

Dalberg

Smarter systems, higher stakes:

AI in credit, benefits, and housing support



MAY 2026

Smarter systems, higher stakes: AI in credit, benefits, and housing support

This working paper identifies AI-enabled services that could materially influence economic mobility for low- and moderate-income (LMI) communities in the U.S. over the next decade. AI is increasingly shaping the systems that affect how LMI communities access financial services, public benefits, and housing support. This paper does not attempt to cover the full range of these changes. Instead, it focuses on select examples in these three areas where AI may strengthen the delivery of public, private, and philanthropic services for LMI communities, while also examining the risks these approaches raise.



Table of Contents

1.	Introduction	4
2.	Methodology	5
3.	Findings	6
4.	Reflections and Next Steps	18

1. Introduction

Artificial intelligence systems are increasingly being used by financial institutions and public agencies to support how people access information, services, and decisions related to their financial wellbeing. Among AI's many capabilities, five are particularly relevant to advancing economic opportunity:

- i. **navigation**, which helps users find and understand relevant information within complex systems such as benefits portals or financial products;
- ii. **personalized content generation**, which produces plain-language explanations, multilingual translations, and responses tailored to the user's inputs, context, or administrative record;
- iii. **targeting**, which uses predictive analytics to direct resources toward those most likely eligible, at risk, or positioned to benefit;
- iv. **matching**, which pairs users with suitable opportunities, products, or services; and
- v. **decision support**, which synthesizes data to inform human decisions without replacing human judgment.

In practice, these capabilities often work in combination. A benefits chatbot, for example, may simultaneously generate plain-language guidance, navigate program rules, personalize responses to a household's circumstances, and support a caseworker's decision-making. Together, these functions enable AI to analyze large volumes of complex data, support or partially automate routine tasks, and tailor assistance to individual needs at scale.¹ The extent and value of these benefits depend on system design, data quality, human oversight, and the institutional context in which the tools are deployed.

Across the ecosystem, actors express both optimism and caution about the potential for AI to advance economic mobility. Some see AI as a powerful lever to enhance systems and improve information sharing, yet they also raise concerns about uneven readiness and capacity for adoption.² Moreover, as these technologies expand, they also introduce new risks, among them infringement on privacy, increased potential for exploitation, bias, and opacity in decision-making, and inequities in access. The question is no longer whether AI will influence key systems that shape economic mobility but rather how, for whom, and under what safeguards. Its design, governance, and deployment in the coming years will determine whether it narrows or deepens existing gaps in access and inclusion.³

Recognizing these opportunities and risks, major philanthropic and private investments are accelerating experimentation and policy research at the intersection of AI and economic mobility. NextLadder Ventures, a new initiative backed by Ballmer Group, Gates Foundation, Stand Together, Valhalla Foundation, and John Overdeck; Humanity AI, a newly formed philanthropic coalition, Google.org's AI Opportunity Fund; and the GitLab Foundation's AI for Economic Opportunity Fund are among the largest efforts to date.^{4, 5, 6, 7} While these investments and emerging responsible AI frameworks and regulations⁸ demonstrate momentum, investment in AI applications designed to serve public and social goals remains a small fraction of overall private investment in AI.

1 [BIS, *Intelligent financial system: how AI is transforming finance* \(2024\)](#)

2 [Federal Reserve Bank of Boston, *Facing uncertainty: In AI visit, Boston Fed president hears from those confronting it, reducing it* \(2025\)](#)

3 Expert interviews

4 [Gates Foundation, *National philanthropies join forces to advance economic mobility and access to opportunity in America* \(2025\)](#)

5 [Google Public Policy, *Google.org's \\$75M AI Opportunity Fund for the United States* \(2025\)](#)

6 [MacArthur Foundation, *Humanity AI Commits \\$500 Million to Build a People-Centered Future for AI* \(2025\)](#)

7 [GitLab Foundation, *Catalyzing transformative economic opportunities through advanced AI technologies*](#)

8 See, for example, [California's SB 53](#), [EU AI Act](#), [Workers First Initiative on AI](#), [Black Tech Agenda](#)

2. Methodology

This working paper was developed through a mixed-methods process that combined a landscape scan, expert input, and targeted deep-dive analysis. We began with a bottom-up scan of philanthropic, private-sector, and public “big bets” at the intersection of artificial intelligence and economic mobility in the U.S., reviewing a wide range of initiatives and examples across four pathways to economic mobility:

- i. employment and workforce development,
- ii. small business and entrepreneurship,
- iii. promoting access to credit, investment, and financial services, and
- iv. housing and neighborhood revitalization.

We then grouped these initiatives and examples into 25 use cases, and conducted a rapid assessment against three criteria (impact, feasibility, and risk) to identify three promising use cases for deeper analysis. These selected use cases were then developed into deep-dive case studies, integrating desk research and expert insights. This analysis is directional rather than exhaustive and the inclusion of specific examples and case studies in this paper does not imply endorsement by Dalberg Advisors.

3. Findings

Our analysis of active AI applications and institutional interest elevated three high potential use cases where AI can enhance existing services to accelerate economic mobility:

1. **AI underwriting and fair credit access:** Machine-learning models and alternative data approaches can help lenders more efficiently assess repayment capacity for individuals with limited credit history, expanding safe access to financial products.
2. **AI-enabled public benefits and service delivery:** AI tools can streamline eligibility determination, automate recertification, and reduce administrative burden to shorten time-to-benefit and decrease error rates.
3. **AI-powered predictive analytics for eviction prevention:** Predictive analytics can help direct limited funds (e.g., rent assistance and eviction-prevention support) to households most likely to face hardship.

Across these areas, examples range from scaled service-delivery platforms to early pilots and proof-of-concept applications. The following deep dives examine the evidence base, illustrative examples, and implementation risks within each use case.

DEEP DIVE #1:

AI underwriting tools can enhance how financial institutions assess creditworthiness, with the potential to expand inclusion and improve risk assessment when paired with strong fairness and explainability safeguards.

