inform decisions about housing and jobs. Even when the information is accurate, it often does not account for a person's present circumstances. And the distillation of a person's experiences into a score or recommendation leaves little reason for decision-makers to inquire further.

**AI makes it harder for individuals harmed to contest decisions** either formally or informally. The quickest ways for people to contest decisions—conversations with the decision-maker, formal complaint processes, or, in the case of government decisions, administrative hearings—are generally made more difficult and less effective by AI.

**AI defies existing political, market, and legal accountability mechanisms.** Despite clear failures of public and private AI uses, governments, landlords, employers, and others continue to develop, buy, and use technologies that fail repeatedly in similar ways. Absent fundamental changes to incentive structures around AI usage, better outcomes for low-income people will be impossible.

# Actionable Recommendations

Part three of this report expands upon the below recommendations.

**Build the capacity** for integrated, multidimensional advocacy among low-income communities and their allies to immediately resist AI-related harms, including the loss of health insurance, housing, or work. Litigation, community activation and organizing, public education, and strategic narrative advocacy all are vital tools to give low-income people a fair chance to fight ground-level harms now. Existing organizations that work alongside low-income people, such as legal aid programs, community health centers, hunger relief agencies, and community organizing groups, all provide an existing institutional framework for harm reduction that strategic investment can bolster.

**Create the infrastructure** to connect low-income people and their allies to AI policy discourse. People exposed to AI-related harms have a right to have a say in its use. Doing so requires, among other things, understandable and actionable summaries of key AI policy ideas for affected communities and their allies, platforms to share information about AI policy and public participation opportunities, story-telling training and support, connections to journalists and media platforms, and peer-to-peer information sharing.

**Push for laws** that regulate AI use and better protect economic stability, with such efforts backed by a more empowered constituency of low-income people and their allies. Ideas for regulating AI include banning AI for some uses, requiring better government contracting practices, promoting transparency, ensuring more robust and enforceable procedures for the development and implementation of AI, incorporating formal mechanisms for the meaningful participation of affected communities and their allies, imposing liability on the vendors who develop and sell AI, and regulating how data can be collected, maintained, used, and sold. To better protect economic stability, the adoption or expansion of affirmative rights to health care, benefit programs, fair pay, paid leave, child care, job stability, housing stability, and privacy could improve the material conditions of low-income people and curb harmful uses of AI.

**Invest in enforcement** of existing and, once passed, new laws that limit harmful AI use. Laws matter mostly when they are enforced. But unaffordable legal services, understaffed government enforcement agencies, industry capture of regulators, collapsing media, and judge-driven roll-back of the ability of private citizens to enforce laws all undermine effective enforcement. Providing needed resources and otherwise spurring more assertive enforcement actions can make a difference now.

**Develop a positive vision** for AI technology. All AI is not inherently bad. While the chances for AI to directly improve the lives of low-income people are limited by structural forces, there are possibilities that might not require the entire reorganization of society. For example, state governments already know the information needed to enroll most eligible low-income people in public benefit programs. But they would need proper AI eligibility and enrollment systems built by competent entities and administered by competent staff, all of whom face enforceable accountability mechanisms. Outside of such uses, there are also possibilities for AI to help low-income communities assert their rights, exercise their power, or be involved in the development and building of AI that reflects their priorities.

# Defining Artificial Intelligence and AI Exposure

## Artificial Intelligence

This report takes a broad view of artificial intelligence to encompass technologies that span different eras and levels of technical sophistication. Here, AI can include complex algorithms based on advanced statistical modeling, simple algorithms that automate decision trees, process automation, data matching, data mining, surveillance technologies that involve computerized analysis, like facial recognition, different forms of machine learning, large language models, and generative AI.

The key feature in the AI systems discussed in this report is some form of automated decision or recommendation—that is, one made by a computer system—based on data about a person or other people similar to that person. When applied to low-income people, such decisions are usually made with little human oversight.

The report's broad approach to AI aligns with international bodies, including the Organisation for Economic Co-operation and Development, or OECD, which defines AI as follows:

> *An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.*[3]

Some people take issue with this broad definition of AI, instead preferring one encompassing only the latest sophisticated systems that appear to learn on their own and make decisions based on changing criteria that the AI systems themselves cannot always identify. Ultimately, though, the technical complexity of the system is not what determines the harm. Older or less complex systems might work differently, but, to the people subject to their decisions, the systems are just as inscrutable and do just as much harm as the newest technologies. That is because, at bottom, the different kinds of technologies share common roots and purposes centered on political and market forces that are hostile toward low-income people.

Certainly, these technologies were not built to advance the interests of low-income people. As a matter of course, these communities do not participate in the technologies' development, are not consulted about its use, have no formal power to stop them before the harm is done, and cannot meaningfully opt out of their use.

