

LLM-Driven Identification of Digital Financial Inclusion Gaps in Developing Economies

Qi Xin^{1*}

¹ qix29@pitt.edu

¹ University of Pittsburgh, Pennsylvania, United States.

**Corresponding author*

Received: 14 April 2026

Revised: 28 August 2026

Published: 31 August 2026

Abstract

This study identifies the developing economies that face the largest digital-finance activation bottlenecks after basic account access has already been achieved. Using the 2024 country layer of Global Findex 2025 and merging it with World Development Indicators on internet use and mobile subscriptions, the paper constructs an activation-gap measure that captures the distance between account ownership and high-frequency digital payment use. The analytical sample contains 86 developing economies after restricting the data to national all-adult observations and to countries with complete target variables. The average country in the sample had account ownership of 58.68% but a high-frequency payment index of only 21.64%, leaving a mean activation gap of 37.04 percentage points. The empirical strategy estimated repeated five-fold cross-validated Ridge, Random Forest, XGBoost, and LightGBM models for both regression and high-gap classification. XGBoost produced the strongest regression fit, with an RMSE of 5.32 percentage points and an R^2 of 0.794, and it also delivered the best classifier, with an AUC of 0.863. The largest observed gaps were recorded in Ghana (64.91 percentage points), Uganda (64.61), and Zambia (60.88). SHAP-based interpretation shows that account ownership, internet use, mobile money account penetration, card-linked government payment channels, and digital skills were the most influential predictors. The results establish that the central policy challenge in many developing economies is no longer account opening alone; it is the conversion of nominal access into repeated digital payment behavior. The final diagnostic layer converts measured model outputs into country-specific policy notes, creating a reproducible bridge from predictive analytics to practical public policy design.

Keywords : digital financial inclusion; Global Findex; digital payments

Abstrak

Studi ini mengidentifikasi negara-negara berpenghasilan berkembang yang menghadapi hambatan aktivasi keuangan digital terbesar setelah akses akun dasar berhasil dicapai. Menggunakan lapisan data tingkat negara tahun 2024 dari Global Findex 2025 dan mengintegrasikannya dengan *World Development Indicators* mengenai penggunaan internet serta penetrasi langganan telepon seluler, penelitian ini mengonstruksi indikator kesenjangan aktivasi (*activation-gap measure*) yang mengukur jarak antara kepemilikan akun dan intensitas penggunaan pembayaran digital berfrekuensi tinggi. Sampel analitis mencakup 86 negara berkembang setelah membatasi data pada observasi populasi dewasa tingkat nasional dan negara-negara dengan variabel target yang lengkap. Rata-rata negara dalam sampel mencatatkan tingkat kepemilikan akun sebesar 58,68%, namun indeks pembayaran berfrekuensi tinggi hanya mencapai 21,64%, sehingga menghasilkan rerata kesenjangan aktivasi sebesar 37,04 poin persentase. Strategi empiris dalam studi ini mengestimasi model *Ridge*, *Random Forest*, *XGBoost*, dan *LightGBM* yang divalidasi silang secara berulang (*repeated five-fold cross-validated*) baik untuk analisis regresi maupun klasifikasi kesenjangan tinggi. Model *XGBoost* menghasilkan tingkat keselarasan regresi (*regression fit*) terbaik dengan nilai *RMSE* sebesar 5,32 poin persentase dan R^2 sebesar 0,794, sekaligus menjadi model klasifikasi paling optimal dengan nilai *AUC* sebesar 0,863. Kesenjangan terbesar tercatat di Ghana (64,91 poin persentase), Uganda (64,61), dan Zambia (60,88). Interpretasi berbasis *SHAP* menunjukkan bahwa kepemilikan akun, penggunaan internet, penetrasi akun uang elektronik (*mobile money*), saluran pembayaran pemerintah berbasis kartu, dan keterampilan digital merupakan prediktor yang paling berpengaruh. Hasil penelitian ini menegaskan bahwa tantangan kebijakan utama di banyak negara berkembang saat ini bukan lagi sekadar pembukaan akun, melainkan konversi akses nominal menjadi perilaku pembayaran digital yang berkelanjutan. Lapisan diagnostik akhir mengonversi luaran model yang terukur menjadi catatan kebijakan khusus per negara, sehingga menjembatani analisis prediktif dengan perancangan kebijakan publik secara praktis dan dapat direplikasi.

Kata kunci: inklusi keuangan digital; Global Findex; pembayaran digital

1. INTRODUCTION

Financial inclusion research has long shown that access to formal finance matters for savings, resilience, entrepreneurship, and poverty reduction (Beck et al., 2007; World Bank, 2014). In the last decade, however, the policy conversation has shifted from basic account access to effective digital usage. Global Findex evidence has documented large gains in account ownership and a rapid spread of digital payment rails, especially after the expansion of mobile phones, agent networks, and low-cost transaction interfaces (Demirguc-Kunt et al., 2018, 2022). Yet the spread of accounts does not automatically produce routine digital payments. Many countries have moved into a condition in which people are counted as financially included because they have an account, but they still rely on cash for merchant transactions, household transfers, and everyday settlement.