— Problem Statement

Nearly one in eight U.S. adults—about 32 million people—are “credit invisible” or “unscorable,” lacking sufficient credit history to access affordable financial services and often relying instead on high-cost alternatives such as payday loans.⁹

— Impact Potential

Access to credit, particularly for LMI households, can provide a financial safety net that helps individuals and families manage major expenses, stay afloat during financial hardships, invest in their future, and pursue opportunities such as homeownership, education, and entrepreneurship.¹⁰

However, traditional credit scoring systems often underserve LMI consumers. Financial institutions typically rely on indicators like credit history and repayment records to determine creditworthiness, which can leave out borrowers without an established credit profile.¹¹ These consumers are limited from accessing financial products like mortgages or credit cards and may turn to informal borrowing or high-cost alternatives such as payday loans which can have steep fees and interest, further limiting their ability to build credit and financial stability.¹²

⁹ Consumer Financial Protection Bureau, Technical correction and update to the CFPB’s credit invisibles estimate (2025)

¹⁰ Opportunity Insights, Credit Access in the United States (2025)

¹¹ Oliver Wyman, Financial inclusion and access to credit (2022)

¹² Urban Wire, Adopting Alternative Data in Credit Scoring Would Allow Millions of Consumers to Access Credit (2021)

AI-enabled underwriting tools that use cash-flow data, rent and utility payments, and transaction histories have demonstrated early gains in expanding access. These models aim to generate a more holistic picture of creditworthiness by capturing indicators of financial reliability that traditional credit files overlook.

not paired with independent oversight and clear accountability standards. This shift toward proactive bias monitoring could strengthen accountability within compliance functions as AI becomes more embedded in credit decisioning.¹⁴

This expansion is clear in market data. The overall AI underwriting segment was valued at \$2.6 billion in 2023, while the emerging AI agents in the financial services market (which includes intelligent systems that use large language models and machine learning to automate customer service, fraud detection, compliance, and credit decisioning) reached \$490.2 million in 2024 and is projected to grow nearly ninefold to \$4.5 billion by 2030.^{15,16} This scale-up signals that, in the private sector, automated credit evaluation is moving beyond pilot stages towards broader adoption. Within five years, consumers seeking small business, personal, or short-term digital loans outside traditional lending channels may increasingly interact with AI credit agents.

Together, alternative underwriting approaches and fairness-monitoring tools could help bring some consumers without traditional credit scores into the formal credit system. If implemented responsibly and validated over time, these approaches may expand access to mainstream financial products and support greater household stability.¹⁷ If successfully scaled, this could reshape how households access financial services, making algorithmic credit assessments a first touchpoint with the formal financial system rather than a gatekeeping mechanism.

HO Early Bets and Examples in Practice

Several market actors are testing alternative underwriting approaches and fairness-focused models that show early promise to expand credit access:

- **Tilt**, a consumer credit card issuer, uses algorithmic underwriting based on 250+ non-traditional data sources to assess thin-file and no-file consumers. Tilt has matched over 5 million customers with products suited to their needs.¹⁸
- **Nova Credit**, in partnership with Mastercard Open Banking, uses consumer-permissioned financial data to help immigrants, gig workers, and other thin-file borrowers qualify for credit. Through its Credit Passport, Nova Credit translates international credit histories from more than 20 countries into U.S.-equivalent credit reports. Meanwhile, its Cash Atlas tool allows applicants to securely share real-time banking data, such as income, spending,

¹³ [Urban Wire](#), *Including Rental Payment History in Underwriting and Credit Scores Could Expand Access to Credit* (2025)

¹⁴ [Toreini, E., Mehrnezhad, M. & van Moorsel, A.](#), "Fairness as a Service (FaaS): verifiable and privacy-preserving fairness auditing of machine learning systems", *Int. Journal of Information Security* 23 (2024)

¹⁵ [Market.us](#), *Global AI in Underwriting Market Report* (2024)

¹⁶ [Grand View Research](#), *AI Agents in Financial Services Market* (2025)

¹⁷ [OECD](#), *Generative artificial intelligence in finance* (2023)

¹⁸ [Tilt](#), *About* (2026)

and savings patterns, through open-banking connections, giving lenders a more complete and current view of financial stability. By embedding advanced analytics and behavioral pattern recognition into existing underwriting workflows, the platform demonstrates how open-banking and data-driven tools can expand responsible credit access for consumers traditionally excluded from the financial system.¹⁹

- **Zest AI** provides lenders with AI-automated underwriting software to assess and approve applicants. According to company materials, its machine learning technology analyzes thousands of variables, which the company claims to enable lenders to evaluate borrower risk with two to four times greater accuracy than traditional models. With this technology, Zest AI enables institutions to automate lending decisions and significantly reduce processing time.^{20, 21}
- **FairPlay** offers “fairness-as-a-service” tools that detect and correct bias in automated AI-enabled decisions.²² For example, FairPlay’s Agentic Assurance Platform will stress-test the output of AI agents to identify vulnerabilities and automatically generate documentation aligned with leading risk management frameworks.²³

These examples are illustrative of the kinds of approaches emerging in the market, but in many cases the strongest available evidence remains company-reported or based on partner case studies rather than independent long-term evaluation.

