北京大學
PEKING UNIVERSITY

No. 01
JUN. 2026

PDF Observer

Public Development Finance

Artificial Intelligence Empowerment of Public Development Financial Institutions

Public Development Finance Research Program at Peking University

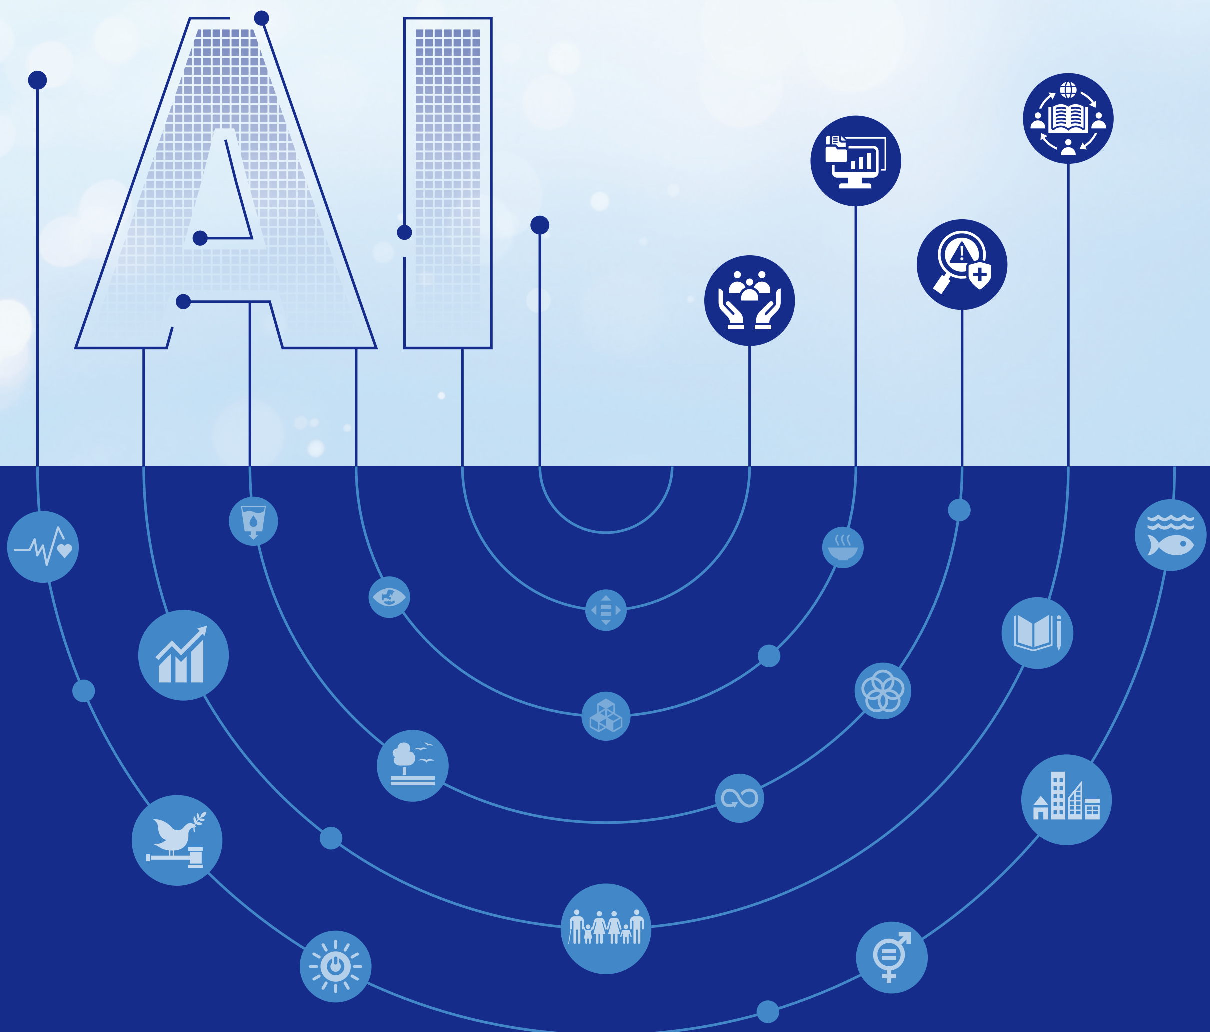

Endorsements

Xiaolan FU

Professor of Technology and International Development at the University of Oxford, Founder of OxValue.AI, a member of the United Nations High-level Advisory Board on Economic and Social Affairs

"This report systematically maps how PDFIs are deploying AI across four operational domains, balancing technological promise with sober governance analysis on human oversight and institutional learning. Highly recommended."

Rémy RIOUX

Chairman of Finance in Common

"This flagship report delivers a practice-grounded analysis of how public development banks can harness AI across core operational domains. An essential strategic roadmap for the Finance in Common community as it navigates the opportunities and governance challenges of AI adoption."

Lan XUE

Cheung Kong Chair Distinguished Professor at Tsinghua University, Chair of China's National Expert Committee on AI Governance

"As AI reshapes global governance, this pilot report examines concrete AI applications by PDFIs while squarely addressing algorithmic bias, transparency, and accountability. A foundational reference for policymakers and scholars in development finance."

Public Development Finance Observer Series is a flagship intellectual product initiated by the Public Development Finance Research Program at Peking University, designed for researchers, practitioners, and policymakers in the field of finance and development. It aims to synthesize the latest development of public development financial institutions, inspire peer practitioners by analyzing how solutions are piloted to solve common challenges, and equip readers with curated evidence and actionable insights.

The report was produced by the team of Public Development Finance Research Program (the "Program") at INSE and SHH at Peking University. The opinions expressed in this paper are those of the author(s) and the Program. It is therefore published under the sole responsibility of its author(s) and the Program.

The colors, boundaries, and other information on figures and tables do not reflect any judgements on the legal status of the territories.

Please send any comments or questions to the corresponding author, Jiajun Xu, at jiajunxu@pku.edu.cn.

Artificial Intelligence (AI) Usage Disclaimer

This report was prepared by the Public Development Finance Research Program at Peking University. Generative AI tools—including large language models—were used solely as assistive aids for literature retrieval, case summarization, and initial drafting of select sections. All AI-generated content was reviewed, fact-checked, and substantively revised by human researchers. The authors assume full responsibility for the accuracy, interpretation, and conclusions presented herein. No AI system was used to make analytical judgments, policy recommendations, or decisions about case selection or exclusion. Any errors or omissions remain solely those of the authors.

Contents

Acknowledgements

The inaugural issue of the Public Development Finance Observer (hereinafter "the PDF Observer") is authored by a research team from the Public Development Finance Research Program (hereinafter "the Program") at Peking University. The team is led by the Program's Principal Investigator, **Jiajun Xu** (Tenured Associate Professor at Peking University), who is responsible for conceptualization, methodology, review and verification of the writing, and supervision.

The full-time research team comprises **Shujun Han, Yixiao Huang, Yulu Tian, and Jiachen Xu**. Shujun Han was the task lead and bore the primary responsibility for collecting news data using digital technologies, as well as developing methodologies for identifying relevant news, with the support of Jiachen Xu. The writing of the manuscript was contributed by Shujun Han, Yulu Tian, and Yixiao Huang, each of whom drafted specific sections. This research project was coordinated by **Beina Xia**.

The completion of this issue would not have been possible without the support and contributions of Professor **Jinbing An** and Dr. **Yang Zou**, who provided technical guidance and support on the use of digital technology for data collection and analysis. Moreover, the team also wishes to express sincere gratitude to Professor **Xiaolan Fu** from the University of Oxford, Professor **Justin Yifu Lin** from Peking University, Mr. **Rémy Rioux** (Chairman of Finance in Common) and Professor **Lan Xue** from Tsinghua University for their valuable and constructive comments.

Abbreviations

ADBC	Agricultural Development Bank of China
AfDB	African Development Bank (Group)
AI	Artificial Intelligence
AVA	AI + Verified Analysis
BPF	Banco Português de Fomento
CDU	Crime Detection Unit
DCRM	Displacement Crisis Response Mechanism
DRC	Democratic Republic of the Congo
FHFA	Federal Housing Finance Agency [United States]
GPT	Generative Pre-trained Transformer
HF	Korea Housing Finance Corporation
IED	Institute for Economic Development [World Bank Group]
MDB	Multilateral Development Bank
ML	machine learning
NABARD	National Bank for Agriculture and Rural Development
NLP	natural language processing
OECD	Organization for Economic Co-operation and Development
PDF	portable document format
PDFI	Public Development Financial Institution
PMJDY	Pradhan Mantri Jan Dhan Yojana
RAG	retrieval-augmented generation
UNHCR	United Nations High Commissioner for Refugees

Executive Summary

This exploratory report examines how Public Development Financial Institutions (PDFIs) are beginning to adopt artificial intelligence (AI) across four operational domains: inclusive and tailored access, execution quality, proactive risk management, and integrated knowledge management and institutional learning. Unlike commercial banks that use AI primarily to maximize profits, PDFIs face distinct mandates to balance developmental impact with financial viability.

The report documents concrete applications across four domains. For inclusive and tailored access, National Bank for Agriculture and Rural Development uses multilingual AI voice agents to help rural low-literacy clients access banking services, and the African Development Bank's ENNOVA platform maps personalized entrepreneurial opportunities for youth and women. For execution quality, Freddie Mac and Korea Housing Finance Corporation employ AI scoring and data infrastructure to compress mortgage approval bottlenecks, while the Agricultural Development Bank of China's smart digital cloud counter connects physical and digital channels to reduce redundant manual input. For proactive risk management, Fannie Mae's AI-powered Crime Detection Unit scans for mortgage fraud before loans are finalized, while the World Bank has developed a refugee inflow forecasting model that predicts displacements four to six months in advance. For integrated knowledge management and institutional learning, Banco Português de Fomento's internal Fomento.AI assistant and the World Bank's AI + Verified Analysis platform transform fragmented information into usable knowledge and close the knowledge-to-action gap.

