Inequality of Opportunity in Territorial Financial Access: Evidence from Peru, 2002–2025

José Antonio Mendoza Sánchez*

May 22, 2026

Abstract

Peru's banking agents reached 95.1% of districts by 2025, yet only 27.6% have a full-service financial office. I apply the Human Opportunity Index framework to 1,873 Peruvian districts observed from 2002 to 2025, estimating the D-index and decomposing it via Shapley values across six territorial circumstances. The office D-index falls from 0.619 to 0.300 but plateaus after 2012; the any-access D-index falls to 0.021. Density is the largest contributor to office-level inequality (28.1%), followed by poverty (21.3%), natural region (15.8%), altitude (13.8%), remoteness (11.5%), and indigenous language (9.5%). The Peruvian expansion equalized basic transactional access through agents while leaving full-service financial infrastructure concentrated in denser, less-poor, lower-altitude territories.

Keywords: inequality of opportunity; financial inclusion; banking deserts; Human Opportunity Index; D-index; Peru.

JEL codes: D63, G21, O16, R12.

1 Introduction

A large literature links formal finance to economic mobility. Galor and Zeira [1993] show theoretically that credit constraints can perpetuate inequality across generations. Jayaratne and Strahan [1996] demonstrate that branch deregulation in the United States accelerated economic growth. Demirguc-Kunt and Levine [2009] argue that well-functioning financial systems can expand opportunity by weakening the link between inherited wealth and economic outcomes. Yet financial inclusion depends not only on individual demand or financial literacy. It also depends on the territorial distribution of financial infrastructure. In countries with large geographic, ethnic, and income disparities—such as Peru, with its coastal cities, Andean highlands, and Amazonian lowlands—the place where people live can determine whether basic financial services are physically available at all [Beck et al., 2007].

*Comments welcome. Correspondence: jose.mendozas@pucp.edu.pe.

Peru offers a sharp case for studying this question. Its 1,873 districts span coastal cities at sea level, Quechua-speaking highland communities above 4,000 meters, and Amazonian settlements reachable only by river. Over the past two decades, Peru's financial system expanded substantially—from roughly 150 districts with a financial office in 2002 to over 500 by 2025—yet most of this geography remained unserved by full-service institutions.[1] I use this setting to measure inequality of opportunity (IOp) in territorial financial access, asking: to what extent is access to a financial office predicted by predetermined territorial circumstances rather than being distributed independently of them?

The conceptual starting point is the equality-of-opportunity framework developed by Roemer [1998] and synthesized by Roemer and Trannoy [2016], which decomposes inequality into a component attributable to circumstances beyond individual control and a component attributable to effort or choice. In the territorial context, the six circumstances listed above—poverty, altitude, density, remoteness, natural region, and ethnolinguistic composition—are inherited or historically persistent features of local opportunity structures, not individual effort variables. If they systematically predict whether a district has formal financial infrastructure, then financial access is an opportunity distributed unequally across circumstance-defined types. The measurement strategy uses the Human Opportunity Index (HOI) framework developed by Paes de Barros et al. [2009] for binary access outcomes in Latin America: I estimate a logit model of access, compute the D-index from predicted probabilities, and decompose it via Shapley values to identify which circumstances contribute most to unequal access.

Three arguments support treating territorial financial access as an opportunity in the normative sense. First, Peru's National Financial Inclusion Policy [Gobierno del Perú, 2019] explicitly treats financial access as a policy objective analogous to basic services such as water, electricity, and education. The SBS publishes district-level inclusion indicators and monitors financial deserts as a policy problem, not merely a market outcome. Second, a body of causal evidence shows that physical access to financial institutions generates real effects on saving, borrowing, and economic activity. Burgess and Pande [2005] show that India's social banking program reduced poverty in districts that received branches. Bruhn and Love [2014] exploit the entry of Banco Azteca in Mexico and find effects on employment and income, especially in low-access areas. Third, the HOI framework was designed precisely for basic services that are preconditions for participation: water, electricity, schooling, sanitation [Paes de Barros et al., 2009, Molinas Vega et al., 2010]. Financial access shares this structure—it is a precondition for formal economic life, not a final outcome such as income or wealth.

A second outcome—whether the district has any financial access point, including banking agents (*cajeros corresponsales*)—makes visible the distinction between basic access and full-service access. Agents, introduced under SBS regulatory frameworks around 2005–2008 [Superintendencia de Banca, Seguros y AFP del Peru, 2013], expanded into bodegas, pharmacies, and retail outlets

[1] Throughout, "financial office" refers to any office of a regulated institution: commercial banks, CMACs, CRACs, EDPYMES, financieras, and Banco de la Nación. "Banking agent" refers to *cajeros corresponsales* of any regulated institution.

across the territory. By 2025, 95.1 percent of districts have at least one access point, and the any-access D-index falls to 0.021—near zero. But the office D-index plateaued at 0.300 after 2012. Peru is not a simple case of financial exclusion: it is a case of layered expansion, where agents nearly solved the extensive margin while leaving full-service infrastructure—credit evaluation, term deposits, advisory services—concentrated in denser, lower-altitude, less-poor districts.

For the IOp literature, the Peruvian data make it possible to apply the HOI framework to financial infrastructure—a domain where the D-index reveals inequality levels well above those documented for water, electricity, or schooling. For the financial inclusion literature, the analysis reframes the standard question: not just how much access exists, but how equally it is distributed across circumstance-defined groups, and how that distribution changed as Peru's system expanded over two decades.

The rest of the paper is organized as follows. Section 2 reviews the related literature. Section 3 describes the data. Section 4 presents the empirical framework. Section 5 reports results and robustness checks. Section 6 discusses implications and limitations. Section 7 concludes.

2 Related Literature

Three bodies of work inform the analysis: the measurement of inequality of opportunity, the Human Opportunity Index for binary access outcomes, and the empirical literature on physical financial access and banking deserts. I review each and identify where the Peruvian financial-access data open a space not covered by existing work.

The theory of equality of opportunity, developed by Roemer [1998] and synthesized by Roemer and Trannoy [2016], distinguishes between inequality that arises from circumstances beyond individual control and inequality that arises from effort or choice. In empirical applications, researchers partition observed inequality into a component attributable to predetermined circumstances—such as parental education, gender, race, or birthplace—and a residual. The circumstance-related component is interpreted as inequality of opportunity. A key insight from Ferreira and Gignoux [2011] is that empirical estimates capture only a partial picture of total IOp, because the researcher observes a subset of all relevant circumstances. For continuous outcomes measured with decomposable inequality indices (such as the mean log deviation), adding circumstances weakly increases measured IOp, yielding a formal lower-bound property. This result has made the "lower bound" interpretation standard in the literature, though as I discuss below, it does not automatically extend to all settings.

The empirical IOp literature has grown rapidly since the pioneering work of Bourguignon et al. [2007], who use parametric decomposition to measure IOp in Brazil, finding that parental education and race are the dominant circumstances. Lefranc et al. [2009] develop non-parametric stochastic-dominance tests and apply them to income in France, offering an alternative to the parametric approach that avoids distributional assumptions but requires stronger data. Checchi and Peragine [2010] provide an influential application to Italy, comparing parametric and non-parametric approaches and finding substantial IOp in earnings. Brunori et al. [2013] extend the analysis to cross-country comparisons, documenting that IOp accounts for a significant share of total inequality in most countries and that higher IOp is associated with lower intergenerational mobility. A common thread across these studies is that IOp is measured at the individual level, with circumstances defined at the individual or household level. I depart from this convention by using the *district* as the unit of observation and territorial characteristics as circumstances—an extension that is defensible (place of residence shapes the opportunity set available to individuals) but that requires careful interpretation, as discussed in Section 6.