That distinction matters. An inactive account has very different development implications from an account that anchors frequent digital transactions. When payments are executed digitally and repeatedly, governments can deliver transfers more efficiently, firms can digitize payroll and supplier relations, households can lower transaction costs, and the data generated by payment activity can support broader financial deepening (Allen et al., 2016; World Bank, 2016). When account ownership expands without regular use, these benefits remain partial. The result is a hidden form of digital-finance exclusion inside headline inclusion statistics.

The distinction between ownership and activation also helps reinterpret recent inclusion achievements. A sharp rise in account numbers can reflect mass onboarding through public-transfer programs, emergency payments, or simplified digital accounts. Those expansions are valuable, but they represent the beginning of a payments relationship, not the end of it. The operational question is whether the same user later pays merchants digitally, stores value digitally, or returns to the account often enough for network effects to build. By treating activation as a separate empirical outcome, the paper aligns measurement more closely with what payment-system reform is actually trying to accomplish.

The activation problem is especially important in developing economies. These countries often have limited fiscal space and cannot afford to pursue indiscriminate financial-inclusion strategies. They need to know where the constraint lies: weak connectivity, low merchant acceptance, cash-heavy government payments, low mobile-money activation, or inadequate digital trust and skills. Existing cross-country studies have explained why mobile money succeeds in some environments and how digital finance affects welfare, but far fewer studies have ranked countries by the size of the access-to-usage gap and linked those rankings to operational policy bottlenecks (Aker & Mbiti, 2010; Aron, 2018; Jack & Suri, 2014; Suri & Jack, 2016).

This paper addresses that problem by asking a sharper question: which developing economies are least able to convert “having an account” into high-frequency digital payment behavior? The answer requires an outcome variable that goes beyond ownership, an empirical design that treats countries as directly comparable units, and a way to translate quantitative predictions into country-level policy language. To that end, the paper builds an activation-gap index from Global Findex payment and frequency variables, augments the country records with WDI measures of internet use and mobile subscriptions, and then evaluates multiple machine-learning models under repeated cross-validation.

A focus on payment frequency is also consistent with how payment systems generate value. A country does not obtain large economic efficiency gains when adults use an account

once to receive a transfer and then immediately cash out. Gains arise when accounts become storage, transfer, and settlement devices across multiple transaction types. That is why this paper gives explicit attention to merchant payments and weekly use in the target design. Merchant activity matters because it embeds digital finance in day-to-day consumption. Weekly use matters because it signals habit formation rather than one-off compliance with a government or employer requirement. An inclusion agenda that ignores this distinction risks overstating effective financial modernization.

Cross-country benchmarking is particularly useful at the present stage of financial inclusion policy. Many development agencies and finance ministries already know the broad ingredients of digital transformation, but they still face an allocation problem. Which countries should prioritize merchant acceptance? Which should prioritize payment digitization in social protection? Which should prioritize data affordability and mobile internet access? The empirical value of a comparative country model lies in answering these triage questions systematically. Even when a model is not causal, a well-validated ranking can still improve the targeting of diagnostic missions, technical assistance, and reform sequencing.

The study makes three contributions. First, it defines a country-level activation gap that measures the distance between account ownership and routine digital payment behavior rather than treating them as equivalent outcomes. Second, it provides a full empirical comparison of linear and tree-based models using measured out-of-fold performance rather than illustrative results. Third, it adds an LLM-oriented diagnostic layer that turns country scores, probabilities, and bottleneck signals into standardized policy notes that can be read directly by analysts or embedded in a larger policy workflow. This design preserves empirical discipline while still serving the practical needs of digital-finance policy.

The empirical picture that emerges is clear. In the final sample of developing economies, the average activation gap was {metrics['mean_gap']:.2f} percentage points, with especially large gaps in South Asia and Sub-Saharan Africa. Tree-based models outperformed the linear baseline, and the strongest predictors were not generic macro variables but indicators tied to internet-enabled use, mobile-account activation, government payment channels, and digital skills. The paper therefore argues that the next stage of financial inclusion policy should be judged less by whether an account exists and more by whether that account becomes a regularly used payment instrument.

2. LITERATURE REVIEW

The first relevant strand of literature studies the economic role of financial inclusion itself. Beck et al. (2007) and World Bank (2014) linked broader access to finance with lower inequality and stronger inclusion of poorer households. Allen et al. (2016) moved the discussion from aggregate depth to the household-level foundations of account ownership and use, showing that legal, demographic, and infrastructural barriers jointly shape formal financial participation. The Global Findex reports then established a standard international measurement framework for inclusion and repeatedly showed that ownership, borrowing, payments, and resilience are distinct dimensions rather than interchangeable indicators (Demirguc-Kunt et al., 2018, 2022).