🔗 Risks

While early bets hold strong potential to advance financial opportunities, adoption remains constrained by limited regulatory clarity and vendor control over proprietary models and datasets. Lower- and under-resourced lenders often lack the capacity to vet or adapt these

While early bets hold strong potential to advance financial opportunities, adoption remains constrained by limited regulatory clarity and vendor control over proprietary models and datasets.

models safely, creating a two-tier innovation landscape. As a result, these institutions may face regulatory constraints, internal bureaucracy, or narrow interpretations of fiduciary duty and may not be equitably equipped to understand, experiment, or adopt emerging innovations. These challenges may lead to an uneven market and widening information gaps between players, limiting the reach of innovation across credit markets.²⁴ Without strong safeguards for

fairness, transparency, and accountability, AI-enabled alternative data and assessment tools could end up entrenching the very inequities they aim to reduce.²⁵ Key risks include:²⁶

- **Algorithmic opacity:** Outsourcing credit decisions to AI can obscure responsibility when errors or discrimination occur. Without consistent human validation, transparency about how models work, and clear pathways for recourse, affected borrowers may have no way to contest outcomes or understand why they were denied. The fact that many ML-based scoring systems function as a “black box” complicates regulatory oversight – making it difficult for regulators and institutions to attribute liability, audit decisions, or ensure meaningful explanation to consumers.^{27, 28}

19 [MasterCard, Case Study: Nova Credit \(2025\)](#)

20 [Zest AI, Smart, More Fair, Efficient AI-Automated Underwriting \(2025\)](#)

21 [Zest AI, Zest AI Ranked as a Fastest-Growing Company in North America on the 2024 Deloitte Technology Fast 500™, 2025](#)

22 [FairPlay, Smarter AI: Better Decisions \(2025\)](#)

23 [PRNewswire, FairPlay and Arva AI Partner to Advance Powerful and Compliant Agentic AI in Financial Services \(2025\)](#)

24 Expert interviews

25 [Cavazos, The Impact of Artificial Intelligence on Lending: A New Form of Redlining?, Texas A&M Journal of Property Law](#)

26 Expert interviews

27 [Open Journal of Social Sciences, Artificial Intelligence and Machine Learning in Credit Risk Assessment: Enhancing Accuracy and Ensuring Fairness \(2024\)](#)

28 [GSC Advanced Research and Reviews \(GSCARR\), Legal accountability and ethical considerations of AI in financial services \(2024\)](#)

- **Inadequate communication of risks to investors:** Early AI-enabled credit underwriting use cases in particular highlight the need for rigorous communication standards in algorithmic decision-making. For example, in 2024, Upstart, an AI-enabled credit underwriting platform designed to expand credit access, reported approving about 43% more borrowers than traditional credit models.²⁹ However, Upstart has faced investor scrutiny for allegedly overstating its model's ability to account for fluctuating macroeconomic conditions.³⁰ Although related regulatory investigations concluded without enforcement action, a shareholder class-action lawsuit was marked as ongoing as of June 2025.³¹
- **Privacy violations:** Serious privacy violations have also arisen in the sector. International precedents underscore the severity of this risk. For example, across the more than 40 loan apps operating in Kenya—many of which leverage smartphone data for predictive underwriting—instances have occurred in which lenders shared users' contacts information for the purposes of debt collection. In response, Kenya passed stringent data privacy legislation and is investigating numerous companies.³² These cases underscore the importance of continued validation and governance frameworks for AI credit models.
- **Bias embedded in data and design:** Models can perpetuate discrimination if trained on historical data reflecting unequal access to credit or neighborhood-level disparities.
- **Digital exclusion:** Consumers with limited digital access, literacy, or English proficiency may struggle to engage with AI-enabled tools and processes that are not designed to be accessible to diverse user needs.
- **Institutional capacity and readiness:** Lower- and under-resourced lenders and CDFIs may lack the data infrastructure, technical expertise, and resources to evaluate or implement AI tools safely, while unclear or inconsistent compliance standards for AI in credit decisioning create uncertainty for traditional institutions. Together, these gaps may slow responsible adoption and risk concentrating innovation among a limited set of tech-agile actors.
- **Capital allocation bias:** Despite growing interest in AI for financial inclusion, venture return expectations and limited rigorous evidence risk skewing capital toward high-growth plays, leaving LMI-focused solutions underfunded. This dynamic reinforces the case for flexible experimentation capital discussed in Section IV.

If left unaddressed, these risks could deepen the divide between institutions with strong digital capacity and those serving LMI borrowers most directly, reinforcing structural inequities in the credit access ecosystem.

DEEP DIVE #2:

AI-enabled benefits systems can enhance how government programs deliver assistance—reducing administrative burden, improving accuracy, and ensuring support reaches eligible families faster.

Problem Statement

Each year, an estimated \$140 billion in federal benefits goes unclaimed in the U.S. as millions of eligible individuals miss out on critical supports due to fragmented benefits systems, complex eligibility rules, and burdensome documentation requirements.³³

²⁹ Upstart, *2024 Access to Credit Report* (2024)

³⁰ Stanford Law School Securities Class Action Clearinghouse, *Upstart, Inc. Case Summary* (2025)

³¹ Ibid.

³² TechCrunch, *Google clamps down on illegal loan apps in Kenya, Nigeria* (2022)

³³ Biden White House Archives, Office of Information and Regulatory Affairs, *Tackling the Time Tax* (2023)

Impact Potential

Federal assistance programs such as the Supplemental Nutrition Assistance Program (SNAP), Medicaid, and the Earned Income Tax Credit (EITC) help underserved households meet essential needs by providing food assistance, health coverage, and income support.^{34, 35, 36} Together, these supports form a critical safety net for LMI families. Despite their importance, an estimated 12% of eligible individuals in the U.S. do not receive SNAP benefits.³⁷ For a family of four, this means \$12,000 less to spend on food each year.³⁸

A national survey found that 40% of people who were eligible for SNAP but did not participate cited paperwork as the primary reason they did not apply, underscoring how administrative burdens prevent millions from receiving assistance.