The literature distinguishes between ex-ante and ex-post approaches to IOp measurement [Fleurbaey and Peragine, 2013, Ramos and Van de gaer, 2016]. The ex-ante approach compares expected outcomes across types (groups defined by circumstances) before effort is realized. The ex-post approach compares individuals at similar effort levels across types. In the territorial setting of this paper, individual effort is not observed—there is no district-level analogue of individual labor supply or savings decisions—and the unit of observation is the district itself. I therefore adopt an ex-ante approach and interpret the D-index as measuring how much predicted access varies across circumstance-defined groups, without conditioning on effort. Because the outcome here is binary and the measure is the D-index, the monotonicity/lower-bound property established by Ferreira and Gignoux [2011] for continuous outcomes does not automatically apply. I therefore interpret the estimates as *partial* measures of circumstance-related inequality—informative about the six observed dimensions but not necessarily a strict lower bound on total IOp.

The Human Opportunity Index (HOI) was developed specifically for binary access outcomes in Latin America and the Caribbean [Paes de Barros et al., 2009]. The motivation was to measure whether basic services—schooling, clean water, sanitation, electricity—are distributed as opportunities available to all children regardless of their circumstances, or whether access depends on factors such as parental income, location, or household composition. The HOI combines two components: average coverage (the share of the population with access) and a penalty for unequal distribution across circumstance groups (the D-index). When coverage is high and equally distributed, the HOI approaches 100; when coverage is low or concentrated in advantaged groups, the HOI is penalized. Molinas Vega et al. [2010] extend the framework to a broader set of LAC countries and report D-index benchmarks: electricity and primary-school attendance have D-index values of 0.01–0.05 in most countries (near-universal and equally distributed), while sanitation—the most unequally distributed basic service—has D-index values of 0.10–0.20. A notable gap in this literature is that to my knowledge, the HOI has not been applied to territorial financial-access infrastructure. This is surprising, given that access to formal finance shares the structural characteristics of the services for which the HOI was designed: it is a binary precondition for participation (having vs. not having a financial institution nearby), it varies systematically across circumstance groups, and it is recognized as a policy objective by governments across LAC. The Peruvian SBS data—covering all regulated institutions at the district level for 24 years—make it possible to apply the HOI to this domain for the first time.

A separate and well-established literature documents the importance of physical access to financial institutions. At the theoretical level, Galor and Zeira [1993] show that credit constraints can perpetuate inequality across generations by preventing talented but poor individuals from investing in human capital. Demirguc-Kunt and Levine [2008] and Claessens and Perotti [2007] review the empirical channels connecting financial development to inequality, including access to credit, insurance, and payment services. On the causal side, several influential studies exploit quasi-experimental variation in branch presence. Burgess and Pande [2005] study India's social banking experiment, in which the Reserve Bank of India mandated branch expansion into previously unbanked rural areas between 1977 and 1990, and find that the program reduced rural poverty and increased non-agricultural output. Bruhn and Love [2014] exploit the simultaneous opening of Banco Azteca branches in over 800 pre-existing Elektra retail stores across Mexico in 2002 and find significant effects on informal employment, income, and new business formation, especially in municipalities with low prior bank penetration. Jayaratne and Strahan [1996] show that US states that relaxed bank branching restrictions experienced faster economic growth. More recently, Nguyen [2019] provides evidence that credit markets remain local even in the age of digital banking: branch closings in US neighborhoods led to significant reductions in small-business lending, with effects persisting for years after closure. These findings motivate the focus on physical infrastructure in this paper: even as digital channels expand, the evidence suggests that physical presence—offices and agents—remains a determinant of effective financial access, particularly in developing countries where digital adoption is incomplete and cash remains dominant.

The banking deserts literature documents the uneven geography of bank offices, primarily in the United States. Hegerty [2016] maps commercial bank locations across US metropolitan areas and identifies banking deserts—areas with no branch within a reasonable distance. Ergungor [2010] links branch presence to credit access in low- and moderate-income neighborhoods, finding that proximity to a branch significantly predicts mortgage lending. This literature typically defines "deserts" by physical distance (travel time or radius) rather than administrative boundaries, and it focuses on branch closures in mature financial systems rather than on the expansion frontier in developing countries. Importantly, the banking deserts literature does not frame access through an inequality-of-opportunity lens: it documents spatial gaps without asking how much of the gap is attributable to predetermined circumstances versus other factors.

What connects these three bodies of work is the Peruvian case: the IOp/HOI framework provides the measurement tool, the causal evidence on branch presence motivates treating physical access as an opportunity, and the banking deserts concept defines the policy problem. Peru's SBS district-level data, covering all regulated institutions from 2002 to 2025, make it possible to document how the distribution of financial access across circumstance-defined groups evolved as the system expanded—first through offices, then through agents.

3 Data

The empirical dataset is a district-year panel for Peru covering 2002–2025. The core financial-access variables come from administrative data reported by Peru's financial supervisory authority, the Superintendencia de Banca, Seguros y AFP [SBS; Superintendencia de Banca, Seguros y AFP del Peru, 2026]. The SBS data cover *all* regulated financial institutions: commercial banks (*banca múltiple*), municipal savings banks (*cajas municipales de ahorro y crédito*, CMACs), rural savings banks (*cajas rurales*, CRACs), small-enterprise development institutions (EDPYMES), finance companies (*financieras*), and the state-owned Banco de la Nación. This comprehensive coverage is important: in 2025, only 257 districts have a commercial bank office, but 517 districts have at least one office when all regulated institutions are included. The CMACs (such as Caja Huancayo, Caja Arequipa, and Caja Piura) and Banco de la Nación are particularly important for territorial reach, operating in many highland and rural districts where commercial banks are absent.

Two binary outcomes are constructed. The first—**financial office**—equals one if district d has at least one office of any regulated financial institution in year t. The second—**any access**—equals one if the district has at least one financial office *or* banking agent (*cajero corresponsal*) in year t. Banking agents were introduced in Peru around 2005–2008 under the regulatory framework established by SBS Resolution 6285-2013 [Superintendencia de Banca, Seguros y AFP del Peru, 2013]. Under this framework, agents can process deposits, withdrawals, loan disbursements and repayments, utility payments, government transfers (including the Juntos and Pensión 65 conditional cash transfer programs), and simplified account openings. They do not typically provide credit evaluation, safe-deposit services, or complex advisory products. Agent data come from the SBS district-level agent panel, compiled from annual inclusion-finance reports.

The data do not include *cooperativas de ahorro y crédito* (COOPAC), which have been supervised by SBS since Law 30822 (2018) but are not yet integrated into the SBS financial-access panel used here.