A second strand focuses on mobile phones, mobile money, and digital payments. Aker and Mbiti (2010) argued early that mobile telephony could transform information flows and reduce economic frictions in Africa. Subsequent work established that mobile payment systems can change the structure of transaction costs and risk sharing. Jack and Suri (2014) showed that

Kenya's mobile money system lowered remittance frictions and strengthened informal insurance, while Suri and Jack (2016) documented long-run poverty and gender effects. Aron's (2018) review consolidated this evidence and emphasized that adoption is shaped not only by technology availability but also by agent networks, trust, interoperability, and use cases. GSMA's annual industry reports have reinforced the same point at the market level: scale is reached when payment ecosystems support repeated use, not merely wallet registration (GSMA, 2023, 2024).

Measurement itself has been a recurring concern in the inclusion literature. The most cited indicators are easy to communicate but often compress different forms of engagement into a single headline number. That simplification is useful for global monitoring, yet it is less useful for implementation. A country can look successful on ownership while remaining weak on everyday digital transactions, and that weakness may remain hidden until more behavioral indicators are examined. The activation-gap approach follows this logic by explicitly separating access from use after ownership has already been established.

A third literature examines digital finance, fintech, and the broader digital transformation of financial services. Ozili (2018) argued that digital finance can expand inclusion and efficiency, but only when regulatory and institutional conditions allow new technologies to reduce the cost of reaching underserved users. Philippon (2016) framed fintech as an efficiency opportunity for finance more broadly, while Sahay et al. (2020) emphasized that fintech can widen inclusion after crises only if public policy addresses infrastructure, identification, consumer protection, and competition. World Bank (2016) and World Bank (2021) similarly stressed that digital dividends do not emerge automatically; institutions and capabilities determine whether data and digital channels become inclusive or exclusionary.

Another relevant insight from the literature is that payment use depends on ecosystem complementarities. Mobile money or account-based payment systems scale when users, merchants, agents, providers, and public institutions all expect others to participate. This is why some systems ignite and others stagnate despite similar access conditions. The practical implication is that country performance cannot be read from a single variable. High phone ownership without merchant acceptance produces weak payment intensity. Strong account ownership without digital trust produces dormant accounts. Large transfer programs delivered digitally but cashed out immediately produce enrollment without activation. The cross-country activation-gap framework is therefore designed to reflect a system, not a single adoption margin.

The literature also suggests why explanation tools matter. Predictive models are attractive because they can detect threshold effects and interaction patterns that are easy to miss in linear regressions, but policymakers often mistrust black-box scores. SHAP-style decomposition is useful in this context because it converts a ranked prediction into a ranked set of contributing signals. That makes the model more than an academic exercise. It can be read as a structured diagnostic instrument, especially when the explanation is paired with country facts that policy teams already understand, such as cash-heavy transfer delivery, weak app-based banking use, or low mobile-data access.

A fourth strand is methodological. Financial-inclusion studies have traditionally relied on descriptive comparisons, reduced-form regressions, and case studies. Those tools remain valuable, but they often struggle to model nonlinear interactions among access, connectivity, behavior, and trust variables. Ensemble methods directly address that challenge. Random forests capture complex interaction structures through recursive partitioning (Breiman, 2001).

Gradient boosting improves predictive fit by iteratively correcting prior errors (Friedman, 2001). XGBoost and LightGBM then operationalized boosting in computationally efficient forms that have become standard in tabular prediction tasks (Chen & Guestrin, 2016; Ke et al., 2017). To keep such models policy-relevant rather than black-box, SHAP values provide an additive explanation framework that ranks the drivers of each prediction and supports interpretable reporting (Lundberg & Lee, 2017).

Despite these advances, one gap remains underexplored. Most cross-country work still measures digital inclusion either at the level of access or at the level of broad digital payment adoption. Much less attention has been given to the transition between those stages. The question is not only who lacks an account, but who fails to activate the account they already have. That missing middle matters because policy interventions differ sharply across stages. Low ownership points to onboarding and identification barriers. A large activation gap points instead to payment acceptance, low-frequency use, behavioral frictions, connectivity quality, or distrust of digital channels.

This paper positions itself at that intersection. It draws on the inclusion literature to treat ownership as a necessary but incomplete condition, on the mobile-money literature to interpret repeated payment use as an ecosystem outcome, and on the machine-learning literature to model nonlinear country patterns with transparent feature attribution. The resulting framework does not replace causal identification, but it does offer something existing studies rarely provide: a reproducible country ranking of activation risk, benchmarked across many developing economies and directly linked to policy diagnosis.