3 Organisation for Economic Co-operation and Development, "OECD AI Principles Overview" (n.d.), available at https://oecd.ai/en/ai-principles (last accessed October 21, 2024). U.S. policymakers, including the White House and Congress, have offered similarly broad definitions. See National Artificial Intelligence Initiative Act of 2021, Public Law 116–283, 117th Cong., 1st sess. (January 1, 2021), 15 U.S.C. 9401(3), available at https://uscode.house.gov/view.xhtml?req=(title:15%20section:9401%20edition:prelim); The White House, "Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," October 30, 2023, available at https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/.

While this report adopts a broad definition of AI, it, with a few exceptions, limits its scope to situations where AI makes or recommends decisions that directly touch the lives of low-income people. This guiding principle distinguishes it from, for example, early computer database systems used to administer public benefits. Those earlier systems offered a central repository for information relevant to determine a person's eligibility for services or benefits, but a caseworker would be charged with applying program rules to the system's information to make a decision.

This focus on decisional AI excludes some hot topics in AI, such as large language models and generative AI. Though alluring, these excluded topics do not tend to immediately shape the material circumstances of low-income people. Also left uncovered are ways in which various actors, including government agencies, use AI in the hopes of making purely internal processes more efficient. Such efforts only affect low-income people indirectly. Other topics, including the environmental impact of AI, do shape low-income people's circumstances but do not involve the application of AI to make decisions. Thus, these areas are not analyzed in depth in this report.

Finally, the report does not extensively discuss several areas where AI is used to make decisions: the criminal legal system, the immigration system, tax enforcement, child support enforcement, credit, health care treatment (independent of insurance), and voting. These areas are critically important to any consideration of AI's impact on low-income people, but they are not covered here because they are too far outside of my current expertise or have been covered robustly by others.

# AI Exposure and Methods

Although scholars and advocates have identified where and how AI impacts the lives of marginalized people in compelling, groundbreaking ways,[4] nobody has yet comprehensively explained and quantified for a broad public audience the extent of AI's reach in the lives of low-income people.

To do so, this report relies on my ground-level expertise as an anti-poverty advocate with roughly 20 years in the field (13 years as a lawyer and 6 years as a nonlawyer or a lawyer-in-training), insights from trusted colleagues, a robust body of existing reporting, scholarship, and policy explainers, and reasonable applications of existing statistical data. No new data was collected for this report. Some of the available data is older (e.g., surveys dating to 2015 or 2017) and thus likely underestimates the present use of AI.

The report's use of extant statistical data is fairly basic and does not attempt any complicated mathematical processes. I have explained my assumptions in reaching the figures offered.

4 See, for example, Benefits Tech Advocacy Hub, "Case Study Library" (2024), available at https://www.btah.org/case-studies.html; Ruha Benjamin, Race After Technology (Cambridge, UK: Polity Books, 2019); Danielle Keats Citron, "Technological Due Process," Washington University Law Review 85 (6) (2008), available at https://journals.library.wustl.edu/lawreview/article/id/6697/; Virginia Eubanks, Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor (New York City: St. Martin's Press, 2019); Michele Gilman, "Poverty Lawgorithms" (New York City: Data & Society, 2020), available at https://datasociety.net/wp-content/uploads/2020/09/Poverty-Lawgorithms-20200915.pdf; Safiya Noble, Algorithms of Oppression: How Search Engines Reinforce Racism (New York City: NYU Press, 2018); Cathy O'Neil, Weapons of Math Destruction (New York City: Penguin Random House, 2016); Brian Jordan Jefferson, Digitize and Punish: Racial Criminalization in the Digital Age (Minneapolis, MN: University of Minnesota Press, 2020); Rashida Richardson and Amba Kak, "Suspect Development Systems: Databasing Marginality and Enforcing Discipline," University of Michigan Journal of Law Reform 55 (2022), available at https://repository.law.umich.edu/mjlr/vol55/iss4/4/; Mizue Aizeki, Matt Mahmoudi, and Coline Schupfer, eds., Resisting Borders and Technologies of Violence (Chicago: Haymarket Books, 2024).

Importantly, some data is simply not available, particularly for private uses of AI by actors such as landlords and employers. Data about government use is mixed. Information about state AI systems for the most prominent or controversial usages, including for Medicaid, is easier to find than for usages that have not given rise to widespread public attention or court challenges. Information about usage by sub-state jurisdictions, such as sheriffs' offices, police departments, school districts, or counties charged with child welfare administration, is rarer.

Of course, opacity is a feature of AI systems and not a bug. Governments and businesses that use AI and developers who sell it regularly refuse to disclose the mechanics of how their AI systems work, the statistical modeling underlying it, projections about its likely impact, or assessments of its post-implementation impact. They often claim that such information is a proprietary trade secret or that divulging such information would allow people to manipulate the system. Meanwhile, some governments do not even keep such data.