Six district-level circumstance variables are merged from official sources. The *poverty rate* comes from the 2018 INEI district poverty map. *Altitude* in meters above sea level comes from georeferenced census coordinates. *Population density* is computed from 2017 census population and district area. *Remoteness* is proxied by the Haversine (great-circle) distance to the largest district in each department.[2] *Natural region* is classified as Coast, Sierra, or Amazon from the SBS panel. *Indigenous-language share* is the proportion of individuals whose mother tongue is Quechua, Aymara, or another native language, pooled from ENAHO household surveys 2012–2023 and aggregated to the district level; 240 of 1,873 districts required province-level imputation due to limited ENAHO coverage (see Appendix Table A8 for full documentation). Altitude and natural region are correlated by construction (Peru's Coast, Sierra, and Amazon zones correspond roughly to altitude bands), but both are retained because natural region captures ecological and institutional differences beyond elevation. The Shapley decomposition handles correlation by averaging marginal

[2]Haversine distance is an imperfect proxy for actual accessibility, especially in the Andes and Amazon where road distance and travel time differ substantially from great-circle distance.

Table 1: Empirical dataset

Component	Description
Unit of analysis	District-year
Period	2002–2025
Sample	1,873 districts; 44,952 district-years
Main outcome	At least one financial office
Secondary outcome	At least one access point (office or agent)
Source	SBS administrative data (all regulated institutions)
Circumstances	Poverty, altitude, density, remoteness, natural region, indigenous language

Table 2: District characteristics by financial-office access, 2025

Group	N	Poverty	Pop.	Density	Alt.	Dist.	Indig.	Sierra	Selva
No office	1,356	37.4	5,006	74.9	2,477	95.0	36.0	71.2	15.9
At least one office	517	25.5	49,072	1395.3	1,581	75.5	21.4	43.9	18.2

Notes: Poverty is district poverty rate (%). Pop. and density from 2017 census. Alt. in meters. Dist. in km. Indig. is share with native mother tongue (pooled ENAHO 2012–2023). Sierra and Selva are % of districts.

contributions over all possible orderings of entry.

The raw SBS panel contains 75 duplicate district-year observations concentrated in 2008–2013, which are dropped retaining the first occurrence. The complete-case estimation sample contains 1,873 districts and 44,952 district-year observations. Table 1 summarizes the dataset. Table 2 compares districts with and without a financial office in 2025. Districts without an office are poorer (37.4 vs. 25.5 percent poverty rate), much smaller (5,006 vs. 49,072 mean population), less dense (74.9 vs. 1,395 per km^2), located at higher altitude (2,477 vs. 1,581 meters), more remote (95.0 vs. 75.5 km to departmental hub), have higher indigenous-language shares (36.0 vs. 21.4 percent), and are more likely to be in the Sierra (71.2 vs. 43.9 percent).

4 Empirical Framework

The empirical strategy combines three elements: a logit model that maps district circumstances into predicted access probabilities, the D-index that summarizes how unequally those probabilities are distributed, and a Shapley decomposition that attributes the D-index to individual circumstances. This section describes each element and discusses the key methodological choices.

Access model. For each benchmark year (2002, 2007, 2012, 2017, 2022, 2025), the probability that district d has at least one financial office is modeled as a logit:

$$\Pr(Y_{dt} = 1 \mid C_d) = \Lambda(\alpha_t + C_d' \beta_t), \tag{1}$$

where Y_{dt} is the binary outcome, C_d is the vector of observed circumstances, and $\Lambda(\cdot)$ is the logistic cumulative distribution function. The circumstance vector includes poverty rate, log altitude, log density, log distance to the departmental hub, and indigenous-language share—all standardized to mean zero, unit variance across the pooled panel—plus natural-region dummies (Coast and Amazon, with Sierra as the omitted category). The estimated probability $\hat{p}_{dt}$ represents the predicted access probability associated with district d's circumstances in year t.

The logit is estimated separately for each benchmark year rather than pooled across time. This allows the coefficients—and thus the mapping from circumstances to predicted probabilities—to evolve over time, which is appropriate because the relationship between, say, density and bank presence may change as the financial system expands. A practical consequence is that "types" (groups of districts with similar predicted probabilities) are defined year-specifically: a district classified as high-probability in 2002 may not retain that status in 2025 if the coefficients shift. The logit serves an instrumental role: its coefficients are not interpreted causally, and the model's purpose is to generate predicted probabilities that feed into the D-index. Model fit is assessed via the area under the ROC curve (AUC) and McFadden's pseudo-R^2 (reported in the Results section and Appendix).

D-index and Human Opportunity Index. The D-index, developed by Paes de Barros et al. [2009], summarizes how unequally access is distributed across circumstance-defined groups. Let $\bar{p}_t$ denote average access coverage in year t and w_d district weights summing to one. The D-index is:

$$D_t = \frac{1}{2\bar{p}_t} \sum_d w_d \, |\hat{p}_{dt} - \bar{p}_t|. \tag{2}$$

Intuitively, D_t measures the fraction of access opportunities that would need to be reallocated across circumstance groups to achieve equal access. If all groups have the same predicted probability, $D = 0$. If access is perfectly sorted by circumstances, D approaches its maximum. The Human Opportunity Index combines coverage with this penalty:

$$HOI_t = \bar{p}_t(1 - D_t). \tag{3}$$

The HOI rises when coverage increases *and* when that increase is distributed more equally across circumstance groups. This distinction is central to the paper: aggregate coverage statistics cannot differentiate between broad-based expansion and expansion concentrated in already-advantaged territories. The D-index can.

Two weighting schemes capture distinct normative objects. In the *unweighted* specification ($w_d = 1/N$), each territory is treated as an opportunity unit. This perspective asks: does the financial system reach all of Peru's administrative geography equally? A small district of 500 inhabitants and a large district of 500,000 receive equal weight. In the *population-weighted* specification ($w_d \propto$ district population, with the logit also estimated using population weights for consistency), each person is the unit. This asks: does the probability of living in a served district depend on circumstances? The unweighted D-index emphasizes territorial equity; the population-weighted D-index emphasizes the share of the population affected by circumstance-related inequality. Both perspectives are policy-relevant, and I report both.

Shapley decomposition. To identify which circumstances contribute most to the D-index, I use a Shapley-value decomposition [Wendelspiess Chavez Juarez and Soloaga, 2014]. Let S denote a subset of circumstances and $D(S)$ the D-index obtained from a logit model using only the circumstances in S. The Shapley contribution of circumstance k is:

$$\phi_k = \sum_{S \subseteq K \setminus \{k\}} \frac{|S|!(|K| - |S| - 1)!}{|K|!} \left[D(S \cup \{k\}) - D(S) \right], \tag{4}$$

where $|K|$ is the total number of circumstance groups. Each circumstance's contribution is its average marginal effect on the D-index across all possible subsets of other circumstances. The Shapley decomposition is exact—contributions sum to the full D-index—and path-independent: unlike sequential decompositions, it does not depend on the order in which variables are entered.

With six circumstance groups, the computation requires estimating $2^6 = 64$ logit models per benchmark year. Natural region enters as two dummies (Coast and Amazon) but is treated as a single circumstance group in the Shapley decomposition: its contribution is computed by adding or removing both dummies jointly. This grouping is standard in the HOI literature and avoids inflating the degrees of freedom attributed to a single conceptual variable.