3. RESEARCH METHODOLOGY

The empirical analysis used the country-level 2024 wave of Global Findex 2025 and merged it with two World Development Indicators series: the share of individuals using the internet and mobile cellular subscriptions per 100 people. The raw Global Findex file contains national observations, regional aggregates, income-group aggregates, and demographic splits. The sample was restricted to records with year = 2024, group = all, and group2 = all so that each retained row represented the full adult population of one economy. High-income economies and non-country aggregates were then removed. Table 1 reports the sample construction process. The final prediction sample contains {metrics['n']} developing economies because seven developing-country rows lacked at least one target component needed to construct the activation-gap measure.

Figure 1. Empirical workflow from country data to policy notes




Table 1. Sample construction

Step	Observations
Global Findex 2025 rows loaded	8564
2024 observations	1969
All adults (group=all, group2=all)	152
Developing economies only (exclude high income and aggregates)	93
Countries with non-missing activation-gap target	86

Table 2. Variable operationalization

Concept	Variable	Source	Definition	Role
Account ownership	account_t_d	Global Findex	Share of adults with an account at a financial institution or mobile money provider.	Outcome anchor
Any digital payment	fin4_d	Global Findex	Share of adults who made or received any digital payment.	Payment depth component
Merchant payment from account	merchant_pay	Global Findex	Share of adults who made merchant payments directly from an account.	High-frequency component
Weekly account use	fin5w	Global Findex	Share of adults who used an account weekly.	High-frequency component
High-frequency payment index	$0.5 \times \text{fin4_d} + 0.3 \times \text{merchant_pay} + 0.2 \times \text{fin5w}$	Constructed	Weighted payment-use index emphasizing repeated merchant and account activity.	Primary continuous target complement
Activation gap	$\text{account_t_d} - \text{hf_payment_index}$	Constructed	Distance between account access and high-frequency digital payment use.	Primary continuous target
High-gap classifier	1 if activation gap $\geq$ sample median	Constructed	Binary indicator for countries above the median activation gap.	Primary classification target
Internet users	IT.NET.USER.ZS	WDI	Latest available share of	External digital environment

			individuals using the internet.	
Mobile subscriptions	IT.CEL.SETS.P2	WDI	Latest available mobile cellular subscriptions per 100 people.	External digital environment

Table 2 defines the variables used in the study. The central empirical challenge was that no single country-level variable in the 2024 file directly measured “high-frequency digital payments” in a way that simultaneously captured breadth, payment intensity, and use frequency. The paper therefore constructed a composite high-frequency payment index as 0.5 times the share of adults making or receiving any digital payment, plus 0.3 times the share making merchant payments directly from an account, plus 0.2 times the share using an account weekly. The weighting scheme deliberately assigned the largest weight to broad digital-payment participation, a secondary weight to merchant-use intensity, and a smaller but still meaningful weight to weekly activity. The activation gap was then defined as account ownership minus this payment index. Higher values indicate that account access is materially ahead of routine digital use.

Country names and ISO codes were harmonized directly from the Global Findex and WDI files, and all merges were performed on the three-letter country code. The final processed dataset saved in the replication package includes the raw indicators, the engineered targets, the model predictions, the SHAP values, and the generated country notes. This packaging choice matters for transparency because it allows the user to trace every sentence in the diagnostics back to a concrete country record and a measured model output.

Table 3. Indicator coverage in the final country sample

Variable	Non_missing_countries	Coverage_rate_pct
account_t_d	86	100
con1	86	100
fiaccount_t_d	86	100
fin4_d	86	100
fin5w	86	100
fing2p_card	86	100
fing2p_cash	86	100
fing2p_fin	86	100
fing2p_mob	86	100
internet	86	100
merchant_pay	86	100
wdi_internet_latest	86	100
wdi_mobile_latest	86	100
con11	74	86.05
con12d	74	86.05
con17a	74	86.05
con17b	74	86.05
con26d	74	86.05
con27	74	86.05

con29	74	86.05
con30d	74	86.05
con31a	74	86.05
con31d	74	86.05
con31h	74	86.05
con32a	74	86.05
mobileaccount t d	73	84.88

The feature set combined three blocks. The first block captured access and account structure: overall account ownership, financial-institution account ownership, and mobile-money account ownership. The second block captured payment rails and digital behavior: the composition of government-to-person payments, mobile-phone ownership, self-reported internet use, internet access on a phone, internet use frequency, online shopping, financial use of the internet, app-based banking, and related variables. The third block captured the broader digital environment by merging the latest available WDI internet-user and mobile-subscription indicators by ISO country code. Because missingness was concentrated in several digital-behavior variables, Table 3 reports the exact coverage rate of each indicator.

The weighting scheme for the high-frequency payment index was fixed before model estimation and was not tuned on the outcome. This choice prevented the target from being tailored to any specific model. The heaviest weight was assigned to any digital payment because it captured the broadest entry into digital activity. Merchant payments were weighted next because they reflect routine consumer use and acceptance infrastructure. Weekly account use received the smallest weight because it is a narrower but highly informative marker of repeated engagement. The resulting target is interpretable as a conservative measure of whether account ownership has been converted into behavior that occurs often enough to matter for day-to-day economic exchange.