A major barrier to accessing federal assistance programs is the complex and time-consuming application process. Complex eligibility rules and heavy documentation requirements make it especially difficult for older adults and LMI families, many of whom face digital access and literacy challenges, have limited English proficiency (LEP), or lack time and support to complete or even begin applications. A national survey found that 40% of people who were eligible for SNAP but did not participate cited

paperwork as the primary reason they did not apply, underscoring how administrative burdens prevent millions from receiving assistance.³⁹

Beyond SNAP, difficulties with the application process also impact Medicaid and EITC access, which are among two of the nation's largest safety net programs. In Medicaid, administrative hurdles such as complex renewal processes, inconsistent language access, and limited digital support have contributed to coverage losses during redetermination, particularly for LEP individuals. For instance, a study on Medicaid redetermination conducted in Illinois found that LEP respondents were more than five times likely to lose coverage than English-proficient applicants, largely due to challenges understanding renewal materials and completing forms.⁴⁰ Similarly, an estimated one in five workers eligible for the EITC fail to claim the credit each year, often due to confusion about eligibility rules or limited access to tax filing assistance.⁴¹ These administrative hurdles, combined with language, literacy, and digital barriers, have created systemic under-enrollment among LMI communities.

Streamlining application and renewal processes is especially valuable during this current period of fiscal pressure and policy change, as it can help government departments with limited administrative capacity to reach more households efficiently. The One Big Beautiful Bill Act of 2025 added new work requirements for SNAP and Medicaid in the context of budget cuts to the departments administering these programs.⁴² Helping beneficiaries navigate these changes could help millions to retain their benefits.

Digital tools are already demonstrating measurable improvements in helping applicants complete and maintain enrollment in benefits programs, and these gains can be further strengthened through AI-enabled features that improve personalization, accuracy, and responsiveness. For example, GetCalFresh, a digital application platform for California's SNAP program, replaced a lengthy and fragmented process, reducing average application time from

34 [USDA, Supplemental Nutrition Assistance Program \(SNAP\) \(2025\)](#)

35 [Medicaid.Gov, Keeping America Healthy \(2025\)](#)

36 [IRS, Earned Income Tax Credit \(EITC\) \(2025\)](#)

37 [USDA, Reaching Those in Need: Estimates of State SNAP Participation Rates in 2022 \(2022\)](#)

38 [Ibid.](#)

39 [Center for American Progress, How To Address the Administrative Burdens of Accessing the Safety Net \(2022\)](#)

40 [Journal of Immigrant and Minority Health, Medicaid Redetermination and Renewal Experiences of Limited English Proficient Beneficiaries in Illinois \(2021\)](#)

41 [IRS, IRS, national partners launch EITC Awareness Day on 50th anniversary of the Earned Income Tax Credit \(2025\)](#)

42 [Pearl Suite, Medicaid, SNAP, and the One Big Beautiful Bill Act Explained \(2025\)](#)

an hour to about 10 minutes, supporting at least 6.2 million applicants in accessing \$12.8 billion in food assistance, and contributing to a period in which California's SNAP participation rate rose from 66% in 2014 to 81% by 2022.⁴³ Its live and AI-enabled chatbot assisted applicants by answering common questions and directing them to relevant resources, showing that even narrowly focused tools can deliver meaningful support and accuracy without requiring complex or fully autonomous systems.⁴⁴ The platform's design principles of plain language, mobile-first design, and streamlined document uploads have since been incorporated into the state's multilingual, multi-benefit portal, BenefitsCal.com, which is now the primary application site for CalFresh and other state benefits.⁴⁵

Such guided eligibility assistants and multilingual chatbots could help households navigate complex rules and explain required documentation. If scaled further, these tools even have the potential to move us closer to the long-sought after goal of a 'one-stop portal' that clarifies how programs such as SNAP and Medicaid interact. This is an aspiration that has historically proven difficult to realize through technology alone, but that AI-enabled personalization may help advance incrementally.⁴⁶ As these tools can operate through mobile devices or voice-enabled platforms, they also have the potential to improve accessibility for millions of LMI families with limited literacy or broadband access, giving applicants an alternative source of support to in-person assistance. For example, personalized AI support tools could also provide step-by-step, multilingual guidance tailored to a household's circumstances, translate rules into plain language, prefill forms from user-provided data, flag missing documents, simulate eligibility across programs, send deadline reminders by SMS or voice, and connect applicants to a case manager when desired. Meanwhile, frontline worker copilots could close case management gaps by catching missed appointments and prompting timely follow ups, which could improve case resolution. Together, these capabilities could shorten processing times by weeks, reduce application drop-offs by double-digit margins in pilot contexts, and increase take-up among LMI families who otherwise face administrative barriers.⁴⁷ Over time, integrated AI-assisted enrollment systems could evolve into portals that pre-screen for multiple programs and auto-renew benefits when verified data remains unchanged, redefining how households interact with the social safety net. While these functions are promising, evidence of their effect on processing times, drop-off rates, and long-term take-up remains early and should be interpreted cautiously.

If deployed responsibly, AI-enabled eligibility and claims systems have the potential to increase participation among eligible households and help recover a portion of the roughly \$140 billion in public benefits that go unclaimed each year due to administrative barriers.

The administrative complexity of benefits access carries major economic costs, particularly as the federal government seeks to increase operational efficiency.⁴⁸ If deployed responsibly, AI-enabled eligibility and claims systems have the potential to increase participation among eligible households and help recover a portion of the roughly \$140 billion in public benefits that go unclaimed each year due to administrative barriers.⁴⁹ These technologies can also help reduce persistent inefficiencies

across federal and state systems. Outdated and manual processes are estimated to cost U.S. government agencies \$38.7 billion annually, while in 2017, Americans were found to collectively spend 11.5 billion hours each year completing federal paperwork.^{50, 51} Meanwhile, another

43 Code for America, *Simplifying California's Online Application for Food Benefits* (2025)

44 Code for America, *Human-Centered, Machine-Assisted: Ethically Deploying AI to Improve the Client Experience* (2023)

45 Code for America (2025)

46 Expert interviews

47 Expert interviews

48 Center for American Progress, *How to Address the Administrative Burdens of Accessing the Safety Net* (2022)

49 Joint Economic Committee, *Modernizing Benefit Systems Can Improve Lives While Also Saving Time and Money* (2024)

50 U.S. Chamber of Commerce Technology Engagement Center, *Government Digitization: Transforming Government to Better Serve Americans* (2022)

51 Center for American Progress, *How To Address the Administrative Burdens of Accessing the Safety Net* (2022)

study has estimated that public employees spend nearly 30% of their workday on tasks that could be automated, such as checking paperwork.⁵² By reducing administrative burden, AI-driven enhancements could enable a larger share of federal funding to reach LMI households directly instead of being absorbed by overhead costs. During recessions, for example, every dollar spent through SNAP generates an estimated \$1.54 in economic activity, underscoring how benefit delivery can simultaneously support families and boost local economies.⁵³ In this way, such AI-enabled enhancements have the potential to create a more responsive and resilient social safety net that improves outcomes for households, reduces long-term costs, and improves government efficiency.