A well-known property of the Shapley decomposition is that when circumstances are correlated, individual contributions reflect both direct effects and shared variance. In the present application, density and remoteness are correlated at -0.60, and altitude and indigenous language at 0.45 (Table A1 in the Appendix). The Shapley values are well-defined regardless of correlation—they provide a unique, fair attribution—but the interpretation should acknowledge that, for example, density's 28.1 percent share partly reflects variance shared with population and remoteness. I address this concern through a robustness exercise that separates geographic circumstances (altitude, remoteness, natural region) from potentially endogenous ones (poverty, density, indigenous language), allowing the reader to assess how much of the D-index is attributable to hard geography alone.

Inference. Standard errors and confidence intervals for the D-index and Shapley shares are obtained by bootstrap. The baseline specification uses 999 district-level iid replications with replacement. Because districts within the same department may share unobserved spatial shocks, I also report a department-level cluster bootstrap (25 clusters, 999 replications). Both methods are implemented in the replication code with documented random seeds (42 for iid, 43 for cluster); zero bootstrap replications failed in either specification.

Table 3: Logit model of district financial-office access, 2025

Variable	Coef.	SE
Poverty rate (z)	-0.435	0.072
Altitude (log, z)	-0.178	0.102
Density (log, z)	0.828	0.082
Distance to dept. hub (log, z)	0.024	0.070
Coast	0.288	0.262
Amazon	0.526	0.204
Indigenous language share (z)	-0.043	0.074
Constant	-1.279	0.102
Observations	1,873	
Coverage	0.276	
D-index	0.300	
HOI	0.193	
Pseudo-R^2	0.190	
AUC	0.782	

Notes: Dependent variable = 1 if district has at least one financial office. Continuous covariates standardized. Sierra omitted. Log-odds coefficients.

5 Results

Table 3 reports the logit model for 2025. Denser districts are significantly more likely to have a financial office (coefficient 0.828, SE 0.082), consistent with the market-size logic of branch location: institutions open offices where the customer base justifies fixed costs. Poorer districts are less likely to have an office (-0.435, SE 0.072). Altitude has a negative but imprecisely estimated coefficient (-0.178, SE 0.102). The indigenous-language share is small and not individually significant (-0.043, SE 0.074), but contributes through interaction with other circumstances in the Shapley decomposition. Conditional on other covariates, Amazon districts have a positive coefficient relative to Sierra (0.526, SE 0.204), likely reflecting the hub role of Amazon provincial capitals. The model achieves an AUC of 0.782 and a McFadden pseudo-R^2 of 0.190 (Appendix Table A6 reports average marginal effects and diagnostics for 2002, 2012, and 2025; the AUC is 0.922 in 2002, when offices were rare and highly concentrated, declining mechanically as coverage expands).

Table 4 shows the evolution of coverage, the D-index, and the HOI for financial offices across the six benchmark years. Coverage rises from 8.0 percent of districts in 2002 to 27.6 percent in 2025. The D-index falls sharply from 0.619 to 0.304 between 2002 and 2012—a period of rapid office expansion concentrated in provincial capitals and fast-growing urban districts. After 2012, the D-index plateaus around 0.300, and coverage barely increases. Bootstrap confidence intervals confirm this formally: the 2012 D-index is 0.304 (95% CI [0.276, 0.339]) and the 2025 D-index is 0.300 (95% CI [0.275, 0.330]); the difference of 0.004 is statistically indistinguishable from zero. The HOI rises from 0.030 to 0.193. This plateau is a central finding: after 2012, financial offices essentially stopped reaching new districts, and the remaining inequality of opportunity in office access became frozen. Figure 4 visualizes this by showing the net change in districts with access

Table 4: Coverage, D-index, and HOI for financial offices

Year	N	Coverage	Desert	D-index	HOI
2002	1,873	0.080	0.920	0.619	0.030
2007	1,873	0.120	0.880	0.493	0.061
2012	1,873	0.251	0.749	0.304	0.175
2017	1,873	0.263	0.737	0.299	0.185
2022	1,873	0.275	0.725	0.292	0.195
2025	1,873	0.276	0.724	0.300	0.193

Notes: D-index from year-specific logit with six circumstances. HOI = coverage $\times$ $(1 - \text{D})$.

per year—office entries were concentrated in 2003–2012 and effectively ceased thereafter. The plateau coincided with consolidation in the banking sector, the full rollout of the agent regulatory framework [Superintendencia de Banca, Seguros y AFP del Peru, 2013], and the maturation of branch networks.

The introduction of banking agents transforms the picture. Table 5 compares the D-index across the two outcomes. Before agents were introduced, the two series are identical: both outcomes track office presence. After 2008, the trajectories diverge sharply. For the "any access" outcome, coverage rises from 8.0 percent in 2002 to **95.1 percent** in 2025—near-universal territorial coverage. The D-index falls to **0.021**, indicating that access to at least one financial service point is distributed almost independently of territorial circumstances. The HOI reaches 0.931. Figure 1 visualizes the divergence: the office D-index freezes after 2012 at approximately 0.30, while the any-access D-index continues falling toward zero through 2025.

Figure 2 shows the evolution of three access tiers. In 2008, roughly 80 percent of districts had no formal financial presence. By 2025, the landscape has reversed: 27.6 percent of districts have a financial office (Tier 1), 67.5 percent have a banking agent but no office (Tier 2), and only 4.9 percent—91 districts—remain financial deserts with no regulated access point of any kind (Tier 3). Figure 4 shows the dynamics underlying this transformation: office entries were concentrated in 2003–2012 and effectively ceased thereafter, while agent entries continued through 2020. Figure 3 reveals the geographic distribution of the three tiers: offices concentrate along the coast and in departmental capitals, agents fill the intermediate Sierra and Amazon districts, and the remaining deserts cluster in the highest-altitude and most remote territories.

This two-tier pattern is the paper's central empirical finding. It means that Peru's financial infrastructure expansion largely resolved the *extensive margin* of basic access—the question of whether a district has any financial service point at all—but left the *intensive margin* of service quality unresolved. Districts served only by agents can process basic transactions, but their residents cannot access in-person credit evaluation, open complex savings products, or receive financial advisory services without traveling to a district with a full-service office.