Bearing these data limitations in mind, **the report's key metric is AI exposure**, an estimate of the number of low-income people subjected to AI that makes or recommends decisions about their lives. The notion of exposure is appropriate because it captures the risk inherent in AI systems. That is, people who may not be harmed by a given AI system today could potentially be harmed by it at any point in the future when they are subject to it. As with other forms of risky exposures, such as environmental hazards, the concept of AI exposure contemplates that exposure is not voluntary.

Admittedly, the metric of AI exposure concedes something to practicality. Limitations on existing data make it infeasible to determine which of the people exposed to AI are harmed by it in many of the domains covered. Consequently, the report describes the harms, suggests the scale of them, and, where possible, tries to quantify them.

Importantly, the number of people exposed to AI is a point-in-time count, likely to lead to systematic undercounting of the number of people who have ever been exposed to or hurt by AI-based decision-making and its attendant risks. For example, the AI exposure metric in Medicaid concludes that everyone currently receiving Medicaid is exposed to AI. But there are many people who received Medicaid in the past and were exposed who are

not counted now because they were terminated by it, have not sought to reenroll, or are now ineligible. Likewise, people who apply for or receive benefits in the future will be exposed to it. As people shift into and out of low-income status, their level of direct exposure changes.

I hope that candid discussions of data limitations and the challenges of quantifying AI exposure and AI harms invite ideas about both what to quantify and how to do so. Ideally, people with more robust statistical skills can reach better estimates in future iterations, and advocates, researchers, and journalists can help to fill in knowledge gaps about where AI-based decision-making is happening, who is developing and using it, and who is being helped or hurt by it.

Before moving on, a note on possible information bias. Admittedly, instances where AI was implemented with less harm may generate less attention than obviously harmful instances. It is therefore possible that available information skews toward the harmful and does not fully represent neutral, or even beneficial, uses of AI to make or recommend decisions about the lives of low-income people in the areas covered by this report.

That said, I am attuned to the subject matter, regularly seek out reports or articles on it, and am in frequent contact with advocates who work on the ground and on the policy level, academics, government officials, and industry officials. The report does not intentionally omit any known AI used to make or recommend decisions about the lives of low-income people on grounds that it would make AI look less harmful.

# Introduction to TechTonic Justice

In the age of AI, the hallmarks of daily life—home, work, school, stores, health clinics, hospitals, and government offices—are changing. TechTonic Justice serves the people these changes are leaving behind.

For some people, AI may make life more convenient. People can write a work report more quickly, grab groceries without stopping to pay, or have a doctor talk to them without typing out notes. And, in a few cases, the technologies may lead to something more than convenience, such as safer cars or more effective medical screenings and treatments.

Yet for people struggling to make ends meet, AI has made it harder to fulfill the most urgent needs. When looking for an apartment, renters meet landlords who use AI to set rents higher or score potential tenants based on false or outdated information. At work, employees face bosses using AI to decide who to hire, to set their pay, to make their schedules, or to keep them from talking to others about making their workplace safer or fairer. At school, AI labels students as “high risk,” and teachers and principals either

write them off as failures or get the police involved when they have never been involved with any crime. At government offices, AI wrongfully denies people eligible for health care, food assistance, unemployment insurance, or disability benefits or unfairly accuses them of fraud, all at the agency's direction.

Even within families, parents might have government social workers show up after AI decides that they are at high risk of neglecting their children. People experiencing domestic violence may struggle to escape abusers with AI-powered tools used to coerce them or to falsify evidence.

In short, AI is inescapable for low-income people, and it is making an already-difficult fight for stability harder by wrongly denying opportunities at every turn. This is not the future I want. I would rather fight for a world where people have what they need and where technology helps people—especially people living in poverty—to have healthy food, safe housing, fair work, and good schools.

Communities facing poverty, discrimination, or other kinds of exclusion have always fought back against injustice. Along the way, leaders and allies emerge: organizers coordinating action and building momentum, lawyers fighting in the courts, advocates raising money and awareness, journalists documenting the struggle, artists capturing the spirit, and protesters putting their bodies on the line. Working as a movement, with much time and toil, they often win changes and make things better, at least for a while. Even in defeat, they challenge common notions of what is possible, calling forth bolder visions for brighter days. They are artists of advancement and agents of imagination, mobilizing the resources that everyday people need to oppose the powerful interests perpetuating an unjust status quo.

Now, the biggest tech companies in the world are developing AI that they, banks, investors, landlords, employers, and governments are using to rewrite the ways we live out our daily lives. They promise efficiency, economic revolution, and scientific advancement at some unspecified point in the future that, according to them, is sure to come soon. But we are still waiting. And even if those developments do happen, who will benefit? Rising tides might lift all boats, but they do not change rafts into yachts.