However, the evidence base remains thin. Most cases are descriptive system launches rather than evaluative impact studies, and many promising applications are still in pilot or pre-operational stages. The report's central thesis—that AI holds potential for empowering PDFIs to better fulfill public policy objectives—is promising but requires stronger empirical support from operational deployments with published metrics.

Despite the thin evidence, governance risks are already clear. AI models can replicate historical inequalities, operate as opaque "black boxes," and inadvertently exclude marginalized populations if poorly governed. The report therefore argues that PDFIs must move beyond isolated pilot projects led by IT departments and establish institution-wide governance frameworks. PDFIs must govern AI responsibly by investing in transparent data infrastructure and accountable vendors, building AI literacy across all staff, keeping humans accountable for consequential decisions while using AI only as decision support, testing high-risk applications in regulatory sandboxes before full deployment, and measuring the effectiveness in addressing market failures—not just efficiency.

Looking ahead, success hinges on whether PDFIs can govern AI with transparency, fairness, and human accountability. The technology is moving fast, but the organizational and governance work will determine whether AI widens or narrows inclusive access, strengthens or weakens accountability, and deepens or erodes trust in PDFIs. The Public Development Finance Observer Series will continue to track this domain and provide updated analysis as operational evidence matures.

I Introduction

1.1 Research Objective

Artificial intelligence (AI) has advanced rapidly and is now used across manufacturing, logistics, health care, public administration, and finance. Previous waves of IT adoption in banking focused on automation, electronic record-keeping, and improved connectivity (e.g., mobile banking), mainly lowering the cost of data storage and transmission. Recent literature highlights that AI, particularly machine learning (ML), introduces fundamentally different capabilities: it substantially reduces the cost of information processing, allowing financial institutions to process vast amounts of unstructured data (such as text and images) and capture complex, nonlinear relationships that traditional econometric models cannot (Agrawal et al. 2018; Financial Stability Board 2017; Mirestean et al. 2021). Empowered by these cognitive and predictive capabilities, AI in the financial sector now supports dynamic fraud detection, credit assessment, customer service, and regulatory compliance (Fuster et al. 2019; Rasheed et al. 2026). Yet AI deployment introduces ethical, technical, and institutional challenges—ranging from algorithmic bias and lack of transparency to data privacy concerns and regulatory fragmentation. These challenges are especially salient in public sector contexts, where trust, accountability, and equity are crucial (World Bank 2025a). As the Bali Fintech Agenda and Organization for Economic Co-operation and Development (OECD) Principles on AI affirm, digital innovation can advance inclusion and efficiency, but only when anchored in sound regulation, privacy protection, and financial stability safeguards (International Monetary Fund and World Bank Group 2018; OECD 2019).

Although a substantial number of researchers have examined AI adoption in commercial banking, its operational use by public development financial institutions (PDFIs) remains largely unexplored. Commercial banks typically operate in data-rich environments, utilizing AI to optimize retail credit scoring, predict short-term consumer behavior, or execute high-frequency trading based on vast historical datasets (Cao 2022; Chatrabhuj et al. 2025; Mention 2019; Rasheed et al. 2026). In contrast, PDFIs operate under distinct mandates—balancing developmental impact with financial viability in often data-poor environments—meaning their AI adoption trajectory may differ from that of commercial lenders. Yet evidence remains fragmented, dispersed across institutional news releases, project documents, and speeches rather than consolidated in a stand-alone report.

The present report fills that gap by mapping how PDFIs are beginning to deploy AI, not as a profit tool, but as an operational instrument to better fulfill public policy objectives. The objective is exploratory: to identify patterns of use, the problems AI is being deployed to solve, and the accompanying risks and governance challenges.

1.2 Data Sources and Justification for Case Selection

This report uses the definition of PDFIs that aligns with the seminal work by Xu et al. (2021). To be qualified as a PDFI, an entity needs to meet the following five qualification criteria: (1) being a stand-alone entity, (2) deploying fund-reflow-seeking financial instruments (such as loans, equity investments, and guarantees) as main products and services, (3) funding sources going beyond periodic budgetary transfers, (4) proactive public policy orientation, and (5) government steering of corporate strategy. In addition to public equity funds and public guarantee funds, public development banks (PDBs) are the major category of PDFIs. Apart from multilateral and subnational development banks, national development banks form the bulk of the PDB family. PDFIs play a crucial role in achieving sustainable development and structural transformation in both developed and developing countries. Together, they represent more than 550 institutions from 155 countries with an estimated US$23 trillion in assets.[1]

To ground our analysis, in this report we employ a three-step case selection methodology. Initially, the research team used a large language model to conduct an initial screening of all news disclosed on the official websites of PDFIs covering Q2 2025 to Q1 2026 to identify the relevant ones on AI deployment. This was followed by a manual validation process designed to filter out general digitalization announcements. Cases were retained if they demonstrated concrete AI deployment in internal functions, public-facing interfaces, or core operational workflows (such as risk and compliance). Ultimately, this process identified 20 institutions with clear, self-reported AI applications. They include 5 multilateral development banks and 15 national development banks from both developed countries and developing countries. Their official mandates range from flexible mandates to specific mandates (such as rural and agricultural development and social housing). We recognize that this methodology relies on publicly available information and may underrepresent institutions whose AI applications remain strictly internal, as well as those whose documentation is only available in local languages or less accessible formats. Finally, we classify AI deployment by PDFIs into four categories, namely, AI for inclusive and tailored access, operational execution, proactive risk management, and integrated knowledge management and institutional learning. We select representative cases with sufficient information available for case analysis in the present report.

1.3 Key Findings: What Challenges PDFIs Face and How AI Helps

PDFIs face four distinct challenges that AI is beginning to address.

First, PDFIs face a compelling challenge on how to turn formal service availability into practical and tailored usability for underserved populations, because marginalized applicants remain excluded by complex administrative procedures, language barriers, and complex service navigation. AI-powered multilingual conversational interfaces and personalized service guidance can help to enhance inclusive and tailored access.

Second, PDFIs grapple with administrative workflows and information processing systems. Unlike commercial

[1] Public Development Banks and Development Financing Institutions Database, initiated by PKU and co-founded by PKU and AFD with the collaborating institution of FERDI. Database DOI: https://doi.org/10.18170/DVN/VLG6SN. http://www.dfidatabase.pku.edu.cn.

lenders focused on standardized credit scoring, these institutions must process vast quantities of unstructured, nonstandardized documentation to ensure adherence to environmental and social standards, procurement integrity, and developmental impact. This document-heavy, labor-intensive workflow generates prolonged approval cycles and high transaction costs, which can be alleviated through intelligent document processing and automated verification engines. These technologies drastically reduce transaction costs and approval cycles while preserving public accountability and fiduciary oversight.

Third, PDFIs' risk-management architectures often remain constrained by a narrow focus on internal financial exposures, such as borrower credit default and interest rate volatility. These are inadequate for capturing compounding risks—such as climate-induced disasters, fiscal exposure, and social instability—where delayed risk signals can close the window for preventive intervention. By deploying predictive ML models and real-time geospatial analytics, PDFIs can synthesize diverse external data signals—such as localized satellite imaging and macroeconomic indicators—to transition from reactive mitigation to proactive, ex-ante anticipatory interventions before losses materialize.

Last but not least, PDFIs suffer from information fragmentation. Invaluable historical project evaluations, policy reports, and regulatory precedents remain trapped in decentralized departmental silos and unstructured formats. The inability to synthesize this information prevents institutions from converting raw data into actionable knowledge. This epistemic deficit not only inflates search costs but also impedes evidence-based policy formulation and institutional learning. Deploying unified AI-powered knowledge platforms that utilize semantic search and retrieval-augmented generation (RAG) can automatically index and link unstructured documents across different departments, allowing staff to instantly retrieve and synthesize historical project lessons.

Yet it is worth noting that AI deployment by PDFIs comes with potential risks and pitfalls. The reliance on opaque algorithmic models ("black boxes") may compromise the transparency and accountability required of public institutions. Moreover, in developing country contexts characterized by uneven digital infrastructure, AI models trained on unrepresentative data may inadvertently perpetuate historical biases or exacerbate digital divides (World Bank 2021).

The report proceeds as follows: The introduction establishes the research objective, methodology, and key findings of the four crucial roles of AI including inclusive and tailored access, operational execution, risk management, and knowledge management. The core of the report is organized into four analytical chapters (II–V), each dedicated to one of these four domains, providing a conceptual framework on challenges and solutions followed by illustrative case studies. The conclusion synthesizes the findings, collates cross-cutting governance challenges, and offers policy implications for PDFIs seeking to harness AI to better fulfill public policy objectives.

II

AI for Enhancing Inclusive and Tailored Access

In this section we examine how AI can enhance inclusive and tailored access by reducing barriers for underserved populations and providing personalized support for youth and women to foster innovation, illustrated through cases from National Bank for Agriculture and Rural Development (NABARD) and the African Development Bank (AfDB).

2.1 The Nature of Client Access and Service Accessibility Challenges

Problems arise when services formally exist but remain difficult to discover, navigate, or use for rural households, small-town residents, low-income clients, farmers, informal workers, first-time borrowers, or other underserved groups. Two closely related challenge areas define this problem. First, client access in the narrow sense: users face practical barriers such as distance from branches, weak connectivity, low digital literacy, or language constraints. Second, service accessibility after entry: users struggle with complex conditions, fragmented information, or unclear next steps even once contact with the institution has been established.

2.1.1 Access Barriers in Underserved, Rural, and Small-Town Markets

The first challenge concerns whether underserved, rural, and small-town clients can obtain services through channels they can actually use. In many PDFIs, the problem is not that financial services are entirely absent, but that conventional delivery channels remain difficult to access for the users the institution is mandated to serve. Rural households, nonliterate users, informal workers, small farmers, low-income families, and first-time borrowers may face unstable connectivity, limited digital skills, weak documentation, or long distance from physical branches. As a result, a service may formally exist while remaining practically out of reach. Therefore, the challenge is about the last-mile inclusion rather than mere channel provision.