Importantly, AI cannot by itself fix the deeper structural barriers that limit benefits access, including restrictive eligibility rules, fragmented program design, underfunded agencies, and uneven coordination across federal, state, and local systems. At best, these tools can make flawed systems easier to navigate and administer. Their long-term value will therefore depend on whether they are paired with broader policy and process reforms.

Early Bets and Examples in Practice

The examples below illustrate different approaches emerging in the field. Many remain in pilot stages or are documented primarily through partner or organization-reported evidence rather than independent evaluation.

- **Nava Labs** is developing and piloting AI-powered tools to help government staff and benefit navigators verify eligibility more efficiently and reduce manual review time, enabling teams to better identify families eligible for programs such as SNAP, Medicaid, and WIC. In 2025, the organization received a \$1.5 million grant from Google.org to build a suite of generative AI agents that navigate multiple databases and benefit portals to complete applications with caseworker oversight. The initiative, in partnership with Imagine LA and First 5 Riverside County, aims to improve caseworker efficiency, reduce administrative burden, and accelerate benefit delivery across California within three years. The project builds on Nava Labs' earlier pilot of an AI-powered chatbot for Imagine LA's Benefit Navigator tool, which showed that well-designed AI systems can streamline processes while keeping human caseworkers at the center of service delivery.^{54, 55}
- **GetCalFresh's** digital platform and AI-enabled chatbot assist California SNAP applicants by answering common questions and directing them to resources, demonstrating how simple, targeted tools can provide meaningful support at scale. The platform's design principles have since been incorporated into the state's multilingual, multi-benefit portal, BenefitsCal.com, which now serves as the primary application site for CalFresh benefits.⁵⁶
- **Propel, Inc.** received a \$300,000 grant from the Agency Fund to build and evaluate an AI copilot to help SNAP households navigate eligibility, paperwork, and renewals, and an AI chat assistant that offers state-specific and multilingual guidance to restore benefits, improving recovery times and easing workloads for state agencies.⁵⁷ It has also piloted AI tools to complement GetCalFresh's digital platform and reduce benefit interruptions, including a diagnostic flow that helps users identify and resolve missed deposits.⁵⁸
- **ConsiliumBots** designs and implements AI bots to help families connect to public programs and services in the U.S. and across Latin America. Its platforms, including Domus and Decidiendo un Futuro Mejor, use AI to organize complex data and provide personalized guidance, making information about benefits, educational opportunities, and other re-

52 [Joint Economic Committee](#), (2024)

53 [Peter G. Peterson Foundation](#), *How Are Safety Net Programs Helping the Response to the Coronavirus Pandemic?* (2020)

54 [Nava](#), *Case Study: Experimenting with AI-powered tools in public benefits* (2024)

55 [Nava](#), *Nava Labs receives Google.org grant to further AI tools with Imagine LA and Riverside County* (2025)

56 [Code for America](#), (2025)

57 [Propel](#), *Propel's mission: To build technology that strengthens the social safety net in America* (2025)

58 [Dave Guarino](#), *Using AI to help SNAP recipients diagnose and restore lost benefits* (2025)

sources easier to navigate. The organization is also developing AI-powered feedback tools that deliver alerts, message nudges, and recommendation reports through WhatsApp, email, or mobile apps to help families make informed decisions and close information gaps that limit access to support.⁵⁹

Collectively, these pilots suggest that modest automation can yield tangible results, cutting administrative workload while improving user satisfaction and retention.

🔗 Risks

While promising, recent research suggests limited current deployment of AI-enabled public benefits such as SNAP.⁶⁰ Furthermore, the nonprofit and tech sectors have recently faced setbacks in efforts to use AI to connect LMI families with public benefits. For example, Benefits Data Trust (BDT), a Philadelphia-based nonprofit founded in 2005, sought to streamline access to benefits using AI and data-driven outreach.⁶¹ Despite receiving a \$20 million unrestricted grant from MacKenzie Scott in 2022, BDT suddenly shut down in mid-2024, citing financial difficulties.⁶² To some observers in the philanthropic sector, BDT's collapse indicated the need to reform the broader public benefits system to prioritize accessibility, rather than relying solely on technological solutions from the nonprofit and private sector.⁶³ Thus, while pilot programs signal an opportunity to enhance public benefit systems, their success will depend on parallel investments in policy, data infrastructure, and human-centered design that make these systems easier to navigate and sustain beyond any single technological solution.