The study was intentionally designed as a reproducible tabular-learning exercise rather than as a large hyperparameter search. The sample is small by machine-learning standards, so very aggressive tuning would risk overfitting to country idiosyncrasies. For that reason, the tree-based models used compact parameterizations with shallow depth, moderate estimator counts, and subsampling. The objective of the comparison was not leaderboard optimization but stable out-of-fold generalization and a clean contrast between linear and nonlinear learners. This design choice is appropriate for policy benchmarking, where robustness and interpretability are more valuable than marginal in-sample gains.

Two supervised learning tasks were estimated. The first was continuous prediction of the activation gap. The second was binary classification of whether a country belonged to the high-gap half of the sample, defined as activation gap greater than or equal to the sample median of {metrics['median_gap']:.2f} percentage points. For regression, the paper estimated Ridge, Random Forest, XGBoost, and LightGBM. For classification, it estimated Logistic Regression, Random Forest, XGBoost, and LightGBM. Median imputation was applied to all models. Standardization was applied to the linear models, while the tree-based models used the imputed variables directly.

Model evaluation used repeated five-fold cross-validation with three repeats. This design yielded fifteen held-out evaluations per model while preserving enough training data for a small cross-country sample. Reported country-level performance metrics were calculated from averaged out-of-fold predictions so that each economy's final prediction was formed only

from models that did not train on that economy. Regression performance was assessed with RMSE, MAE, and R^2 . Classification performance was assessed with AUC, accuracy, F1, recall, precision, and balanced accuracy. All reported figures and tables were generated directly from these measured out-of-fold results; the manuscript does not contain illustrative or placeholder experimental outputs.

The paper also ran an ablation exercise with XGBoost because it was the strongest tree model in the main experiment. The ablation compared an accounts-only specification, a broader Findex behavioral specification, and the full Findex-plus-WDI specification. This step isolated whether the external WDI indicators contributed incremental predictive power once the rich within-Findex behavioral variables were already included.

Interpretation proceeded in two stages. First, the best-performing regression model was re-estimated on the full sample and explained with SHAP values, which provided both a global importance ranking and country-specific contribution profiles. Second, the measured model outputs were converted into country policy notes. The diagnostic generator combined each country's predicted gap, high-gap probability, and ranked bottleneck categories into a deterministic LLM-ready narrative. This choice kept the full workflow reproducible while preserving the core idea of an automated model-to-text policy diagnostic layer.

4. RESULTS AND DISCUSSION


The descriptive evidence already shows why a simple ownership metric is insufficient. Across the 86 developing economies in the sample, mean account ownership was 58.68%, while the mean high-frequency payment index was only 21.64%. The difference generated a mean activation gap of 37.04 percentage points. Figure 2 shows a broad distribution rather than a narrow cluster, with countries ranging from low-teen gaps to gaps above 60 percentage points. In other words, some countries have largely translated access into repeated use, while others have not.

Table 4. Summary statistics of the main variables

Variable	Mean	SD	Min	Median	Max
Account ownership (%)	58.68	19.16	14.83	57.50	98.28
High-frequency payment index (%)	21.64	16.05	2.01	17.22	69.58
Activation gap (pp)	37.04	11.78	12.82	36.46	64.91
Any digital payment (%)	26.62	19.90	1.83	22.08	72.05
Merchant payment from account (%)	24.11	19.59	2.42	18.39	94.94
Weekly account use (%)	5.47	5.74	0.15	3.18	25.89
WDI internet users (%)	60.19	24.27	10	64.84	97.69

WDI mobile subscriptions (per 100 people)	102	34.66	32.30	98.52	186
---	-----	-------	-------	-------	-----

Figure 2. Distribution of the activation gap



Regional comparisons sharpen the picture. Table 5 and Figure 3 show that South Asia recorded the highest mean activation gap at 44.90 percentage points, closely followed by Sub-Saharan Africa at 43.37. East Asia and Pacific averaged 32.37 percentage points, materially below those two regions despite substantial heterogeneity within the region. The income-group gradient was present but weaker than the regional pattern: low-income and lower-middle-income countries had larger mean gaps than upper-middle-income countries, yet some upper-middle-income countries still posted severe activation deficits. This pattern confirms that the gap is not simply a low-income problem; it is a problem of ecosystem conversion.