2.1.2 Lack of Intelligible and Personalized Service Guidance

The second challenge concerns whether clients and the wider public can obtain continuous, clear, and personalized service guidance once they enter the service environment. Even where initial access barriers are partly reduced, users may still struggle if they encounter fragmented information, complex product conditions, limited language support, or unclear next steps. PDFI services often involve intricate rules on credit products, eligibility criteria, and application processes that are difficult to understand through static webpages, conventional FAQs, or limited-hour hotlines.

2.2 Main Pathways of AI's Support for Inclusive and Tailored Access

AI supports inclusive and tailored access in PDFIs through two main pathways.

First, access barriers can be reduced through voice, language, hybrid physical–digital channels, and public-channel assistants. AI-enabled tools such as voice-based interfaces in local languages, conversational assistants embedded in messaging or call-center channels, and automated document-capture and verification systems matter when they reduce these frictions and make existing services more discoverable, more intelligible, and more usable for underserved users (Galileo Financial Technologies 2025; Sirtaine 2025).

Second, service-accessibility barriers can be reduced through personalized guidance tools that explain conditions and accompany users through complex processes. AI helps translate complex information into personalized, intelligible, and context-specific guidance. AI contributes to tailored access when it helps users understand available services in terms they can relate to, and move through service processes with clearer explanations, fairer assessment, and more transparent avenues for recourse where problems arise (Shetty 2025; World Bank 2021).

2.3 Case Analysis

2.3.1 Voice-Based Access for Rural Low-Literacy Users at NABARD

Case Summary

NABARD, India's national development bank for agriculture and rural development, operates primarily as a wholesale lender channeling funds through cooperative and regional rural banks. Even under India's national financial inclusion mission, many rural accounts remain inactive due to language, literacy, and procedural barriers.

NABARD has deployed multilingual AI voice agents at the cooperative–bank interface, allowing rural clients to check balances, receive alerts, and initiate basic transactions by speaking in local languages. This lowers the procedural distance for users who are physically close to branches but textually distant from digital channels.

NABARD is India's development bank tasked with fostering rural prosperity. To enhance financial inclusion, the Indian government launched a national mission—Pradhan Mantri Jan Dhan Yojana (PMJDY)—to ensure universal access to basic savings and deposit accounts, remittances, credit, insurance, and pensions in an affordable manner.

Even under PMJDY, however, many rural and semi-urban accounts remain inactive, and financial and cyber awareness is limited, indicating that formal account opening has not fully translated into practical service use (Department of Financial Services, Ministry of Finance, Government of India n.d.). Within this architecture, cooperative and regional rural banks function as the main local branches through which PMJDY account holders in rural areas actually interact with the banking system, but many of these users still face language, literacy, and procedural barriers that limit the active use of their accounts.

NABARD operates primarily as a wholesale national development bank funded by central-government capital contributions, market borrowings, institutional deposits, and international credit lines, rather than by retail household savings. Retail outreach is largely channeled through cooperative and regional rural banks, which act as local financial intermediaries serving rural clients on the basis of NABARD's refinancing and on-lending operations (NABARD 2024).

NABARD's deployment of multilingual AI voice agents in cooperative banks directly addresses the access barriers. Rural and nonliterate clients can interact with the banking system by speaking in local languages rather than navigating text-heavy interfaces. These AI agents are intended to allow customers to check balances, receive alerts, and initiate basic transactions or loan enquiries through natural-language interaction linked to PMJDY accounts, thereby turning branch-centered access into a more conversational channel.

The voice-based AI tools mentioned in this case are implemented by the responsible institution within the broader policy and infrastructure frameworks promoted by the Government of India. NABARD is piloting a voice-based AI assistant to deliver information and services to rural clients, building on India's broader AI and digital public infrastructure initiatives, as shown in a short video from the AI and AgriTech Innovation Center at the AI4Agri 2026 event featuring NABARD's chairman (AI and AgriTech Innovation Center 2026). The broader significance of the case is that it demonstrates how AI can reconfigure the interaction layer itself. By lowering language and literacy barriers for users who are physically close to cooperative institutions but procedurally distant from text-based digital channels, the system helps turn nominal inclusion into more meaningful participation in public finance. In this sense, NABARD illustrates how AI-enabled voice interfaces can expand last-mile access for priority rural users.

2.3.2 ENNOVA Platform and AI-Supported Opportunity Mapping for Youth and Women Entrepreneurs at AfDB

Case Summary

The AfDB's ENNOVA platform uses AI to aggregate and curate entrepreneurial resources—over 1,400 enterprise support organizations, 300+ business resources, mentorship tools, and a job portal—across Africa. The platform is designed to lower informational and coordination barriers for youth and women entrepreneurs, helping them translate diffuse information into clearer pathways for skills, finance, and market access.

As a regional development bank aiming to build a "prosperous, inclusive, resilient and integrated" Africa (African Development Bank Group 2025a), AfDB has framed youth employment and women's economic empowerment as central to Africa's structural transformation, given the continent's rapidly expanding labor force and persistent shortages of decent jobs. Within AfDB's Jobs for Youth in Africa agenda (African Development Bank Group 2016) and wider structural-transformation priorities, the Bank's Innovation and Entrepreneurship Lab developed ENNOVA as an AI-powered knowledge and exchange platform for Africa's entrepreneurial ecosystem with a special focus on empowering youth and women. ENNOVA operates as a one-stop hub connecting entrepreneurs, enterprise support organizations, investors, and other ecosystem actors who often lack reliable information and networks (African Development Bank Group 2025b).

ENNOVA is designed to address the service-accessibility and information-navigation challenges. The platform aggregates and curates content through AI-supported knowledge functions, including a directory of more than 1,400 enterprise support organizations, over 300 business resources on funding and training, events and news feeds, mentorship and networking tools, a job portal, and a document library of AfDB and partner knowledge products. AfDB communications emphasize that the platform aims in particular to equip youth and women entrepreneurs with tailored opportunities, information, and connections that they might otherwise struggle to obtain, thereby lowering informational and coordination barriers to entering innovation and venture pipelines.

The ENNOVA case illustrates how AI can operate at the level of opportunity mapping and ecosystem navigation. By using AI to organize, recommend, and surface relevant resources and contacts across a fragmented continental landscape, ENNOVA helps young and women entrepreneurs translate diffuse information into clearer pathways for skills, finance, and market access.

2.4 Takeaways from Case Analysis

The two case studies demonstrate that AI enhances inclusive access not by replacing human judgment, but by acting as an intelligible bridge between underserved users and complex financial systems. First, AI lowers access barriers by meeting users where they are: NABARD's multilingual voice agents overcome literacy and language obstacles, allowing rural clients to interact with banks through natural conversation rather than text-heavy interfaces—a clear example of making services more discoverable and usable. Second, AI improves service accessibility by providing personalized guidance through fragmented information landscapes: AfDB's ENNOVA platform curates and maps entrepreneurial resources, helping youth and women entrepreneurs translate diffuse opportunities into actionable pathways. Critically, neither system removes the human from the loop. The voice agent handles routine transactions but would escalate complex requests (like loan approvals) to human staff, whereas ENNOVA informs but does not automate final decisions. The overarching takeaway is that effective AI for inclusion operates at the interaction layer—raising intelligibility, reducing procedural distance, and knowing when to step back and let a human take over.

It is worth noting that the value of user-facing AI is contingent upon robust governance frameworks. Technological convenience, multilingual interfaces, and conversational capabilities do not inherently guarantee equitable outcomes. Voice agents and entrepreneurial platforms must preserve intelligibility, fairness, options for human escalation (i.e., when a computer or automated system cannot solve a problem, it hands the problem over to a real person), and meaningful recourse if automated interaction proves insufficient. In the absence of bias-mitigation measures, human-in-the-loop oversight, and accessibility complaints mechanisms, automated systems may exacerbate exclusionary practices while shielding discriminatory outcomes from institutional scrutiny. For PDFIs, the relevant policy question is therefore not how quickly to scale up user-facing AI, but whether particular deployments expand practical inclusive access in ways consistent with user protection and the institution's development mandate.

III

AI for Improving Operational Execution

In this section we analyze AI's role in improving operational execution, drawing on examples from Freddie Mac, Korea Housing Finance Corporation (HF), and the Agricultural Development Bank of China (ADBC).

3.1 Operational Inefficiency and Challenges in PDFIs

Operational execution refers to the ability to process applications, verify information, coordinate workflows, manage data, and execute decisions in a timely, consistent, and accountable manner. In many institutions, delivery remains constrained by paper-based files, manual verification routines, long approval chains, and costly verification of collateral assets.

3.1.1 Burdensome Document Processing and Collateral Assets Verification

The frontline operations of PDFIs are heavily constrained by the manual and labor-intensive nature of document processing and physical asset verification. PDFIs routinely handle vast volumes of unstructured and dense documentation, ranging from complex financial statements and land-use certificates to regulatory filings. Concurrently, verifying collateral—such as agricultural machinery, leased equipment, or municipal infrastructure—demands physical inspections and site visits. Because these verification tasks rely primarily on manual review and physical verification, this input stage becomes a severe operational bottleneck, driving up transaction costs, extending processing timelines, and introducing high risks of human error and inconsistency across geographically dispersed operations.

3.1.2 Process Fragmentation and Decision Latency

The second challenge lies in the prolonged and opaque approval workflows within credit, leasing, and specialized lending operations. In many PDFIs, processing an application requires sequential and multilayered reviews: document completion, risk assessment, policy-eligibility checks, collateral review, authorization, and internal sign-off. These layered procedures result in increased processing times, diminished workflow visibility, and inflated operational costs for both institutions and clients. In operational terms, case handlers lack real-time visibility into application status and pending actions, a problem exacerbated when files are routed across different departments or informal communication channels. The core challenge is to preserve essential prudential regulation while eliminating redundant loops, accountability blind spots, and unnecessary delays in workflow.