As agencies enhance systems and adopt automated tools, ensuring that these technologies protect sensitive data, operate transparently, and remain inclusive for all users will be critical. Risks include:⁶⁴

- **Data protection and privacy:** Expanding data sharing across agencies or using administrative records to train models raises the potential for misuse or unintended disclosure of personal information. Sustained public trust will depend on strong privacy safeguards, transparent data governance, and clear consent standards.
- **Accessibility and inclusion:** If automated systems are not designed for multilingual use,⁶⁵ mobile access, or low-connectivity environments, they risk excluding the populations that benefits programs aim to serve. Accessibility must be embedded from the outset, rather than retrofitted later.
- **Administrative readiness and oversight:** Many agencies lack the technical capacity, governance frameworks, or procurement processes needed to evaluate and monitor AI tools. Without these structures, innovation may outpace oversight, resulting in uneven performance, accountability gaps, or compliance risks.
- **Professional competency, ethical use, and workforce impact:** Social work practitioners using AI tools must be trained to critically interpret outputs, recognize inaccuracies or “hallucinations,” and understand ethical boundaries around client data, particularly when working on sensitive cases. Without clear guidelines, informed consent protocols, and appropriate workload adjustments, automation can shift cognitive burdens onto caseworkers, leading to burnout, reduced service quality, and diminished trust. Agencies should pair training and ethical frameworks with role and caseload recalibration to ensure AI enhances, rather than strains, professional practice.⁶⁶

59 [Consilium Bots, What We Do \(2025\)](#)

60 [American Public Human Services Association, AI-Powered SNAP Modernization \(2025\)](#)

61 [Center for High Impact Philanthropy, Benefits Data Trust \(2022\)](#)

62 [Chronicle of Philanthropy, Why \\$20 Million From MacKenzie Scott Couldn't Save This A.I. Nonprofit \(2024\)](#)

63 [Pew Research Center, Benefits Data Trust's Closure Should Prompt Us to Rebuild the Flawed Public Benefits System \(2024\)](#)

64 Expert interviews

65 [Code for America, Benefits Playbook: Designing Human-Centered Applications \(2024\)](#)

66 [Khan, Fatima, Qureshi, et al., Drawbacks of Artificial Intelligence and Their Potential Solutions in the Healthcare Sector \(2023\)](#)

- **Over-automation and loss of human judgement:** Fully automated systems can miss contextual nuance, especially in complex or exceptional cases. Sector experts have noted that human review by caseworkers will remain essential to ensure fairness, correct errors, and preserve applicants’ right to appeal or explain their circumstances.
- **Sustainability and dependence on technology:** Technological fixes cannot compensate for underlying policy and process fragmentation. The closure of BDT, despite significant philanthropic support, highlights the risks of overreliance on stand-alone technology solutions. Without broader policy and process reforms, even well-designed AI systems may struggle to achieve sustainable impact or remain financially viable.⁶⁷

If these risks are not managed, AI-driven enrollment could widen rather than close access gaps, creating a digital divide in public benefits participation between states and populations with unequal digital infrastructure.

DEEP DIVE #3:

AI-enabled housing stability systems use predictive data to target rental support, streamline case management, and strengthen the infrastructure that prevents eviction and financial distress.

Problem Statement

More than three million U.S. households face eviction filings each year. Many at-risk renters miss assistance because fragmented data systems, delayed communication, and limited administrative capacity prevent rental relief or legal aid from reaching households before they lose their homes.^{68, 69}

Impact Potential

Eviction remains one of the most disruptive shocks LMI renters can experience, often triggering job loss, credit deterioration, disruptions to education for children, and long-term housing instability. Each year, landlords file an estimated 3.6 million eviction cases across the U.S, disproportionately impacting LMI communities.⁷⁰ In many jurisdictions, such as New York City, eviction filings have returned to or exceeded pre-pandemic levels, reflecting both persistent affordability pressures and uneven access to early intervention programs.⁷¹

Each year, landlords file an estimated 3.6 million eviction cases across the U.S, disproportionately impacting LMI communities. In many jurisdictions, such as New York City, eviction filings have returned to or exceeded pre-pandemic levels, reflecting both persistent affordability pressures and uneven access to early intervention programs.

To promote housing stability during the pandemic, the federal Emergency Rental Assistance (ERA) program distributed roughly \$46 billion nationwide, providing essential relief to millions of renters.⁷² However, evaluations and implementation reviews suggest many local ERA programs struggled to identify and assist households before crises escalated. Fragmented data systems, limited capacity,

67 [Pew Research Center](#), (2024)
 68 [Eviction Lab](#), *Who is Evicted in America* (2023)
 69 [Yale Tobin Center for Economic Policy](#), *Eviction and Poverty in American Cities* (2024)
 70 [Eviction Lab](#), (2023)
 71 [New York City Comptroller](#), *Evictions Up, Representation Down* (2025)
 72 [U.S Department of the Treasury](#), *Emergency Rental Assistance Program* (2025)

burdensome documentation requirements, and slow review processes often meant aid reached renters only after eviction filings were already underway.^{73,74}

AI-driven tools could help shift eviction prevention from a reactive to a proactive system by enabling earlier identification of households at risk and faster delivery of support. Predictive analytics can integrate data from courts, landlords, and utility providers to identify households likely to face eviction within the next 30 to 60 days, using early signs of housing distress such as missed rent or rising outstanding payments.⁷⁵ This type of early-warning capability could allow social service agencies to prioritize outreach, connect tenants with mediation or legal aid, and deploy emergency rental support before court filings occur.⁷⁶ If paired with reliable data, strong privacy safeguards, and effective intervention pathways, these capabilities may help make housing assistance more timely and better targeted. Importantly, prediction alone does not prevent eviction. Any impact depends on whether households identified as at risk can be connected to timely mediation, legal aid, rental assistance, or other effective support.

Beyond targeting, AI could strengthen how limited housing assistance funds are allocated. Machine learning models could analyze application patterns, payment histories, and neighborhood-level indicators to help agencies distribute aid based on need and potential impact rather than on a first-come-first-served basis.⁷⁷ Automated document parsing and case-management tools could streamline verification and eligibility review, reducing administrative bottlenecks and accelerating aid disbursement.⁷⁸ Together, these capabilities could make housing assistance more timely, equitable, and fiscally efficient, helping local governments reduce displacement, prevent homelessness, and stabilize communities over the long term.

Critically, the quality of the predictive outputs is only as high as the quality of the underlying data. Researchers at the California Policy Lab (CPL) underscored that predictive modeling efforts require, at minimum, data on the outcome being predicted (e.g., homelessness), at least two sources of data on risk factors, and that these data must be linked and updated regularly. While many jurisdictions have access to this level of data, integrating sources and enabling regular updating capability are often key barriers.