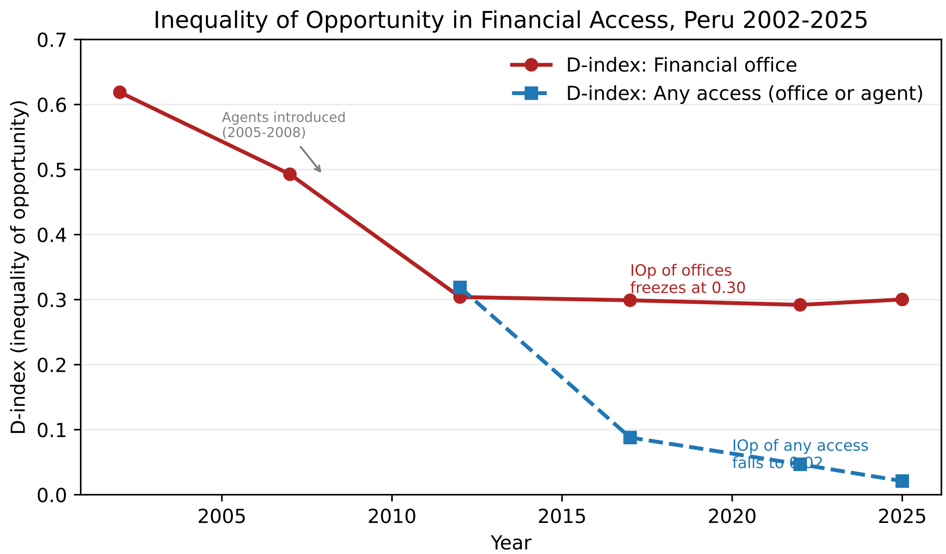

Figure 1: D-index for financial offices and any financial access, 2002–2025

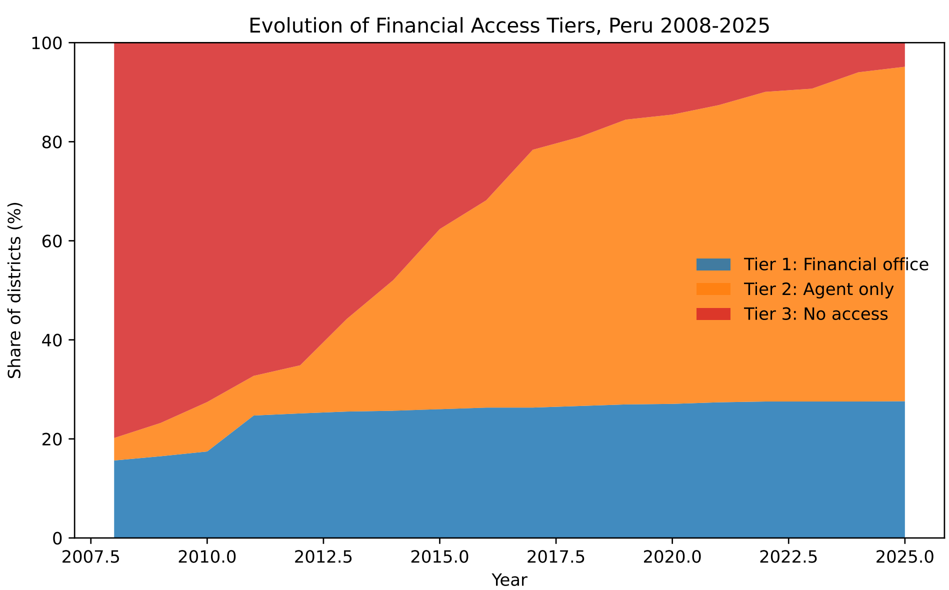

Figure 2: Evolution of financial access tiers, Peru 2008–2025

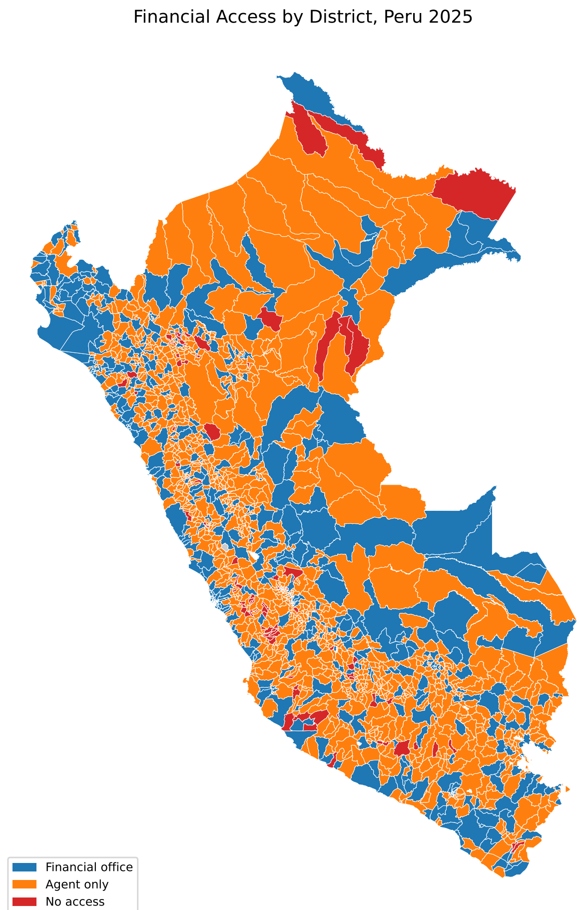

Figure 3: Financial access by district, Peru 2025. Blue: financial office. Orange: agent only. Red: no access.

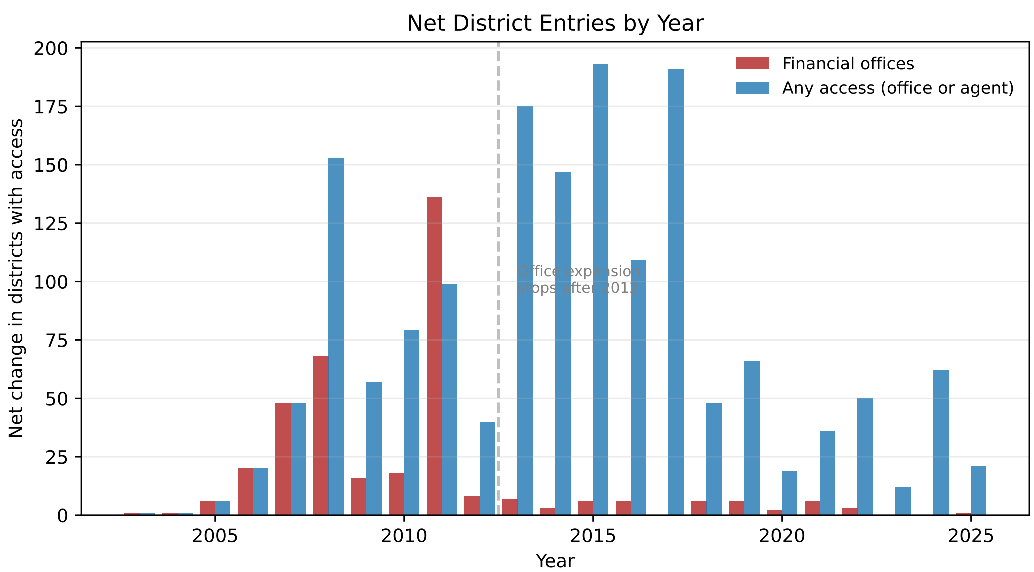

Figure 4: Net change in districts with financial access per year

Table 5: D-index comparison: financial offices vs. any access

Year	Cov.	D	HOI	Cov.	D	HOI
2002	0.080	0.619	0.030	0.080	0.619	0.030
2007	0.120	0.493	0.061	0.120	0.493	0.061
2012	0.251	0.304	0.175	0.349	0.319	0.238
2017	0.263	0.299	0.185	0.784	0.088	0.715
2022	0.275	0.292	0.195	0.901	0.047	0.859
2025	0.276	0.300	0.193	0.951	0.021	0.931

Notes: Left panel: financial offices. Right panel: any access (office or agent). Before agents were introduced (∼2008), the two outcomes are identical.