3.2 Main Pathways of AI's Support for Improving Operational Execution

AI can enhance operational execution in PDFIs through two main pathways that correspond directly to the challenges identified previously.

First, document-processing and asset-monitoring tools—such as optical character recognition, intelligent document processing, computer-vision models, and automated valuation systems—reduce the amount of repetitive manual work required to verify documents and check collateral or financed assets. These tools keep evidence trails structured and traceable, while allowing experts to review and, where necessary, correct or override AI-generated outputs. They shift execution from repetitive manual file handling and low-frequency site visits to more structured, traceable workflows, while retaining expert validation for high-risk cases.

Second, workflow-oriented tools—including AI-enabled scoring and pre-assessment models, automated underwriting support, and unified workflow platforms—help convert fragmented, multistep approval chains into more visible and coordinated processes. In these systems, case status, pending actions, and responsibilities become easier to track across units. This moves execution from fragmented approval chains to more visible and timely processes, while preserving human responsibility for consequential decisions and maintaining process transparency.

3.3 Case Analysis

3.3.1 AI-Enabled Execution in Housing-Finance Institutions: Freddie Mac and HF

Case Summary

Freddie Mac and HF both operate under public-policy mandates to provide liquidity and stability in housing finance. Freddie Mac uses AI scoring and pre-assessment models (Loan Product Advisor, Income Calculator, Asset and Income Modeler) to compress high-friction decision nodes within mortgage approval chains, shortening cycles while leaving final responsibility with human underwriters. HF's AX Strategy 2026–2028 focuses on building a big-data platform that integrates housing-finance data accumulated since its establishment—including mortgage loans, mortgage-backed securities, and housing guarantees—with the stated goal of reducing data-access time by up to 80 percent, and on deploying AI agents for automated document checks and case routing (Korea Housing Finance Corporation 2025).

Freddie Mac and HF occupy central positions in their respective housing-finance systems while operating under explicit public-policy mandates. Freddie Mac's stated mission is "to provide liquidity, stability and affordability to the U.S. housing market in all economic environments" (Freddie Mac n.d.). Under its founding legislation, HF's purpose is to "contribute to the promotion of national welfare and the development of the national economy by facilitating the supply of long-term and stable housing funds" (Republic of Korea 2003, art. 1). Both institutions thus combine market-facing operations with responsibility for the stability and accessibility of housing finance.

In operational terms, both institutions face execution frictions that are characteristic of housing-finance and mortgage-related activities. Freddie Mac must support lenders in processing large volumes of mortgage applications that involve heterogeneous income types, complex documentation, and prudential eligibility checks. HF manages extensive datasets on mortgage loans, mortgage-backed securities, and housing guarantees accumulated over time, and needs to make these data usable across customer service, risk management, and internal decision-support functions.

Freddie Mac uses AI to compress specific high-friction decision nodes within mortgage approval chains instead of redesigning entire workflows. Its Loan Product Advisor is an automated underwriting system that assesses whether a loan is eligible for purchase by Freddie Mac and provides lenders with structured feedback on documentation completeness, risk characteristics, and required follow-up actions. Complementary tools such as the Income Calculator and Asset and Income Modeler standardize the assessment of borrower income, assets, and employment, including nonstandard and self-employed income categories.

Functionally, these AI scoring and pre-assessment models, automated underwriting support, and end-to-end digital processing manage to shorten approval cycles by reducing repeated manual recalculation of income and collateral-related data, make decision criteria more transparent to lenders, and support more consistent application of underwriting rules across large volumes of loans, while leaving final responsibility for consequential decisions with human underwriters. This shows how AI can address execution problems in the approval chain without weakening prudential review.

HF’s AX (AI Transformation) Strategy 2026–2028 tackles execution frictions primarily through data-infrastructure reform. The strategy prioritizes the construction of a big-data platform that integrates housing-finance data accumulated since HF’s establishment—including mortgage loans, mortgage-backed securities, and housing guarantees—with the expectation of reducing data-access time by up to 80 percent. On top of this platform, HF plans to upgrade its scenario-based chatbot into a generative-AI-based intelligent chatbot, establish an AI Contact Center capable of 24/7 customer consultation, enhance its in-house developed generative AI model, labelled as HFGPT, and deploy AI agents to automate repetitive internal tasks such as standard document checks and simple case routing (Korea Housing Finance Corporation 2025).

These initiatives reorganize historical and operational data into a coherent evidence base, aiming to make it easier for staff to retrieve the historical record of mortgage loans, mortgage-backed securities, and housing guarantees across systems, monitor risk indicators, and respond to customer queries using consistent information. In this way, data infrastructure becomes a direct component of execution capacity rather than a background IT concern.

Taken together, Freddie Mac and HF show how housing-finance institutions use AI to address execution bottlenecks at different points in the mortgage and guarantee life cycle. Freddie Mac focuses on compressing high-friction assessment nodes inside approval chains, while HF builds a data foundation that enables faster retrieval, monitoring, and service handling across products and channels. Both approaches illustrate that AI creates operational value when it reduces friction at specific points of evidence processing and workflow coordination, rather than simply adding new technical features.

3.3.2 Smart Counter Integration at ADBC

Case Summary

The ADBC, China's policy bank for agriculture and rural development, launched a "smart digital cloud counter" built on distributed architecture using cloud computing and large AI models. The system connects on-site counter operations with remote authorization, voice interaction, intelligent Q&A, and digital-human functions across physical counters, self-service terminals, mobile banking, and internet banking. This allows an instruction started in one channel to be continued in another without reentering data, reducing redundant manual input and making approval steps more visible.

The ADBC is China's policy bank for supporting agriculture, rural development, and related strategic financing priorities. Its main task is to play a pivotal role in the rural financial system by relying on the support of national credit to support key areas and weak links in agriculture and rural areas and promote sustained and healthy economic and social development (Agricultural Development Bank of China n.d.).

In its branch and counter operations, ADBC must process large volumes of routine instructions through branch counters, remote approvals, self-service terminals, mobile banking, and internet banking. Before the smart-counter initiative, many of these instructions were handled mainly through manual counter work. Staff at physical counters had to enter transaction data, check supporting documents, request remote approvals for higher-risk or higher-value transactions, and update core-banking systems, often using different interfaces for different product lines. Self-service and online channels were only loosely connected to these processes, so customer instructions could be duplicated or reentered, and staff in different locations could not always see the same status of a case in real time. Processing a single instruction—such as changing repayment terms, registering a guarantee, or drawing down a loan—could therefore involve repeated data entry, several phone calls for authorization, and additional counter visits by the client.

To reduce this fragmentation, ADBC launched its "smart digital cloud counter" (Li Yan and Deng Guoping 2025) as a new generation of counter-operation and service system built on a distributed architecture using cloud computing and large AI models. The system provides a single platform on which counter staff, remote approvers, and digital channels work with the same transaction record and case status. It connects on-site counter operations with remote authorization, voice interaction, intelligent Q&A, and digital-human functions, and links physical counters, self-service terminals, mobile banking, and internet banking so that an instruction started in one channel can be continued in another without reentering all the data. In practical terms, when a customer submits an instruction through a self-service device or mobile app, the relevant fields are prefilled for counter staff, and remote approvers can review the same record directly in the system instead of relying on separate emails or calls.

The smart digital cloud counter changes concrete execution steps rather than only adding a new frontend tool. It cuts redundant manual input, reduces the number of times a case is passed between counters and back-office units, and makes it easier to use the same eligibility and verification rules across channels. Identity-verification and product-eligibility checks are encoded as machine-readable rules in the system interface, so staff receive explicit prompts when information is missing or inconsistent instead of relying on memory of internal manuals. AI-enabled Q&A and digital–human functions answer common questions and guide customers through standard procedures, which frees counter staff to spend more time on exceptions and higher-risk transactions. Remote authorization steps are automatically recorded in system logs, so it is clear who approved which operation, at what time, and on the basis of which data.

3.4 Takeaways from Case Analysis

Based on the analysis of Freddie Mac, HF, and ADBC, the key takeaway is that AI improves operational execution in PDFIs by reducing friction at specific bottlenecks rather than fully automating human judgment. Across all three cases, AI addresses two core operational challenges: burdensome document and collateral verification, and fragmented, latency-prone approval workflows. Freddie Mac's AI scoring tools (Loan Product Advisor, Income Calculator) compress high-friction decision nodes within mortgage approvals by standardizing income and asset assessment, while leaving final underwriting responsibility with human experts. HF's AI transformation strategy prioritizes building a unified big-data platform to cut data-access time by up to 80 percent, upon which AI agents handle routine document checks and case routing—demonstrating that data infrastructure is a prerequisite for execution efficiency. ADBC's smart digital cloud counter connects physical counters, remote authorization, self-service terminals, and mobile banking into a single platform, allowing an instruction started in one channel to continue in another without redundant data entry, while making approval steps visible and traceable.

The common thread is that AI creates operational value when it reduces repetitive manual work, makes workflows visible and coordinated, and preserves human accountability for consequential decisions—shifting execution from fragmented, paper-heavy processes to structured, traceable, and channel-integrated operations.

IV

AI for Proactive Risk Management

In this section we explore AI applications for proactive risk management, highlighting fraud detection at Fannie Mae and refugee inflow forecasting at the World Bank.

4.1 Risk Management Challenges for PDFIs

Because socioeconomic costs compound exponentially under reactive responses, PDFIs must pivot to anticipatory action enabled by early warning systems. Compounding this challenge, the conventional reliance on siloed databases and static appraisals neglects cross-sectoral interactions. By failing to synthesize real-time signals, risks may be mispriced.

4.1.1 Risks with Wider Public Consequences Require Earlier Action

PDFIs encounter distinct risks. Whereas private entities mitigate idiosyncratic portfolio losses, PDFIs operate to address market failures and mitigate risks with profound socioeconomic externalities. If institutions react only after vulnerabilities materialize, socioeconomic costs compound exponentially. Thus, PDFIs should pivot from reactive shock-absorption to anticipatory action. Early warning systems are indispensable: though they cannot eliminate systemic risks, they secure the critical operational lead time required to intervene before developmental outcomes deteriorate.