The vision of a world remade by AI leaves low-income people behind. Each time it does, whether it's canceled health insurance, lower wages, or lost housing, low-income people have few places to turn. They are struggling financially and do not have the money to wait things out. Most lawyers are too expensive. And, while community-focused social services and legal aid programs do their best with limited resources, these organizations do not currently have the knowledge, capacity, or influence to push back at scale against AI and the complex issues that come with its use.

TechTonic Justice exists to fill this gap. We strengthen local justice movements—the legal aid lawyers who serve low-income folks at no charge, front-line advocates in places such as community health clinics, and affected communities—to effectively resist the harms that AI causes. We train, learn from, litigate alongside, organize with, offer strategic assistance to, and build infrastructure to support people needing immediate ways to fight AI systems. And we connect the experiences of people on the ground to the ongoing policy debates about curbing unsafe AI. This way, we can ensure that the people who must live with the harmful decisions that AI makes have some say in how that AI is used.

I founded TechTonic Justice after spending 13 years as a lawyer working at Legal Aid of Arkansas alongside low-income people to fight various kinds of injustices. I worked with survivors of domestic violence to seek relief from abuse. I supported low-wage workers to get paid the money they earned, fight discrimination, and resist other unfair labor practices. I served disabled kids to get appropriate K–12 education and a meaningful chance to learn. And I fought alongside low-income folks to access the public benefits for which they were eligible but wrongly denied.

At Legal Aid, I saw that my clients kept facing new ways that state government was using AI and related technologies to cut benefits. The harms still haunt me. My first AI-related battle involved Arkansas' Medicaid program, which provides in-home caregiver services to disabled people who would otherwise have to live in nursing facilities. The caregiver helps with daily activities, from getting out of bed and bathing to using the bathroom, eating, and running errands.

In-home care is cheaper for the state and better for people's dignity and independence. But when the state replaced nurses' professional judgement with an algorithm-based system for determining eligibility, people whose conditions had not gotten better faced drastic cuts to their benefits. They were forced to lie in their own waste. They got bed sores from not being turned. They were shut in at home and shut out of their community.

We—the people suffering, me, my Legal Aid teammates, and other allies—fought back in the courts. We activated hundreds of people who depended on the program to learn the key policy details, support one another, talk with reporters, and share what was going on more broadly. And, eventually, we won in the courts and in the legislature, forcing the state to abandon the system. Our victory was recognized as one of the first in the world against government use of harmful AI systems[5] and is still held up six years later as an example of what is possible. TechTonic Justice promises to help make many more successful stories of resistance.

People are at the heart of what I did as a Legal Aid attorney. TechTonic Justice will carry the same spirit. With that in mind, TechTonic Justice offers this document: part report, part warning, part roadmap, and part vision. Put together, it is the manifesto for why we exist and what we are going to do. Join us in the fight!

5 Importantly, legal aid programs and affiliated organizations have a long history dating to the 1970s of litigating successfully against problematic technology used in public benefit programs. Our victory in Arkansas was informed by and built upon this existing legacy. Our campaign also ran alongside similar efforts in other states against similar sorts of systems, with the work of Idaho and North Carolina advocates being particularly inspiring and helpful.

# Report Overview

This report is intended to introduce the broad public to the intersection between people living in or near poverty and the all-encompassing reach of AI. The report does not assume that any reader will know much about these topics beforehand and offers explanatory overviews of all key concepts. Hopefully, the report also will be useful to people with existing expertise in these topics.

The report has three substantive parts and an appendix.

**Part One: How AI Decides the Lives of Low-Income People.** This section contains the thrust of the report: an issue-by-issue discussion of the ways AI is used to decide key facets of low-income people's lives, with estimates for the number of people exposed to AI. It covers public benefits, employment, housing, K–12 schooling, child welfare, and domestic violence, with passing reference to other harmful AI uses that are not analyzed fully. The report uses some instances of AI use to touch on broader themes that are discussed in more detail in Part Two.

**Part Two: Innovating Injustice.** This section discusses how AI harms are distinct from the human-based decision-making that precedes it, analyzes existing protections for low-income people hurt by AI and their limitations, and highlights efforts of low-income people who have successfully fought back.

**Part Three: The Way Forward.** This section discusses how future accountability and harm-reduction efforts can be strengthened and offers a positive vision for ways that AI could be used or is being used to support low-income people.

**Appendix: Definitions and Methods.** This section provides context to the report's findings and analysis by defining and discussing the concepts of "poverty" and "low-income."

**linkedin.com/company/techtonicjustice**

**@techtonicj**

**@techtonicjustice.bsky.social**

**TechTonic Justice**

**@ techtonicjustice**

**TECHTONIC JUSTICE**

**techtonicjustice.org**