Table 6 decomposes the 2025 office D-index into the contribution of each circumstance. Density is the largest contributor at 28.1 percent of the total D-index. This is consistent with the commercial logic of office location: financial institutions—whether commercial banks, CMACs, or financieras—face fixed costs (rent, staff, security, compliance) that require a minimum customer base to cover. Districts below this threshold simply do not attract offices, regardless of other characteristics. Poverty contributes 21.3 percent, capturing the joint effect of lower demand for financial services and lower profitability in poor areas. Natural region accounts for 15.8 percent, reflecting ecological and institutional differences across Peru's three macro-regions: coastal districts benefit from better infrastructure and market integration, while Sierra and Amazon districts face higher operational costs. Altitude contributes 13.8 percent, remoteness 11.5 percent, and indigenous-language composition 9.5 percent. The indigenous-language coefficient in the logit is small and not individually significant (−0.043, SE 0.074), but the Shapley decomposition captures its contribution through correlation with other circumstances: districts with higher indigenous-language shares tend to be poorer, higher-altitude, and more remote, and the Shapley procedure attributes to indigenous language the portion of the D-index that is uniquely or jointly explained by this variable across all possible orderings.

The any-access decomposition (Figure 5) reveals a strikingly different pattern. For the near-zero D-index of 0.021, geographic factors—remoteness, altitude, and natural region—become the dominant contributors, while density and poverty shrink in importance. This contrast is economically interpretable. Offices follow market logic: density and poverty determine whether fixed costs can be recovered. Agents, by contrast, have much lower fixed costs (a POS terminal in an existing retail establishment) and can operate in smaller, poorer districts. Once agents have reached into the territory, the only districts that remain unserved are those where geography itself is the barrier: extremely high altitude, extreme remoteness, or deep Amazon isolation. The Shapley decomposition thus tells a story about the *technology* of financial access: heavier infrastructure (offices) is sorted by market conditions, while lighter infrastructure (agents) is sorted by geography.

Figure 6 shows that the composition of office-level inequality is remarkably stable across the 2002–2025 period. Density's share rises gradually from about 25 to 28 percent, while poverty's share declines from about 29 to 21 percent—consistent with a process in which offices expanded

Table 6: Shapley decomposition of the office D-index, 2025

Circumstance	Contribution	Share (%)
Poverty	0.064	21.3
Altitude	0.042	13.8
Density	0.084	28.1
Remoteness	0.034	11.5
Natural region	0.047	15.8
Indigenous language	0.029	9.5

Notes: Contributions sum to D = 0.300. Shares are percentages of total measured circumstance-related inequality.

first into less-poor districts, reducing poverty's relative contribution. Altitude, natural region, and remoteness maintain roughly constant shares throughout. This stability suggests that the *structure* of circumstance-related inequality in office access has not changed even as its level has fallen.

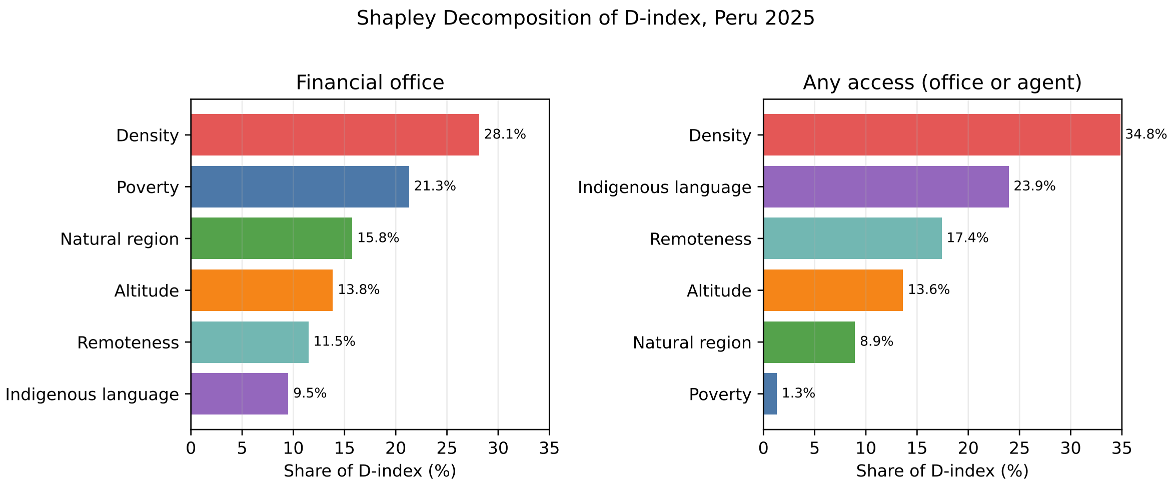

Figure 5: Shapley decomposition: financial offices vs. any access, 2025

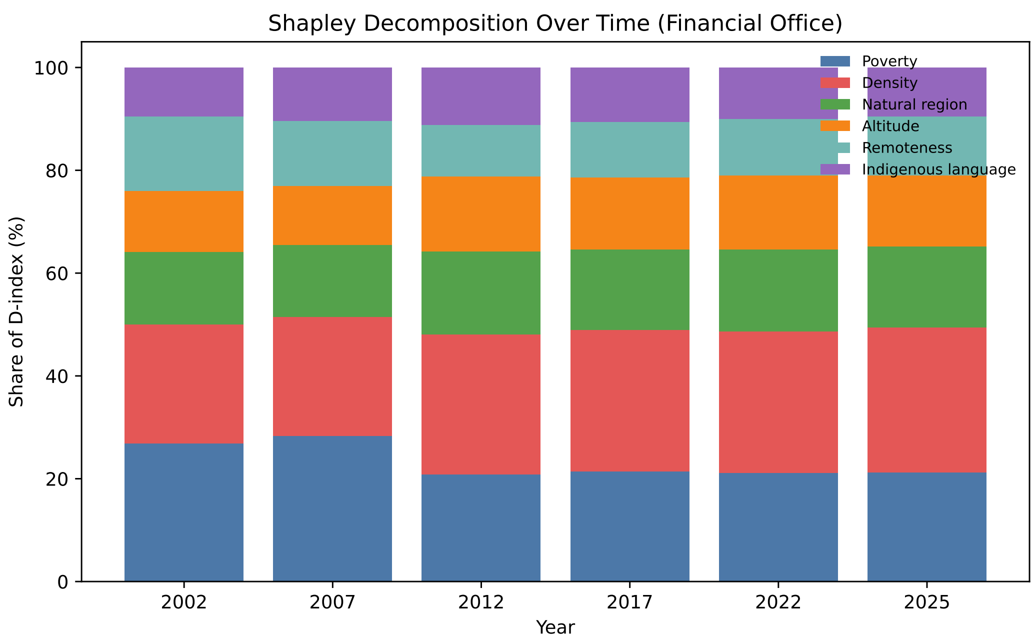

Figure 6: Shapley decomposition over time (financial offices)

The results are robust to several alternative specifications (detailed in the Appendix). Bootstrap inference with 999 district-level iid replications yields a D-index SE of 0.014 and a 95% CI of [0.275, 0.330] (Table A2). A department-level cluster bootstrap (25 clusters, 999 replications) produces wider but qualitatively unchanged results: SE = 0.023, CI [0.261, 0.349] (Table A7). When both the logit and the D-index are weighted by population, the office D-index falls to 0.135 and the any-access D-index to 0.003, with population-weighted any-access coverage of 99.6 percent (Table A3). Adding population as a seventh circumstance, it absorbs 49.1 percent of the D-index, but the other six circumstances retain 50.9 percent (Table A4). A geography-only specification using only altitude, remoteness, and natural region—variables that are fully exogenous and unaffected by reverse causality—produces a D-index of 0.234, which is 77.8 percent of the full-model value, stable across all benchmark years at 77–81 percent (Table A5). This confirms that hard geography alone accounts for roughly four-fifths of the measured inequality of opportunity. Excluding the 240 districts with imputed indigenous-language shares, the D-index is 0.268 (vs. 0.300)—qualitatively unchanged. The poverty rate (2018) and density (2017) are time-invariant across the 2002–2025 panel; the geography-only specification provides a robust alternative for early years where poverty mismatch is a concern (Table A1 reports circumstance correlations; the highest is density–population at 0.637).