HO Early Bets and Examples in Practice

While applications are more nascent in this sector, early pilots of predictive tools to improve housing stability show promise:

- **The Los Angeles County Homelessness Prevention Unit (HPU)** has explored AI-enabled predictive analytics to detect early signs of housing instability and proactively assist residents at risk of eviction. The HPU is run by the Housing for Health division of the Department of Health Services (DHS) and leverages a machine learning model developed by the California Policy Lab (CPL). The tool monitors data from multiple local government agencies, including records of residents who have received mental health crisis care or enrolled in public benefits such as CalFresh. Using machine learning, it identifies Los Angeles residents most at risk of homelessness. Case managers then reach out to these residents to enroll them in a support program that provides several months of case management and helps participants manage \$4,000–\$6,000 in assistance, which can be used for rent or other essential expenses. An early evaluation found that HPU participants were 71% less

73 [Government Accountability Office, Evictions: National Data Are Limited and Challenging to Collect \(2024\)](#)

74 [University of California Berkeley Terner Center for Housing Innovation, Lessons Learned from the Emergency Rental Assistance Program \(2022\)](#)

75 [Harvard Kennedy School Data-Smart City Solutions, Data-Driven Models for Eviction Prevention \(2025\)](#)

76 [California Policy Lab, The Homelessness Prevention Unit: A Proactive Approach to Preventing Homelessness in Los Angeles County \(2024\)](#)

77 [Association for the Advancement of Artificial Intelligence, Preventing Eviction-Caused Homelessness through ML-Informed Distribution of Rental Assistance \(2025\)](#)

78 [National League of Cities, How Local Leaders Can Use AI to Advance Housing Policy \(2024\)](#)

likely to enroll in a homeless shelter within 18 months, though the program's relatively small scale and evolving methodology suggest these results should be interpreted as promising rather than definitive. CPL researchers also noted that because the model predicts clients rather than relying on self-selection, it reduces fraud risk, and can enable departments to deliver assistance faster, without the full layers of upfront eligibility verification. Reducing verification layers comes with its own risks, and departments exploring these solutions will need to balance specificity, sensitivity, and speed in their targeting methods. To date, the program has served nearly 1,500 people, with 86% of participants retaining their housing upon completion.^{79, 80, 81}

- Penn State College of Information Sciences and Technology researchers in collaboration with the Child Poverty Action Lab** have developed AI models that estimate, by location and scale, the number of renters at risk of eviction to inform policy development and resource distribution. Their “Weakly-Supervised Aid to Relieve Nationwide Eviction Rate (WARNER)” model identifies neighborhoods with disproportionately high concentrations of households at risk of eviction by using socioeconomic characteristics associated with housing instability and tags satellite data, such as indicators of gentrification. This information is then used to train a machine learning model that detects similar eviction filing hotspots in other areas. Recognizing that reliable data on housing instability is often limited or difficult to collect, the model incorporates alternative data sources associated with eviction to enable more targeted reporting that can be scaled and adapted across different geographies in the U.S. Their second model, the “Multi-view Model Forecasting the Number of Tenants at Risk of Formal Eviction (MARTIAN),” projects how many tenants may be vulnerable to eviction at specific points in the future. Similar to the first model, it draws on data from available eviction filing records, the U.S. Census’ American Community Survey, and labor and employment statistics to estimate the number of tenants who may face eviction in each census tract. The research team partnered with the Child Poverty Action Lab to test both models using a real-world dataset from Dallas County, where eviction records are more comprehensive and accessible. The AI models outperformed some traditional approaches by as much as 36%.^{82, 83, 84}

Beyond these examples, two promising data integration initiatives could serve as platforms for AI-enabled prediction in the future:

- New York City’s Housing Data Coalition** and civic-tech partners maintain public data dashboards (e.g., Right to Counsel’s Eviction Crisis Monitor and Association for Neighborhood and Housing Development’s Displacement Alert Project) that integrate New York City housing information, such as eviction filings and building violations to support local knowledge and make data actionable for housing justice groups.^{85, 86}
- The Eviction Lab at Princeton University** has compiled the first nationwide dataset of eviction filings and built interactive tools for policymakers and communities to explore eviction risk and housing instability, laying important groundwork for the development of early-warning systems that combine administrative data and public records. The research center’s analysis of the prevalence, causes, and consequences of eviction in the U.S. aims to highlight how significant of an issue eviction is across the U.S. and demonstrate how cases are unevenly distributed both across and within states.^{87, 88}

79 [California Policy Lab](#), (2024)

80 [NPR](#), *Los Angeles is using AI to predict who might become homeless and help before they do* (2023)

81 [California Policy Lab](#), *Early Outcomes from the Los Angeles County Homelessness Prevention Unit* (2025)

82 [Tabar, Jung, Yadav, Chavez, Flores, and Lee](#); *WARNER: Weakly-Supervised Neural Network to Identify Eviction Filing Hotspots in the Absence of Court Records* (2022)

83 [Tabar, Jung, Yadav, Chavez, Flores, and Lee](#); *Forecasting the Number of Tenants At-Risk of Formal Eviction: A Machine Learning Approach to Inform Public Policy* (2022)

84 [Penn State University](#), *Artificial intelligence models aim to forecast eviction, promote renter rights* (2023)

85 [Displacement Alert Project](#), *About Us* (2025)

86 [Housing Data Coalition](#), *About Us* (2025)

87 [Eviction Lab](#), (2023)

88 [Eviction Lab](#), *About* (2025)

🔗 Risks

While early bets demonstrate promising emerging applications of AI and data-driven tools, most remain in pilot or proof-of-concept stages and have not yet been scaled or institutionalized.⁸⁹ Local governments and nonprofit partners often lack the data infrastructure, technical expertise, and institutional capacity needed to manage, validate, and continuously improve predictive models.⁹⁰ In many jurisdictions, data relevant to housing instability are fragmented across agencies, inconsistently updated, or unavailable in machine-readable formats, limiting model reliability and comparability.⁹¹ Privacy and ethical considerations, particularly regarding the use of sensitive administrative and location data, further constrain the extent to which AI-enabled prediction tools can be deployed.⁹² Weak data-sharing agreements between courts, landlords, and social service agencies also make it difficult to establish integrated systems that can both detect risk and enable timely intervention.⁹³ As a result, current AI-enabled eviction prediction efforts remain largely experimental and localized, with significant challenges to scaling in a safe, equitable, and evidence-based manner.