4.1.2 Risk Signals Remain Fragmented

Global South nations are increasingly exposed to what some analysts call a “polycrisis”—a convergence of climate change, geopolitical volatility, supply chain disruptions, and food insecurity that manifests as compound risks (Tooze 2022). However, the traditional diagnostic tool kits employed by PDFIs are underequipped to capture these complex circumstances. The core vulnerability of these frameworks lies in their reliance on siloed databases and static, manual project appraisals, which make it difficult to capture how systemic risks emerge from interactions across sectors and spread through interconnected systems.

When Multilateral Development Banks (MDBs) conduct project appraisals through traditional manual reviews of isolated loan files, they overlook the correlation matrices between these variables. A project may appear fiscally viable on a stand-alone balance sheet, yet its actual risk of default rises exponentially when a localized drought (climate risk) triggers food price inflation (macroeconomic volatility), which in turn sparks civil unrest (social risk), ultimately disrupting project operations and cash flows. Consequently, by failing to synthesize these cross-sectoral signals in real time, risks may be mispriced.

4.2 Main Pathways of AI’s Support for Proactive Risk Management

Three AI functions are especially relevant for PDFIs to make proactive risk management.

First, integrating dispersed data. AI can make risk signals visible by combining loan documents, transactions, climate indicators, conflict records, spatial data, and text-based information. This moves institutions from manual file-by-file review to risk-based prioritization, with AI outputs treated as investigative leads rather than final decisions.

Second, detecting unusual patterns. ML can identify anomalies across large datasets and help teams prioritize investigations or resources. This shifts work from reactive case handling to proactive screening.

Third, embedding early warnings into workflows. Forecasting models can widen the time available for action when they connect to procurement, budgeting, and disbursement processes. This enables earlier preparation while maintaining formal approval procedures.

4.3 Case Analysis

4.3.1 Proactive Detecting Mortgage Fraud at Fannie Mae

Case Summary

Fannie Mae, a government-sponsored enterprise supporting US housing finance, launched an AI-powered Crime Detection Unit (CDU) in May 2025 in partnership with Palantir Technologies. The CDU integrates disparate datasets—transaction records, property valuations, loan application files, behavioral indicators—into a unified data science platform that scans for anomalous patterns before mortgages are finalized. In a benchmark test involving four historical loan files, the system identified relevant anomalies in approximately 10 seconds—a process estimated to require two months of manual review. The CDU is designed as a "human-in-the-loop" system: AI flags anomalies, but human investigators decide whether to act.

Fannie Mae, established in 1938, is a government-sponsored enterprise in the United States housing finance system. It does not lend directly to homebuyers; instead, it buys mortgage loans from lenders, packages them into mortgage-backed securities, and sells them to investors, thereby supporting market liquidity and stability. Fannie Mae qualifies as a PDFI because it fulfills a clear public policy mandate: supporting liquidity, stability, and affordability in the US housing finance system through the purchase and securitization of mortgage loans. Although it operates through a corporate legal structure, its strategic direction and risk exposure remain subject to substantive government steering under the conservatorship and regulatory oversight of the Federal Housing Finance Agency (Fannie Mae n.d.a; Federal Housing Finance Agency 2022).

Fannie Mae is highly exposed to credit and operational risks stemming from mortgage fraud. Historically, Fannie Mae's fraud mitigation relied on reactive, ex-post controls, including lender self-reporting and manual forensic audits by its Financial Crimes Team. Although Fannie Mae initiated localized AI and ML pilots for property valuation and NLP between 2019 and 2021, fraud detection remained bottlenecked by data silos and the labor-intensive nature of identifying nonlinear anomalies across massive, multisource document pools (Fannie Mae n.d.b; Federal Housing Finance Agency Office of Inspector General 2022).

Fannie Mae deployed a platform-based risk intelligence framework by partnering with Palantir Technologies to launch the AI-powered CDU on May 28, 2025 (Fannie Mae 2025). The CDU integrates disparate datasets—including transaction records, property valuations, loan application files, and behavioral indicators—into a unified data science platform. By scanning for anomalous patterns before mortgages are finalized or integrated into Fannie Mae's portfolio, the system shifts the institution's risk management posture from reactive, postpurchase investigation to ex-ante risk prevention. The CDU is designed as a "human-in-the-loop" decision support system (Bracken 2025; Fannie Mae 2025). The AI model does not make binding legal or credit decisions; instead, it discerns anomalies to trigger human investigation. During an initial benchmark test involving four historical loan files, the Palantir-enabled system identified relevant anomalies in approximately 10 seconds—a process estimated to require two months of manual human review. Although this test indicates substantial efficiency gains in document comparison, the platform's performance at scale remains unproven. Peer-reviewed evidence regarding its long-term false-positive rates (how often legitimate cases are incorrectly flagged as suspicious), precision-recall metrics (whether the flagged cases are accurate and whether the system detects the fraud cases that are actually present), and overall reduction in credit losses has not yet been publicly disclosed.

In summary, an AI-supported platform can search for hidden relationships across large datasets and flag priority cases before losses enter the portfolio. Investigators remain central: AI can direct them to files, patterns, and cases that warrant closer review, but human investigators must decide whether a signal justifies further action. The Federal Housing Finance Agency's (FHFA) Advisory Bulletin 2022-02, revised in May 2025, allows AI and ML to support business and operational functions but requires institutions to manage model risk, bias risk, data risk, information security, privacy, third-party vendor risk, and business continuity risk (Federal Housing Finance Agency 2025).

4.3.2 Forecasting Refugee Inflows at the World Bank

Case Summary

The World Bank developed an AI-powered refugee forecasting model to predict registered refugee inflows from South Sudan and the Democratic Republic of Congo into Uganda four to six months in advance. The model analyzes more than 90 variables across conflict events, food prices, climate conditions, news, and social media. Validation on historical data (2014–2023) achieved predictive accuracy above 80 percent in terms of absolute percentage error: 9.2 percent for the Democratic Republic of the Congo (DRC) (120-day horizon) and 12.16 percent for South Sudan (180-day horizon) at an 80 percent confidence level. As of March 2026, the model was not yet operational; the next stage involves near-real-time forecasting and integration with Uganda's Displacement Crisis Response Mechanism.

The World Bank is a leading multilateral development bank that provides financing, policy advice, technical support, and knowledge services to developing countries, in line with its mission to end extreme poverty and boost shared prosperity on a livable planet (World Bank Group n.d.). Forced displacement is not only a humanitarian challenge but also a long-term development challenge. Large refugee inflows can place sustained pressure on schools, health facilities, water systems, local budgets, and relations between refugees and host communities. The World Bank developed an AI-powered refugee forecasting model that analyzes conflict, economic, climate, food-price, spatial, news, and social-media variables to forecast inflows four to six months ahead, supporting service planning and financing preparation.

Uganda has long received refugees from South Sudan and the DRC. Its relatively open refugee policy allows many refugees, under certain conditions, to access land, employment opportunities, and public services, so refugees and host communities share schools, clinics, water systems, and local infrastructure (Mearns et al. 2025). This creates pressure when refugee inflows rise quickly: classrooms become crowded, clinics approach capacity, and water points serve more users. Because classrooms, health facilities, and water points require budgeting, procurement, construction, and staffing, a response that starts only after arrivals can reduce service quality and increase competition between refugees and host communities.

The World Bank supported Uganda in establishing the Displacement Crisis Response Mechanism (DCRM), which provides prearranged funding to refugee-hosting areas when population inflows place pressure on education, health, and water services. DCRM reduces coordination costs because funding pools, trigger conditions, and eligible uses are defined before a crisis; however, it remains mainly responsive because it releases funds when predefined thresholds are crossed (Mearns et al. 2025). The AI model adds the missing anticipatory layer: thresholds detect pressure after it emerges, whereas predictive signals seek to identify pressure before it fully materializes.

The World Bank's AI-powered refugee forecasting model focuses on registered refugee inflows from South Sudan and the DRC into Uganda and forecasts movements four to six months in advance (World Bank 2025b). It analyzes registered refugee movements from 2014 to 2023 using arrival data maintained by the Government of Uganda and the United Nations High Commissioner for Refugees (UNHCR), together with more than 90 variables across conflict events, food prices, climate conditions, news, and social media (Structured Data Analytics 2024; Structured Data Analytics 2025).

The model combines objective conditions and public perceptions. Conflict events, food prices, and climate conditions may shape migration decisions, whereas language patterns in news and social media can indicate how people perceive those changes. By combining these signals, the model generates risk information for policy decisions, but it does not release funds automatically. AI moves the risk-management window forward: when governments can anticipate likely pressure four to six months before arrivals, local authorities gain time to prepare water points, expand health facilities, and construct classrooms (World Bank 2025b).

Validation results show practical promise but require caution. The model exceeded 80 percent predictive accuracy on data not used during training (Mearns et al. 2025), and the validation also tested six forecast horizons—30, 60, 90, 120, 150, and 180 days—using absolute percentage error to compare predicted and observed registered arrivals. The best-performing DRC forecast used a 120-day horizon and had an APE of 9.2 percent; the best-performing South Sudan forecast used a 180-day horizon and had an APE of 12.16 percent, both reported at an 80 percent confidence level (Structured Data Analytics 2025). These metrics indicate potential, but they do not prove that the model will remain reliable over time or capture every sudden conflict, migration route, or host-area change.

As of March 2026, operationalization remained ongoing. The next stage includes updating recent data, running the model in near-real time, generating forecasts every two weeks, and sharing results through an online platform for anticipatory action (World Bank-UNHCR Joint Data Center on Forced Displacement 2026). Operational use also requires capacity building because governments, development institutions, and humanitarian organizations must translate forecasts into budgeting, procurement, service planning, and resource allocation. The project is therefore an AI-supported forecasting model moving toward near-real-time use, not a fully deployed automated system.