6 Discussion

The Peruvian case is not a simple failure of financial inclusion. By 2025, banking agents had reached 95.1 percent of districts, and population-weighted coverage stands at 99.6 percent. The any-access D-index of 0.021 is lower than the D-index values reported for sanitation (0.10–0.20) and comparable to those for electricity (0.01–0.05) in the HOI literature for Latin America [Paes de Barros et al., 2009, Molinas Vega et al., 2010]. In terms of the *extensive margin*—whether any formal financial service point exists in a given territory—the inclusion challenge is largely resolved. The 91 remaining desert districts represent 4.9 percent of administrative geography and less than 0.5 percent of the population.

The remaining inequality lies not in whether districts have access, but in what kind of access they receive. Two-thirds of Peruvian districts (67.5 percent) are served only by banking agents, which offer basic transactional services—deposits, withdrawals, bill payments, government transfers—but not the full menu available at financial offices: credit evaluation, term deposits, safe-deposit boxes, insurance products, and personalized advisory services. The office D-index of 0.300 captures this quality gap. It means that nearly a third of the access opportunities to full-service financial infrastructure would need to be reallocated across circumstance groups to achieve territorial equality. The office D-index is above the benchmark values reported for sanitation—the most unequally distributed basic service—in the LAC HOI literature. The asymmetry is structural: water, electricity, and schooling operate under public provision mandates that push toward universality; financial offices are located by profit-seeking institutions—Caja Huancayo, BCP, Interbank—whose expansion stops where expected revenue no longer covers fixed costs.

A small high-altitude district in Puno or Huancavelica may not attract a CMAC branch—the expected revenue does not cover fixed costs. This may be an efficient market outcome, but the D-index is not a test of efficiency. It documents the distributional *consequence* of that equilibrium: residents of poor, remote, high-altitude districts face systematically lower access probabilities regardless of whether the allocation is privately optimal. The geography-only robustness check reinforces this interpretation: altitude, remoteness, and natural region—variables that are fully exogenous and unrelated to market dynamics—account for 77.8 percent of the full D-index. Even after stripping out potentially endogenous circumstances (poverty, density, indigenous language), hard geography alone generates a D-index of 0.234. The inequality of opportunity in financial-office access is, at its core, a geographic phenomenon rooted in Peru's extreme territorial heterogeneity.

Several limitations qualify the analysis. First, the logit model does not incorporate spatial dependence. A district's access probability likely depends on whether neighboring districts have offices, through substitution (residents travel to nearby served districts) or agglomeration (financial hubs attract multiple institutions). The department-level cluster bootstrap partially addresses the inference implications of spatial correlation, but it does not correct for potential bias in the logit coefficients from omitted spatial lags. Future work could include a spatial-lag variable (e.g., share of contiguous districts with access) and compare the resulting D-index.

Second, district of residence is not purely a circumstance in Roemer's sense if individuals can migrate. If people systematically move toward districts with financial offices, current residence partly reflects choice rather than predetermined circumstance. However, for the small, remote districts that drive the D-index, the dominant pattern is outmigration, not selective inmigration: people leave these districts, but few choose to move into them. The geography-only specification, which relies exclusively on altitude, remoteness, and natural region—variables entirely unaffected by individual migration decisions—captures 77.8 percent of the D-index, reinforcing the circumstance interpretation.

Third, without an observed "effort" variable, the D-index captures total circumstance-related inequality rather than the component remaining after conditioning on effort. In the territorial setting, potential effort variables might include local government initiatives to attract financial institutions, community savings mobilization, or municipal investment in connectivity infrastructure. These are not observed in the data. The IOp interpretation is therefore that the D-index measures the extent to which predicted access varies systematically with predetermined territorial characteristics—a descriptive-structural quantity that is informative for policy even without a formal effort partition.

Finally, Peru's digital finance ecosystem has grown rapidly and may be reshaping the meaning of "financial access." Yape (BCP) surpassed 15 million users by 2025, and the BCRP's interoperability regulation [Banco Central de Reserva del Perú, 2023] connected multiple platforms. Since 2022, Yape con DNI allows basic peer-to-peer transactions with only a national ID card, without requiring a bank account. These developments narrow the service gap between agent-only and office-served districts for simple transactions. However, digital finance complements rather than substitutes for physical access in important dimensions: accessing credit, making cash deposits, and receiving complex financial services still require physical infrastructure in most Peruvian districts. Moreover, digital adoption is itself unequally distributed across territories, likely correlated with the same circumstances—density, poverty, remoteness—measured here. The analysis captures a necessary condition for full financial participation that remains relevant alongside digital channels, but future work should extend the framework to incorporate digital-access indicators as they become available at the district level.

7 Conclusion

Between 2002 and 2025, Peru's financial system expanded from roughly 150 districts with a regulated financial office to over 500, while banking agents reached 95.1 percent of the territory. Applying the HOI/D-index framework—previously used for schooling, water, sanitation, and electricity—to this expansion reveals a layered process. The office D-index fell from 0.619 to 0.300 by 2012 and then froze: offices stopped entering new districts, consistent with the system reaching a commercially viable frontier. Agents continued the territorial expansion, bringing the any-access D-index to 0.021. The Shapley decomposition identifies density and poverty as the market-size variables that sort office location, while the 91 remaining desert districts are defined by the Andean and Amazonian geography that agents, too, have not yet reached.

The estimates should be read as partial, ex-ante measures of circumstance-related inequality—informative about the six observed circumstances but not a formal lower bound, since the monotonicity property has not been demonstrated for the binary D-index. The policy reading is that Peru achieved near-universal basic access through a combination of Banco de la Nación, CMACs, and private agents, but left a quality gradient between territories with full-service offices and those with only agents. Addressing this gap requires either public incentives for office expansion beyond the market frontier, or regulatory upgrades that allow agents and digital platforms to deliver credit, insurance, and advisory services to the two-thirds of districts that currently lack a physical office.

Several extensions could strengthen and generalize the analysis. First, the framework could be applied to other countries with agent-banking models—such as Colombia (corresponsales bancarios), Mexico (corresponsales), or Brazil (correspondentes bancários)—where comparable regulatory data exist and where the two-tier pattern may also be present. Second, incorporating district-level data on digital transactions (mobile-money volumes, Yape penetration) would allow a more complete measurement of effective financial access in a rapidly digitalizing environment. Third, linking territorial IOp in access to individual IOp in financial outcomes—savings rates, credit access, asset accumulation—would connect the infrastructure dimension studied here to the welfare consequences that ultimately motivate the equality-of-opportunity framework. Fourth, exploiting the temporal dimension more aggressively—for example, using the post-2012 plateau as a natural experiment to study whether agent-only access generates the same economic effects as office access—could move beyond measurement toward causal identification of the quality gap's consequences.