AI tools for eviction prevention and rental assistance have the potential to improve early detection and speed up aid delivery. However, without clear oversight and accountability frameworks, they risk compromising renter privacy, perpetuating bias in decision-making, and reinforcing existing inequities. Risks include:⁹⁴

- **Data privacy and consent:** Predictive models often depend on sensitive rental, income, and court data. Without robust privacy protections and clear consent mechanisms, such systems risk misuse or stigmatization of renters flagged as “high risk.”
- **Algorithmic bias:** Predictive targeting models may replicate historical inequities if trained on biased eviction data or incomplete records. Households with informal rental agreements, recent immigrants, or unreported income may be systematically underrepresented in models.
- **Administrative readiness:** Many local housing agencies lack data infrastructure, AI governance frameworks, or partnerships with researchers to audit model fairness and performance. To date, pilots remain limited, and scaling would require stronger investment and coordination among local governments, housing authorities, research institutions, and philanthropic or federal partners.

89 Expert interviews

90 Desmond, Gromis, Edmonds, Hendrickson, Krywokuski, Leung, and Porton; *Eviction Lab Methodology Report* (2022)

91 Cicero Institute, *Common Deficiencies in Homelessness Program Oversight Across Five U.S. Jurisdictions* (2025)

92 Expert interviews

93 National League of Cities, *The Data Cities Need to Understand & Address Their Eviction Crises* (2021)

94 Expert interviews

4. Reflections and Next Steps

As we look across credit markets, public benefits delivery, and housing, we are encouraged to see early potential to promote access, streamline processes, improve targeting, and personalize support in ways that could materially benefit LMI communities. It is important to note that not all AI applications should be treated alike. Lower-stakes uses, such as navigation, translation, reminders, or document sorting, may be appropriate earlier and at larger scale. Medium-stakes uses, such as caseworker copilots or triage tools, require stronger oversight and human review. Higher-stakes uses, such as credit decisions, eligibility determinations, or predictive targeting for scarce aid, warrant the strongest safeguards, including explainability, appeal pathways, bias testing, and clear human accountability.

In particular, three reflections emerge that may inform how ecosystem actors approach this agenda in the coming years:

1. **In the short term, the current focus seems to be on enhancing existing systems rather than creating new ones.** The three areas of priority AI use cases focus on improving public and financial infrastructures already in place, reducing friction in access, documentation, and decision-making for LMI users.
2. **The pace of innovation remains uneven across the ecosystem.** Tech-agile actors such as AI startups are advancing use cases across multiple domains of financial opportunity at a faster pace than traditional and smaller-scale institutions are currently adapting. This imbalance risks creating fragmented progress and, more critically, could concentrate analytical power and data access among a small number of technology vendors, leaving mission-driven institutions dependent on proprietary tools they cannot fully audit, adapt, or govern. There is an opportunity to strengthen diffusion and coordination to bridge this gap and promote responsible and inclusive innovation.
3. **Governance and data infrastructure that promote fairness, transparency, and consent are foundational to maintaining trust across initiatives.** While AI can enhance the precision and reach of financial opportunity efforts, its benefits will depend on being paired with sound governance, trusted intermediaries, and clear, practical standards for risk, safety, and fairness. Regulatory uncertainty and the lack of formal benchmarks for responsible AI adoption continue to slow progress, highlighting the need for concrete field signals and guardrails, such as equity and fairness metrics and recertification or retention indicators, that translate broad principles into actionable practice.

Building credible market signals and institutional capacity that enable responsible innovation to scale will be critical to realize this potential. Specifically, organizations and funders will need to focus on the following:

- **Data and infrastructure readiness:** Shared, machine-readable data standards can make program guidelines and regulatory criteria easier to operationalize and automate.
- **Consumer protection and accountability:** Clear standards for transparency, consent, and redress are critical to maintaining public trust and ensuring that AI tools enhance, rather than undermine, equitable access to financial opportunity.
- **Consumer readiness and digital confidence:** Building public awareness, digital literacy, and trust in AI-enabled tools will be essential to support informed uptake and ensure that solutions meet the needs and comfort levels of LMI users.

- **Flexible experimentation capital:** Social-sector innovators need dedicated experimentation capital—such as program-related investments, recoverable grants, or structured challenge funds—that allow for iterative learning and tolerate failure, mirroring how venture capital fuels private-sector discovery.
- **Measurements of success:** Establishing clear performance indicators, such as processing-time reduction, benefit-retention rates, model accuracy, and error correction, will help organizations compare approaches and track outcomes.
- **Evidence and independent evaluation:** As pilots proliferate, organizations and funders should distinguish between proof of concept, operational pilots, scaled implementation, and independently evaluated outcomes. Clear evidence standards will be essential to separate promising ideas from approaches that demonstrably improve access, equity, and long-term outcomes.

The impacts and potential of AI on economic mobility in the U.S. are both fast moving and continually evolving. Ensuring that developments and implementation are a net benefit for LMI communities will require continued engagement and dialogue with those communities served, as well as funders, practitioners, and AI developers to ensure these technologies are tested rigorously, governed responsibly, and deployed in ways that genuinely improve equity and opportunity.

www.dalberg.com

Dalberg

Serving those who
serve the world