The case also shows the limits of forecasting. The model focuses on registered refugee inflows, so it cannot fully capture unregistered movements, internally displaced persons, or secondary migration; it also focuses on South Sudan and the DRC, even though conflict patterns, routes, and countries of origin may change (Structured Data Analytics 2025). News and social media can complement traditional statistics, but their value depends on internet access, multilingual processing quality, media bias, and possible manipulation. Finally, model variables identify relationships rather than definitive causal mechanisms. Earlier forecasts can widen the intervention window, but they cannot remove implementation bottlenecks unless institutions connect them to prearranged financing, project execution, and local capacity.

4.4 Takeaways from Case Analysis

Based on the analysis of Fannie Mae's AI-powered fraud detection unit and the World Bank's refugee inflow forecasting model, three key takeaways emerge for proactive risk management in PDFIs.

First, AI enables a shift from reactive to anticipatory risk management by integrating disparate, siloed data sources—transaction records, property valuations, conflict events, food prices, climate conditions, and even news and social media—into unified platforms that can detect emerging threats before they fully materialize. Fannie Mae's CDU scans for anomalous mortgage patterns before loans are finalized, whereas the World Bank's model forecasts refugee inflows four to six months in advance, widening the intervention window for service planning and budget allocation.

Second, both cases reinforce the human-in-the-loop principle established earlier: AI serves as a decision-support tool that flags anomalies or generates probabilistic forecasts, but human investigators and policymakers retain final authority over consequential actions. Fannie Mae explicitly designs its system to trigger human investigation rather than automate binding decisions, and the World Bank's model informs but does not automatically disburse funds.

Third, proactive risk management requires both technological capability and operational integration—forecasts only create value when connected to prearranged financing, procurement workflows, and local response capacity. The World Bank's model complements Uganda's existing DCRM by adding an anticipatory layer to threshold-based funding, whereas Fannie Mae's system embeds early warnings directly into mortgage purchase workflows.

The overarching takeaway is that AI strengthens proactive risk management when it synthesizes cross-sectoral signals into actionable lead time, treats AI outputs as investigative leads rather than final verdicts, and anchors early warnings within existing operational and budgetary processes.

V

AI for Integrated Knowledge Management and Institutional Learning

In this section we investigate how AI can transform knowledge management and institutional learning, featuring the internal assistant Fomento.AI at Banco Português de Fomento (BPF) and the public-facing AI + Verified Analysis (AVA) platform at the World Bank.

5.1 Knowledge Management and Institutional Learning Challenges for PDFIs

5.1.1 Abundant Information Does Not Readily Turn into Usable Knowledge

PDFIs operate in information-rich environments, managing vast repositories of operational manuals, regulatory frameworks, and historical project data. However, these assets frequently suffer from fragmentation across decentralized, non-interoperable databases. This systematic disconnect hampers mobilization of explicit knowledge. The lack of a unified, authoritative entry point generates significant cognitive and temporal search costs for staff, leading to version-control inconsistencies and operational redundancies. To be institutionally "usable," information must transcend passive storage; it must be dynamically retrievable, verifiable, contextually interpretable, and seamlessly integrated into project appraisal and execution workflows.

5.1.2 Static Knowledge Products Are Hard to Turn into Timely Action

Although MDBs act as "knowledge banks" generating extensive flagship reports, country diagnostics, and thematic evaluations, there remains a persistent "knowledge-to-action" gap. Codified knowledge products often remain static and physically separated from active decision-making nodes. Users face transaction costs when attempting to synthesize lengthy, multilingual publications and evaluate their applicability to heterogeneous country, sector, or project contexts. The main challenge for PDFIs is therefore not producing more development research. It is helping users turn existing research into evidence they can apply when designing policies and projects before decisions are made.

5.1.3 Professional Knowledge Depends Too Heavily on Individual Experience

Institutional experience often resides in individual judgment. Expertise is valuable, but overreliance on a small group of people makes an organization vulnerable. A single question may involve financing policies, compliance rules, internal procedures, and historical cases. Experienced staff can navigate these materials quickly; new employees may not. Turnover and cross-departmental work widen the gap. Institutional memory requires searchable and traceable resources that preserve knowledge beyond individual experience.

5.2 Main Pathways of AI's Support for Integrated Knowledge Management and Institutional Learning

PDFIs can build AI-enabled knowledge management capacity through five connected functions.

First, integrating dispersed information into trusted natural-language entry points using semantic search and unified knowledge bases. This lowers search costs and speeds access to information, requiring document classification, version control, and access management.

Second, locating and synthesizing evidence for specific questions through cross-document synthesis, summary generation, and contextual explanations. This makes static documents usable in daily work, requiring source traceability and human review.

Third, improving verifiability through source links, citation mapping, and reasoned abstention (declining to answer when evidence is insufficient). This reduces information risk and strengthens trustworthiness, requiring source curation, audit records, and continuous evaluation.

Fourth, widening access through multilingual question answering, interactive summaries, courses, and podcasts. This lowers language and format barriers, requiring curated sources, localization, and feedback mechanisms.

Fifth, creating feedback loops through user feedback, phased deployment, and dynamic updates. This supports organizational learning and platform iteration, requiring training mechanisms, clear responsibilities, and evaluation systems.

5.3 Case Analysis

5.3.1 Improving Internal Knowledge Management through Fomento.AI at BPF

Case Summary

BPF, Portugal's national promotional bank, launched Fomento.AI in January 2026 as an internal AI platform that processes legal, financial, operational, and regulatory information. The system combines generative AI with semantic search, allowing staff to query internal documents in natural language and receive answers from institutionally verified content. The design emphasizes security, traceability, and alignment with BPF's privacy and compliance rules. As of the report's writing, Fomento.AI had entered phased rollout with power users; measurable long-term effects on time savings, consistency, or institutional memory still require evaluation.

BPF is Portugal's national promotional bank. In 2020, the Portuguese government merged existing institutions into Banco Português de Fomento, S.A. (BPF) (Portuguese Republic 2020).

Portuguese law defines BPF as a financial company (sociedade financeira) with a national promotional mandate to support business finance, investment, innovation, employment, the green transition, and balanced regional development (Portuguese Republic 2020).

BPF's knowledge problem reflects its institutional integration. Its chief technology and operations officer described a fragmented landscape of 24 applications and separate systems, which made authoritative information harder to locate quickly (Mateus 2026). Documents and rules therefore sit across network folders, email accounts, cloud storage, and other sources. This fragmentation raises search costs, creates version inconsistencies, slows onboarding, weakens institutional memory, and can lead staff to different conclusions when they rely on different files or individual experience. Fomento.AI addresses this governance problem by turning dispersed information into trusted, searchable, and reusable resources.

Fomento.AI has operated as BPF's internal AI platform with Glintt Next since January 2026. It processes legal, financial, operational, and regulatory information so staff can query and interpret key content more quickly, safely, and consistently (Banco Português de Fomento 2026; Glintt Next 2026). By combining generative AI with semantic search, the system retrieves relevant files and generates answers from institutionally verified content. The design emphasizes security, traceability, and alignment with BPF's privacy and compliance rules, making Fomento.AI a controlled RAG knowledge assistant rather than an open-ended chatbot (Glintt Next 2026). It lowers search costs by allowing staff to query internal documents in natural language; reduces inconsistencies caused by document-version differences and individual interpretation; and supports institutional memory by making rules, procedures, and past knowledge reusable beyond a small group of experienced employees. The assistant still does not replace staff judgment: employees must verify answers against the business context and consult original documents for consequential decisions.

BPF opened Fomento.AI to a group of early users, referred to internally as "Power Users"; collected feedback in real workflows; and planned a gradual rollout with tutorials, training, usage rules, staff skills development, responsibility for updates, feedback mechanisms, and later comparison of the time required to prepare reports and operational documents before and after deployment (Mateus 2026). Available evidence therefore supports a cautious maturity assessment: Fomento.AI has entered real-world testing and phased rollout, but measurable long-term effects still require evaluation.

BPF keeps the platform in a controlled environment because banks handle contracts, client information, and employee data. Institutional control reduces reliance on external general-purpose chat systems and links knowledge management with privacy, compliance, and operational risk (Mateus 2026). Its long-term value depends on regular document updates, clear version management, role-based access rights, traceable answers, human review for consequential decisions, and feedback-driven improvement of the knowledge base. Together, these conditions can turn an internal tool into a stable knowledge management capability.

5.3.2 Turning Research into Usable Evidence with AVA at World Bank

Case Summary

The World Bank Group Institute for Economic Development (IED) launched AVA, a curated multilingual knowledge assistant that draws on more than 4,000 World Bank and partner research materials in more than 60 languages. AVA is a multi-agent, domain-bounded RAG system with citation verifiability and reasoned abstention—it declines to answer when evidence is insufficient. An April 2026 evaluation of more than 2,200 users across 116 countries showed that sustained use was associated with self-reported time savings of approximately 2.4–3.9 hours per week, with page-level citations cited as particularly helpful for verification (Karnatak et al. 2026).

The World Bank Group IED, established in 2024, connects research with action through regional hubs in Rome, Tokyo, and Washington, D.C., and through partnerships with think tanks, universities, and experts (Institute for Economic Development n.d.a). IED responds to a persistent MDB problem: development research is abundant, but information overload, language barriers, and limited contextual relevance often prevent users from applying evidence in time (World Bank 2026). Policymakers may need to read long PDF reports separately, compare sources manually, and synthesize evidence under time pressure; for non-English users, the cost is even higher. IED therefore presents knowledge through conversations, summaries, courses, podcasts, and regional partnerships, not only static reports.

General-purpose generative AI can summarize information quickly, but unclear sources and unsupported answers weaken trust. AVA responds with a curated entry point for identifying, synthesizing, and verifying evidence. It draws on World Bank reports and selected publications; provides multilingual summaries, cross-topic syntheses, source links, and writing support; and gives users on-demand access to more than 4,000 World Bank and partner research materials in more than 60 languages (Institute for Economic Development n.d.b; World Bank 2026). Unlike open-internet tools, AVA restricts responses to a curated knowledge base, reducing noise and making verification easier. It does not replace original reports or generate policy recommendations automatically.