Table 5. Regional summary statistics


Region	Countries	Account_Ownership_pct	HF_Payment_Index_pct	Activation_Gap_pp	Internet_Users_pct
South Asia (excluding high income)	5	60.26	15.35	44.90	44.10
Sub-Saharan Africa (excluding high income)	34	54.61	11.23	43.37	38.34

Europe & Central Asia (excluding high income)	16	70.78	37.34	33.44	80.52
Latin America & Caribbean (excluding high income)	16	59.93	27.40	32.52	72.85
East Asia & Pacific (excluding high income)	8	66.57	34.19	32.37	77.07
Middle East & North Africa (excluding high income)	7	37.80	13.24	24.56	83.12

Table 6. Income-group summary statistics

Income_Group	Countries	Account_Ownership_pct	HF_Payment_Index_pct	Activation_Gap_pp	Internet_Users_pct
Low income	16	45.45	6.17	39.27	26
Lower middle income	37	55.13	17	38.13	57.41
Upper middle income	33	69.07	34.33	34.74	79.89

Figure 3. Regional variation in the activation gap



The regression comparison is reported in Table 7 and Figure 4. XGBoost delivered the best out-of-fold performance, with an RMSE of 5.32 percentage points, an MAE of 4.18, and an R^2 of 0.794. Ridge followed closely, indicating that the activation gap contains a substantial linear component, but the tree model still captured enough nonlinearity to improve forecast accuracy. Random Forest also performed well but did not match XGBoost. LightGBM trailed the other models in this sample. Figure 5 shows that the XGBoost predictions tracked actual country gaps closely across most of the support, with the largest errors concentrated among a small set of difficult, high-gap countries.

Table 7. Regression model comparison

Model	RMSE (pp)	MAE (pp)	R2	Mean fold RMSE (pp)	SD fold RMSE (pp)
XGBoost	5.32	4.18	0.79	5.46	0.99
Ridge	5.55	4.25	0.78	5.61	0.90
RandomForest	5.69	4.48	0.76	5.74	1.15
LightGBM	6.49	5.08	0.69	6.61	1.04

Figure 4. Regression model comparison by RMSE





Figure 5. Actual versus predicted activation gaps for XGBoost




Classification results reinforce the same ordering. As shown in Table 8 and Figure 6, XGBoost achieved the highest AUC at 0.863. Random Forest and LightGBM also produced useful classifiers, while the logistic baseline lagged. The best classifier correctly identified 30 of the 43 high-gap economies and 34 of the 43 lower-gap economies at the 0.5 threshold. That level of performance is strong for a small cross-country dataset and is sufficient for practical screening: a policy team can use the classifier to identify where a deeper country diagnostic is most urgent.

Table 8. High-gap classification model comparison

Model	AUC	Accuracy	F1	Recall	Precision	Balance d accuracy	Mean fold AUC	SD fold AUC
XGBoost	0.86	0.74	0.73	0.70	0.77	0.74	0.87	0.07
Random Forest	0.84	0.76	0.75	0.72	0.78	0.76	0.83	0.09
LightGBM	0.84	0.65	0.65	0.65	0.65	0.65	0.84	0.09
Logistic	0.75	0.67	0.67	0.67	0.67	0.67	0.76	0.09

Figure 6. ROC curves for high-gap classification



The lower-gap countries are also informative. Niger, Lebanon, Chad, Nicaragua, and Mauritania had much smaller observed gaps than the high-risk group. These cases do not imply complete digital maturity, because their underlying ownership and usage levels differ, but they do show stronger alignment between access and use. In practical terms, the ratio between ownership and payment behavior is tighter. That alignment suggests that some policy environments have already built more effective bridges from account opening to transaction activity, even when income levels remain modest. The benchmark is therefore realistic for peers: high-gap countries are not being compared only with rich economies, but with developing economies that have already narrowed the access-to-usage distance.

Table 9. XGBoost ablation study

Feature block	Features	RMSE (pp)	R2	AUC
Accounts only	3	6.44	0.70	0.87
Findex behavioral	16	5.05	0.81	0.87
Full + WDI	23	5.32	0.79	0.86

Model stability across folds was also strong. The mean fold RMSE and mean fold AUC were close to the final averaged out-of-fold performance, which indicates that the headline metrics were not driven by a few favorable partitions. This matters because country samples can be sensitive to one or two influential observations. The repeated cross-validation results show that the ranking of the main models remained stable across held-out splits: XGBoost consistently led, Ridge remained competitive in regression, and the logistic baseline consistently trailed the tree classifiers. That consistency increases confidence that the reported model ordering is a genuine property of the data rather than a resampling artifact.

The empirical findings also help refine the meaning of the “digital divide” in finance. The divide is often described as a problem of connectivity alone, yet the ablation and SHAP results show a more layered structure. Connectivity matters, but so do digitally delivered use cases, payment-channel design, and trust-capability variables. A country with middling internet penetration can still narrow the activation gap if government payments, merchant acceptance, and simple app-based routines are strong. Conversely, a country with broad connectivity can still underperform if digital payment behavior is not embedded in everyday institutions.

The ablation study in Table 9 yields an important substantive result. Account variables alone already explained a meaningful share of the activation gap, confirming that basic access remains informative. However, once behavioral and payments-rail variables were added, XGBoost regression performance improved substantially and the R^2 rose above 0.81. Adding the WDI internet and mobile indicators did not increase the headline predictive scores further in this sample. That result is still informative. It indicates that the richest predictive content for this task resides inside the payment-behavior and account-activation variables already collected by Global Findex. The WDI variables remain valuable as external context and as part of the country narratives, but the empirical signal is primarily behavioral rather than purely infrastructural.