AVA is characterized as a multi-agent, domain-bounded RAG system with citation verifiability and reasoned abstention built into its design. This architecture aims to keep unverified content out of policy analysis: traceable citations allow users to return to original reports, whereas reasoned abstention prevents unsupported generation when evidence is insufficient or a question falls outside the knowledge base (Institute for Economic Development n.d.b; Karnatak et al. 2026). The design principle is therefore not to remove verification, but to reduce its cost. Users still need to judge whether evidence fits a specific country, sector, or project context.

AVA also sits within IED's wider knowledge service ecosystem. IED curates research, connects communities, adapts knowledge to user needs, and measures results through feedback (World Bank 2026). The ecosystem includes multilingual question answering, "Chat with the Report," interactive summaries, knowledge cards, online learning, and the expert-led podcast series Accelerating Development: Ideas for Impact in seven languages; regional initiatives such as Africa Growth and Opportunity:

Research in Action connect researchers, policymakers, and partners so global evidence can interact with local experience.

An evaluation covered more than 2,200 users from a wide range of institutions across 116 countries, institutions, and professional roles, using system logs, surveys, and 20 interviews (Karnatak et al. 2026). Sustained use was associated with savings of about 2.4–3.9 hours per week. Page-level citations helped users verify evidence quickly, and reasoned abstention clarified the platform's limits (Karnatak et al. 2026). These findings suggest that a curated assistant can reduce the time required to search for, organize, and validate evidence.

The results nevertheless require caution. The time savings are early findings and may not apply to all users or tasks; the research team includes World Bank staff and technical partners, so independent research would strengthen the evidence base. Efficiency gains also cannot replace assessment of content quality or local relevance.

AVA helps users find and use evidence from development research; it does not decide what a government or practitioner should do. Users must still check the cited sources and decide whether the findings fit local institutions and policy objectives. Professionals remain responsible for interpretation and decisions.

5.4 Takeaways from Case Analysis

Based on the analysis of BPF's Fomento.AI and the World Bank's AVA knowledge assistant, three main takeaways emerge for knowledge management and institutional learning in PDFIs.

First, AI addresses the core problem that abundant information does not readily turn into usable knowledge by creating unified, natural-language entry points across fragmented, non-interoperable databases. BPF's Fomento.AI integrates 24 separate applications into a single semantic search and RAG assistant, lowering search costs, reducing version-control inconsistencies, and preserving institutional memory beyond the experience of individual staff members.

Second, AI helps close the knowledge-to-action gap by making static knowledge products dynamic, verifiable, and contextually interpretable within decision-making workflows. The World Bank's AVA system draws on more than 4,000 research materials across 60-plus languages, using a multi-agent RAG architecture with citation verifiability and reasoned abstention—declining to answer when evidence is insufficient—so that policymakers and practitioners can find, verify, and apply evidence before making decisions rather than after. Third, both cases reinforce that AI serves as an assistive layer, not a replacement for professional judgment. Fomento.AI operates as a controlled internal assistant within BPF's privacy and compliance rules, requiring staff to verify answers against original documents for consequential decisions. AVA restricts responses to a curated knowledge base and cannot generate policy recommendations automatically—users must still assess whether evidence fits their specific country, sector, or project context.

The overarching takeaway is that effective AI for knowledge management transforms fragmented, static information into searchable, traceable, and reusable institutional resources, but only when built on source transparency, verifiable citations, reasoned abstention, and clear human responsibility for final interpretation and action.

VI Conclusions

This exploratory report shows that PDFIs are beginning to deploy AI as an instrument to overcome challenges to enhance inclusive and tailored access, operational execution, proactive risk management, and integrated knowledge management and institutional learning.

However, the evidence base remains thin: most cases are descriptive rather than evaluative, and many promising applications are still in pilot or pre-operational stages. The report's main thesis—that AI holds the potential for empowering PDFIs to better fulfill public policy objectives—is promising but requires stronger empirical support from operational deployments with published metrics. The Public Development Finance Observer Series will continue to delve deeper into this domain and provide the latest analysis on a regular basis.

Despite the thin evidence, governance risks are clear: algorithmic bias, reduced transparency, and potential exclusion of marginalized populations. PDFIs should therefore prioritize explainable AI (i.e., an AI system that can tell you why it gave a particular answer, prediction, or recommendation—in terms a human can understand), human oversight, regular bias audits, and inter-institutional collaboration on shared standards.

6.1 Policy Implications

The integration of AI into public development finance is not without profound risks. The very capabilities that make ML powerful—its ability to capture nonlinear relationships at scale—also make it susceptible to algorithmic bias. If poorly governed, AI could inadvertently replicate historical inequalities, directing capital toward data-rich urban centers while further marginalizing rural, data-poor populations. Therefore, the policy implications for PDFIs are clear and urgent. AI adoption cannot remain relegated to isolated pilot projects led by IT departments; it requires an institution-wide governance framework tailored to the ethics of public development finance.

First, PDFIs must prioritize explainable AI. Because PDFI decisions carry severe public consequences, algorithmic outputs should remain transparent and auditable. Human professional judgment must remain the final arbiter in project approval, risk sanctioning, and compliance verification.

Second, PDFIs must regularly audit user-facing and credit-scoring algorithms to ensure they pass the "inclusion test" and do not perpetuate systemic biases against vulnerable demographic groups.

Third, PDFIs should work as a system and collaborate to establish open-source shared datasets and common algorithmic auditing standards. Establishing regulatory sandboxes will allow PDFIs to safely test inclusive AI innovations before full deployment.

6.2 Future Outlook

Looking ahead, the success of PDFIs in deploying

AI hinges on whether they can govern it with the transparency, fairness, and human accountability necessary to harness the power of AI as a force for good.

First, before deploying AI models, PDFIs must invest in data infrastructure and embed procurement discipline to enable procedural transparency and open the black box. This means contracting only vendors that grant audit rights, mandate explainability, and accept liability for harms—alongside building unified data platforms that make model decisions traceable and auditable. Without this dual foundation of infrastructure and procurement rigor, opaque systems will erode public trust rather than strengthen it. For small PDFIs, however, investing in such data infrastructure may be challenging—requiring lighter, incremental approaches or shared regional platforms and collaborative vendor contracts to still achieve transparency without large-scale systems.

Second, they will build internal AI literacy across all staff, not just within IT departments, so that program officers, risk analysts, and frontline staff can effectively query, verify, and challenge AI outputs rather than passively accepting or blindly rejecting them.

Third, they will maintain human accountability for consequential decisions, using AI as decision support rather than a substitute—meaning that loan approvals, fraud determinations, and resource allocations remain human responsibilities, with AI serving as an investigative lead or forecasting assistant.

Fourth, they need to engage with regulators early through sandboxes and pilot frameworks, testing high-risk applications like automated underwriting or forecast-triggered funding in controlled environments before full deployment, thereby containing failures and building evidence for rulemaking.

Fifth, they need to measure what matters—not just time saved or operational efficiency, but the effectiveness of fulfilling their public policy objectives outcomes for underserved populations, fairness across demographic groups, and institutional learning over time.

In a nutshell, the technology is moving fast, but the organizational and governance work will determine whether AI widens or narrows inclusive access, strengthens or weakens accountability, and deepens or erodes trust in PDFIs.

References

African Development Bank Group. 2016. *Jobs for Youth in Africa: Strategy for Creating 25 Million Jobs and Equipping 50 Million Youth 2016–2025.* https://www.afdb.org/sites/default/files/documents/policy-documents/bank_group_strategy_for_jobs_for_youth_in_africa_2016-2025_rev_2.pdf.

African Development Bank Group. 2025a. *African Development Bank Group Ten-Year Strategy 2024–2033: Investing in Africa's Transformation.* Revised April. https://www.afdb.org/en/documents/african-development-bank-group-ten-year-strategy-2024-2033.

African Development Bank Group. 2025b. "African Development Bank Launches ENNOVA AI Platform to Boost Entrepreneurship and Job Creation at Sankalp Summit." Press release, March 21. https://www.afdb.org/en/news-and-events/press-releases/african-development-bank-launches-ennova-ai-platform-boost-entrepreneurship-and-job-creation-sankalp-summit-81959.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence.* Harvard Business Review Press.

Agricultural Development Bank of China. n.d. "About Us." Accessed June 5, 2026. https://www.adbc.com.cn/en/n1060/n1066/index.html.

AIAIC Maharashtra. 2026. "At AI4Agri 2026, Shaji K. V., Chairman of National Bank for Agriculture and Rural Development (NABARD), Shared His Perspective on the Importance of . . ." Instagram video, April 6. Accessed June 5, 2026. https://www.instagram.com/reel/DWOJisKD4jD/.

Banco Português de Fomento. 2026. "BPF lança Plataforma de AI para Colaboradores" [BPF Launches AI Platform for Employees]. News release, January 14. https://www.bpfomento.pt/pt/noticias/banco-portugues-de-fomento-lanca-plataforma-de-ai-para-colaboradores/.

Boukherouaa, El Bachir, Ghiath Shabsigh, Khaled AlAjmi, et al. 2021. *Powering the Digital Economy: Opportunities and Risks of Artificial Intelligence in Finance.* Departmental Paper 2021/024. International Monetary Fund. https://doi.org/10.5089/9781589063952.087.

Bracken, Matt. 2025. "Fannie Mae Partners with Palantir on Mortgage Fraud Detection Tech." *FedScoop,* May 28. https://fedscoop.com/fannie-mae-partners-with-palantir-on-mortgage-fraud-detection-tech/.

Cao, Longbing. 2022. "AI in Finance: Challenges, Techniques, and Opportunities." *ACM Computing Surveys* 55 (3): Article 64, 1–38. https://doi.org/10.1145/3502289.