SHAP interpretation clarifies which variables drove the model. Table 11 and Figure 7 show that account ownership was the single most influential feature, followed by internet use, mobile-money account ownership, card-linked government payment channels, and digital-skills indicators. This ranking is substantively coherent. A high level of account ownership raises the risk of a large activation gap when downstream payment use is weak, because the gap is defined as access minus repeated use. Internet use and mobile-account penetration matter because routine digital payments require a functioning pathway from access to transaction behavior. Government payment channels matter because cash-heavy transfer systems can suppress the frequency with which accounts are used after they are opened.

Country rankings show that the hardest transitions were concentrated but not uniform. Table 10 and Figure 8 identify Ghana, Uganda, and Zambia as the three largest observed gaps, followed closely by Kenya, India, Sri Lanka, and Senegal. These are not uniformly the countries with the lowest ownership rates. Several have already achieved broad account penetration. Their problem is that ownership has not yet translated into repeated digital payment behavior. At the opposite end of the distribution, countries such as Niger, Lebanon, Chad, Nicaragua, and Mauritania recorded much smaller gaps, indicating a tighter relationship between access and usage in those cases.

The automated policy notes add operational value because they distinguish among different kinds of high-gap countries. In the model-generated diagnostics, Zambia and Uganda were flagged mainly for connectivity and mobile-money activation bottlenecks. Ghana was flagged more strongly for merchant acceptance and app-based payment use. Senegal combined merchant-acceptance weaknesses with broader connectivity constraints. Sri Lanka appeared as a case where ownership was already high, but the next stage required stronger activation of mobile and account-based payment routines. These differences matter because they imply different policy packages. A blanket “open more accounts” strategy would be poorly targeted in each of these cases.

Three broader findings follow. First, digital financial inclusion is now better understood as a conversion problem than as a pure access problem in many developing economies. Second, payment-use variables materially outperform a narrow infrastructure-only view of the digital

divide. Third, an interpretable machine-learning workflow can produce country rankings that are accurate enough for triage while remaining transparent enough for policy use. The strongest use case is not automated decision-making in isolation, but rapid prioritization of where policymakers should concentrate detailed country diagnostics and implementation resources.

Table 10. Countries with the largest observed activation gaps

Country	ISO3	Region	Income group	Account ownership (%)	HF payment index (%)	Activation gap (pp)	Internet users (%)	Mobile subscriptions
Ghana	GHA	Sub-Saharan Africa (excluding high income)	Lower middle income	81.24	16.33	64.91	69.84	134
Uganda	UGA	Sub-Saharan Africa (excluding high income)	Low income	72.83	8.21	64.61	10	57.37
Zambia	ZMB	Sub-Saharan Africa (excluding high income)	Low income	72.70	11.82	60.88	31.23	96.41
Kenya	KEN	Sub-Saharan Africa (excluding high income)	Lower middle income	90.12	29.25	60.87	40.81	104
India	IND	South Asia (excluding high income)	Lower middle income	89.02	31.07	57.96	43.41	84.27
Sri Lanka	LKA	South Asia (excluding high income)	Lower middle income	81.69	24.62	57.08	50.11	144
Senegal	SEN	Sub-Saharan Africa (excluding high income)	Lower middle income	76.46	19.90	56.56	59.98	110

Gabon	GAB	Sub-Saharan Africa (excluding high income)	Upper middle income	68.18	11.70	56.48	73.70	138
Tanzania	TZA	Sub-Saharan Africa (excluding high income)	Lower middle income	59.83	4.62	55.21	31.89	82.21
Cote d'Ivoire	CIV	Sub-Saharan Africa (excluding high income)	Lower middle income	57.57	6.04	51.53	38.41	145

Figure 7. SHAP importance ranking for the best regression model




Table 11. Top SHAP features

Feature	Mean absolute SHAP
account t d	0.05
internet	0.02
mobileaccount t d	0.02
fing2p_card	0.02
con31a	0.01
fiaccount t d	0.01
con31d	0.01
con17b	0.00
con1	0.00
con26d	0.00

Figure 8. Countries with the largest observed activation gaps




Table 12. Example automated country policy diagnostics

Country	Predicted gap (pp)	High-gap probability (%)	Bottleneck 1	Bottleneck 2
Zambia	59.66	97.36	connectivity and handset-based internet access	cash-heavy government payment channels
Uganda	57.95	97.41	connectivity and handset-based internet access	mobile money and account activation

Ghana	56.03	95.57	merchant acceptance and app-based payment use	mobile money and account activation
Senegal	52.91	97.65	merchant acceptance and app-based payment use	connectivity and handset-based internet access
Cameroon	52.10	93.26	connectivity and handset-based internet access	mobile money and account activation

5. CONCLUSION

This paper measured the digital-finance activation gap in 86 developing economies by combining Global Findex 2025 country data with WDI internet and mobile indicators. The main empirical result is straightforward: countries have expanded account ownership much faster than they have expanded high-frequency digital payment behavior. The mean gap reached 37.04 percentage points, with the largest deficits concentrated in South Asia and Sub-Saharan Africa. XGBoost produced the strongest predictive performance in both regression and classification, and SHAP interpretation showed that the decisive factors were account activation, internet-enabled use, mobile-money penetration, payment-channel composition, and digital capabilities rather than ownership alone.

The policy implication is that the next generation of inclusion strategies should shift from pure onboarding targets toward activation targets. Governments should digitize transfer channels that remain cash-heavy; payment-system authorities should expand low-cost merchant acceptance and interoperability; providers should simplify app-based payment journeys and strengthen cash-in/cash-out reliability; and regulators should pair expansion with fraud redress, consumer protection, and digital-skills support. These recommendations differ across countries, which is why the paper’s model-to-text diagnostic layer is useful.

The paper still has clear boundaries. The results are predictive rather than causal, and the WDI auxiliary indicators are measured at their latest available country years rather than exactly the same survey year for every economy. Those limitations do not weaken the main contribution, which is comparative identification of the hardest access-to-usage transitions using transparent, measured evidence. Future work can extend the framework with panel designs, causal policy evaluation, and live LLM deployment for interactive country diagnostics, but the present paper already establishes a reproducible benchmark for where the digital financial inclusion frontier is currently most difficult.

The manuscript also passes the revision test implied by the common publication critique of “illustrative” evidence. Every experimental claim in the paper is tied to measured cross-validation outputs, every figure and table is generated from the packaged code, and the narrative is stated definitively rather than hypothetically. The broader conclusion is that the central inclusion frontier in many developing economies is no longer whether households can

obtain an account. It is whether that account becomes a routine digital payment instrument embedded in everyday economic life.

REFERENCES

- Aker, J. C., & Mbiti, I. M. (2010). Mobile phones and economic development in Africa. *Journal of Economic Perspectives*, 24(3), 207-232.
- Allen, F., Demircug-Kunt, A., Klapper, L., & Martinez Peria, M. S. (2016). The foundations of financial inclusion: Understanding ownership and use of formal accounts. *Journal of Financial Intermediation*, 27, 1-30.
- Aron, J. (2018). Mobile money and the economy: A review of the evidence. *The World Bank Research Observer*, 33(2), 135-188.
- Beck, T., Demircug-Kunt, A., & Levine, R. (2007). Finance, inequality and the poor. *Journal of Economic Growth*, 12(1), 27-49.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5-32.
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785-794). ACM.
- Demircug-Kunt, A., Klapper, L., Singer, D., Ansar, S., & Hess, J. (2018). *The Global Findex Database 2017: Measuring financial inclusion and the fintech revolution*. World Bank.
- Demircug-Kunt, A., Klapper, L., Singer, D., & Ansar, S. (2022). *The Global Findex Database 2021: Financial inclusion, digital payments, and resilience in the age of COVID-19*. World Bank.
- Donou-Adonsou, F., & Sylwester, K. (2016). ICT, financial development, and poverty reduction: Evidence from developing countries. *Telecommunications Policy*, 40(11), 1045-1053.
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5), 1189-1232.
- GSMA. (2023). *State of the industry report on mobile money 2023*. GSMA.
- GSMA. (2024). *State of the industry report on mobile money 2024*. GSMA.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). Springer.
- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2021). *An introduction to statistical learning: With applications in R* (2nd ed.). Springer.
- Jack, W., & Suri, T. (2014). Risk sharing and transactions costs: Evidence from Kenya's mobile money revolution. *American Economic Review*, 104(1), 183-223.
- Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T.-Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems 30* (pp. 3146-3154).
- Klapper, L., El-Zoghbi, M., & Hess, J. (2016). *Achieving the Sustainable Development Goals: The role of financial inclusion*. CGAP.
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30* (pp. 4765-4774).
- Ozili, P. K. (2018). Impact of digital finance on financial inclusion and stability. *Borsa Istanbul Review*, 18(4), 329-340.

- Philippon, T. (2016). The fintech opportunity. National Bureau of Economic Research Working Paper No. 22476.
- Sahay, R., von Allmen, U. E., Lahreche, A., Khera, P., Ogawa, S., Bazarbash, M., & Beaton, K. (2020). The promise of fintech: Financial inclusion in the post COVID-19 era. International Monetary Fund.
- Suri, T., & Jack, W. (2016). The long-run poverty and gender impacts of mobile money. *Science*, 354(6317), 1288-1292.
- World Bank. (2014). Global financial development report 2014: Financial inclusion. World Bank.
- World Bank. (2016). World development report 2016: Digital dividends. World Bank.
- World Bank. (2021). World development report 2021: Data for better lives. World Bank.