Chatrabhuj, Kundan Meshram, Umank Mishra, and Upaka Rathnayake. 2025. "Application of Artificial Intelligence in Agri-Tech, Environmental and Biodiversity Conservation."

Array 26: 100412. https://doi.org/10.1016/j.array.2025.100412.

Department of Financial Services, Ministry of Finance, Government of India. n.d. "Scheme Details." *Pradhan Mantri Jan-Dhan Yojana (PMJDY).* Accessed June 5, 2026. https://www.pmjdy.gov.in/scheme.

Fannie Mae. 2025. "Fannie Mae Launches AI Fraud Detection Technology Partnership with Palantir." Press release, May 28. https://www.fanniemae.com/newsroom/fannie-mae-news/fannie-mae-launches-ai-fraud-detection-technology-partnership-palantir.

Fannie Mae. n.d.a. "History." Accessed June 2, 2026. https://www.fanniemae.com/about-us/who-we-are/history.

Fannie Mae. n.d.b. "Mortgage Fraud Prevention." Accessed June 2, 2026. https://singlefamily.fanniemae.com/mortgage-fraud-prevention.

Federal Housing Finance Agency. 2022. "History of Fannie Mae and Freddie Mac Conservatorships." Last updated October 17. https://www.fhfa.gov/conservatorship/history.

Federal Housing Finance Agency. 2025. *AB 2022-02: Artificial Intelligence/Machine Learning Risk Management.* Advisory Bulletin. Originally issued February 10, 2022; revised May 2025. https://www.fhfa.gov/advisory-bulletin/ab-2022-02.

Federal Housing Finance Agency Office of Inspector General. 2022. *Enterprise Use of Artificial Intelligence and Machine Learning.* White Paper WPR-2022-002, September 19. https://www.fhfaoig.gov/WhitePaper/enterprise-use-artificial-intelligence-and-machine-learning.

Financial Stability Board. 2017. *Artificial Intelligence and Machine Learning in Financial Services: Market Developments and Financial Stability Implications.* November 1. https://www.fsb.org/2017/11/artificial-intelligence-and-machine-learning-in-financial-service/.

Freddie Mac. n.d. "About Us." Accessed June 5, 2026. https://www.freddiemac.com/about.

Fuster, Andreas, Matthew Plosser, Philipp Schnabl, and James Vickery. 2019. "The Role of Technology in Mortgage Lending." *The Review of Financial Studies* 32 (5): 1854–99. https://doi.org/10.1093/rfs/hhz018.

Galileo Financial Technologies. 2025. "How Banks Can Improve Customer Support with Conversational AI." October 9. https://www.galileo-ft.com/blog/banks-customer-support-conversational-ai/.

Glintt Next. 2026. "Banco Português do Fomento Invests in Generative AI to Modernize Internal Processes." January 15. https://www.glinttnext.com/banco-portugues-do-fomento-invests-in-generative-ai-to-modernize-internal-processes/.

Institute for Economic Development. n.d.a "AVA: AI + Verified Analysis." Accessed June 3, 2026. https://www.wbginstitute.org/ava.

Institute for Economic Development. n.d.b "Our Work." Accessed June 3, 2026. https://www.wbginstitute.org/about.

International Monetary Fund and World Bank Group. 2018. *The Bali Fintech Agenda.* IMF Policy Paper 18/050. October 11. https://www.imf.org/en/publications/policy-papers/issues/2018/10/11/pp101118-bali-fintech-agenda.

Karnatak, Nimisha, Mohamad Chatila, Daniel Alejandro Pinzón Hernández, Reza Yazdanfar, Michelle Dugas, and Renos Vakis. 2026. "Learning from AVA: Early Lessons from a Curated and Trustworthy Generative AI for Policy and Development Research." In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems,* 1–32. Association for Computing Machinery. https://doi.org/10.1145/3772318.3791062.

Korea Housing Finance Corporation. 2025. "HF to Enhance Accessibility of Housing Finance Services through AI Transformation." Press release, December 29. https://www.hf.go.kr/en/sub05/sub05_01.do?articleNo=599906&mode=view&title=HF+to+Enhance+Accessibility+of+Housing+Finance+Services+Through+AI+Transformation.

Li, Yan, and Deng Guoping. 2025. " 农发行数智云柜投产上线 " [ADBC Smart Digital Cloud Counter Goes Live]. Agricultural Development Bank of China, December 11. https://www.adbc.com.cn/n5/n15/c691917/content.html.

Mateus, Tiago. 2026. "'No Futuro Haverá Certamente Especialistas em Prompts de IA, Que Serão Valiosos em Qualquer Indústria.'" Interview by Tiago Alexandre Pereira. *ECO*, January 28. https://eco.sapo.pt/entrevista/no-futuro-havera-certamente-especialistas-em-prompts-de-ia-que-serao-valiosos-em-qualquer-industria/.

Mearns, Robin, Benjamin Reese, and Chris Mahony. 2025. "How AI Can Support Anticipatory Action to Address Forced Displacement." *Development for Peace* (blog), June 3. https://blogs.worldbank.org/en/dev4peace/how-ai-can-support-anticipatory-action-to-address-forced-displac.

Mention, Anne-Laure. 2019. "The Future of Fintech." *Research-Technology Management* 62 (4): 59–63. https://doi.org/10.1080/08956308.2019.1613123.

National Bank for Agriculture and Rural Development. 2024. *Annual Report 2023–24.* https://www.nabard.org/nabard-annual-report-2023-24.aspx.

OECD. 2019. *Recommendation of the Council on Artificial Intelligence.* OECD Legal Instruments, OECD/LEGAL/0449. https://legalinstruments.oecd.org/en/instruments/oecd-legal-0449.

Policy Commons. n.d. "NABARD: National Bank for Agriculture and Rural Development." Accessed June 5, 2026. https://policycommons.net/orgs/national-bank-for-agriculture-and-rural-development/.

Portugal. Presidência do Conselho de Ministros. 2020. *Decreto-Lei n.º 63/2020, de 7 de setembro: Regula a Atividade e Funcionamento do Banco Português de Fomento, S.A., e Aprova os Respetivos Estatutos. Diário da República,* 1.ª série, no. 174, September 7. https://diariodarepublica.pt/dr/detalhe/decreto-lei/63-2020-142124830.

Rasheed, Muhammad Haroon, Jamshed Khalid, and Muhammad Shahid Rasheed. 2026. "Artificial Intelligence in Finance: Unveiling Foundations, Themes, and Future Research Agendas." *Future Business Journal* 12: 48. https://doi.org/10.1186/s43093-026-00759-z.

Republic of Korea. 2003. *Korea Housing Finance Corporation Act.* Act No. 7030, December 31. Amended through Act No. 18574, December 7, 2021. English translation by the Korea Legislation Research Institute. Accessed June 5, 2026. https://elaw.klri.re.kr/eng_mobile/viewer.do?hseq=60116.

Shetty, Roshan. 2025. “This Is What the New Frontier of AI-Powered Financial Inclusion Looks Like.” *World Economic Forum,* December 9. https://www.weforum.org/stories/2025/12/this-is-what-the-new-frontier-of-ai-powered-financial-inclusion-looks-like/.

Sirtaine, Sophie. 2025. “AI’s Promise: A New Era for Financial Inclusion.” *CGAP,* April 4. https://www.cgap.org/blog/ais-promise-new-era-for-financial-inclusion.

Structured Data Analytics. 2024. *Methodology Report—Uganda: Artificial Intelligence to Forecast Population Movements for Forced Displacement Preparedness.* Prepared on behalf of the World Bank, June 12. https://documents.worldbank.org/en/publication/documents-reports/documentdetail/099061725090035169.

Structured Data Analytics. 2025. *Report on Model and Results—Uganda: Artificial Intelligence to Forecast Population Movements for Forced Displacement Preparedness.* Prepared on behalf of the World Bank, January 14. https://documents.worldbank.org/en/publication/documents-reports/documentdetail/099061725090040236.

Tooze, Adam. 2022. “Welcome to the World of the Polycrisis.” *Financial Times,* October 28. https://www.ft.com/content/498398e7-11b1-494b-9cd3-6d669dc3de33.

World Bank. 2021. *World Development Report 2021: Data for Better Lives.* World Bank. https://doi.org/10.1596/978-1-4648-1600-0.

World Bank. 2025a. *Building Trustworthy Artificial Intelligence: Frameworks, Applications, and Self-Assessment for Readiness.* World Bank. https://doi.org/10.1596/43656.

World Bank. 2025b. “AI-Powered Refugee Forecasting: Preparing for Refugee Movements before They Happen.” Brief, June 3. https://www.worldbank.org/en/topic/fragilityconflictviolence/brief/ai-powered-refugee-forecasting-preparing-for-refugee-movements-before-they-happen.

World Bank. 2026. “Research in Action: Fast, Far, and for All.” Feature story, January 8. https://www.worldbank.org/en/news/feature/2026/01/08/research-in-action-fast-far-and-for-all.

World Bank Group. n.d. “What We Do.” Accessed June 2, 2026. https://www.worldbank.org/ext/en/what-we-do.

World Bank–UNHCR Joint Data Center on Forced Displacement. 2026. “Artificial Intelligence Predictive Modelling for Forced Displacement Preparedness.” March 27. Accessed June 2, 2026. https://www.jointdatacenter.org/activites/ai-predictive-modelling-uganda/.

Xu, Jiajun, Régis Marodon, Xinshun Ru, Xiaomeng Ren, and Xinyue Wu. 2021. “What Are Public Development Banks and Development Financing Institutions?—Qualification Criteria, Stylized Facts, and Development Trends.” *China Economic Quarterly International* 1 (4): 271–94. https://doi.org/10.1016/j.ceqi.2021.10.001.

PEKING UNIVERSITY

PDF Observer

Public Development Finance

Public Development Finance Research Program at Peking University

E-Mail: nsedfi@nsd.pku.edu.cn

Address: Peking University, No. 5 Yiheyuan Road, No. 38 Xueyuan Road, Haidian District, Beijing.

Website: http://www.dfidatabase.pku.edu.cn