References

Banco Central de Reserva del Perú. Disposiciones de interoperabilidad de los servicios de pago, 2023. Circular No. 0003-2023-BCRP.

Thorsten Beck, Asli Demirguc-Kunt, and Maria Soledad Martinez Peria. Reaching out: Access to and use of banking services across countries. *Journal of Financial Economics*, 85(1):234–266, 2007. doi: 10.1016/j.jfineco.2006.07.002.

François Bourguignon, Francisco H. G. Ferreira, and Marta Menéndez. Inequality of opportunity in brazil. *Review of Income and Wealth*, 53(4):585–618, 2007.

Miriam Bruhn and Inessa Love. The real impact of improved access to finance: Evidence from mexico. *Journal of Finance*, 69(3):1347–1376, 2014.

Paolo Brunori, Francisco H. G. Ferreira, and Vito Peragine. Inequality of opportunity, income inequality, and economic mobility: Some international comparisons. *Research on Economic Inequality*, 21:85–115, 2013.

Robin Burgess and Rohini Pande. Do rural banks matter? evidence from the indian social banking experiment. *American Economic Review*, 95(3):780–795, 2005.

Daniele Checchi and Vito Peragine. Inequality of opportunity in italy. *Journal of Economic Inequality*, 8(4):429–450, 2010.

Stijn Claessens and Enrico Perotti. Finance and inequality: Channels and evidence. *Journal of Comparative Economics*, 35(4):748–773, 2007. doi: 10.1016/j.jce.2007.07.002.

Asli Demirguc-Kunt and Ross Levine. Finance and economic opportunity. Technical Report Policy Research Working Paper 4468, World Bank, 2008.

Asli Demirguc-Kunt and Ross Levine. Finance and inequality: Theory and evidence. *Annual Review of Financial Economics*, 1:287–318, 2009. doi: 10.1146/annurev.financial.050808.114334.

O. Emre Ergungor. Bank branch presence and access to credit in low- to moderate-income neighborhoods. *Journal of Money, Credit and Banking*, 42(7):1321–1349, 2010.

Francisco H. G. Ferreira and Jeremie Gignoux. The measurement of inequality of opportunity: Theory and an application to latin america. *Review of Income and Wealth*, 57(4):622–657, 2011. doi: 10.1111/j.1475-4991.2011.00467.x.

Marc Fleurbaey and Vito Peragine. Ex ante versus ex post equality of opportunity. *Economica*, 80 (317):118–130, 2013. doi: 10.1111/j.1468-0335.2012.00941.x.

Oded Galor and Joseph Zeira. Income distribution and macroeconomics. *Review of Economic Studies*, 60(1):35–52, 1993.

Gobierno del Perú. Política nacional de inclusión financiera, 2019. Decreto Supremo No. 255-2019-EF.

Scott W. Hegerty. Commercial bank locations and banking deserts. *Applied Geography*, 71:35–43, 2016.

Jith Jayaratne and Philip E. Strahan. The finance-growth nexus: Evidence from bank branch deregulation. *Quarterly Journal of Economics*, 111(3):639–670, 1996.

Arnaud Lefranc, Nicolas Pistolesi, and Alain Trannoy. Equality of opportunity and luck: Definitions and testable conditions, with an application to income in france. *Journal of Public Economics*, 93(9-10):1189–1207, 2009.

Jose R. Molinas Vega, Ricardo Paes de Barros, Jaime Saavedra Chanduvi, Marcelo Giugale, Louise J. Cord, Carola Pessino, and Amer Hasan. *Do Our Children Have a Chance? The 2010 Human Opportunity Report for Latin America and the Caribbean.* World Bank, 2010.

Hoai-Luu Q. Nguyen. Are credit markets still local? evidence from bank branch closings. *American Economic Journal: Applied Economics*, 11(1):1–32, 2019.

Ricardo Paes de Barros, Francisco H. G. Ferreira, Jose R. Molinas Vega, and Jaime Saavedra Chanduvi. *Measuring Inequality of Opportunities in Latin America and the Caribbean.* World Bank, 2009.

Xavier Ramos and Dirk Van de gaer. Approaches to inequality of opportunity: Principles, measures and evidence. *Journal of Economic Surveys*, 30(5):855–883, 2016.

John E. Roemer. *Equality of Opportunity.* Harvard University Press, 1998.

John E. Roemer and Alain Trannoy. Equality of opportunity: Theory and measurement. *Journal of Economic Literature*, 54(4):1288–1332, 2016. doi: 10.1257/jel.20151206.

Superintendencia de Banca, Seguros y AFP del Peru. Reglamento de canales complementarios de atención al público de las empresas del sistema financiero, 2013. Resolución SBS No. 6285-2013.

Superintendencia de Banca, Seguros y AFP del Peru. Cifras de inclusion financiera, 2026. URL https://www.sbs.gob.pe/Inclusion-financiera-principal/Cifras-de-inclusion-financiera.

Florian Wendelspiess Chavez Juarez and Isidro Soloaga. Iop: Estimating ex-ante inequality of opportunity. *The Stata Journal*, 14(4):830–846, 2014. doi: 10.1177/1536867X1401400408.

Appendix

Table A1: Circumstance correlations, 2025

	Pov.	Alt.	Den.	Rem.	Indig.	Pop.
Poverty	1.000					
Altitude (log)	0.424	1.000				
Density (log)	-0.251	-0.242	1.000			
Remoteness (log)	0.297	0.155	-0.603	1.000		
Indigenous	0.316	0.454	-0.154	0.135	1.000	
Population (log)	-0.193	-0.427	0.637	-0.440	-0.149	1.000

Notes: Log transformations applied to altitude, density, remoteness, and population. Lower triangle only.

Table A2: Bootstrap inference for D-index and Shapley shares, 2025

Circumstance	Share (%)	SE	95% CI
Poverty	21.3	2.3	[16.6, 25.9]
Altitude	13.8	1.4	[11.2, 16.6]
Density	28.1	2.6	[23.1, 33.7]
Remoteness	11.5	1.4	[9.0, 14.3]
Natural region	15.8	2.0	[12.1, 19.9]
Indigenous language	9.5	1.3	[7.1, 12.2]
D-index	0.300	0.014	[0.275, 0.330]

Notes: 999 iid replications, district-level resampling. Percentile-based CIs. Zero failures. Seed = 42.

Table A3: Population-weighted D-index

Year	Cov.	D	HOI	Cov.	D	HOI
2002	0.532	0.299	0.373	0.532	0.299	0.373
2007	0.615	0.238	0.469	0.615	0.238	0.469
2012	0.763	0.141	0.655	0.833	0.118	0.734
2017	0.774	0.139	0.666	0.965	0.020	0.946
2022	0.786	0.134	0.680	0.987	0.007	0.980
2025	0.789	0.135	0.682	0.996	0.003	0.993

Notes: Both the logit model and the D-index aggregation use 2017 census population weights.

Table A4: Shapley decomposition with and without population, 2025

Circumstance	Baseline (%)	With pop. (%)
Poverty	21.3	11.7
Altitude	13.8	7.2
Density	28.1	12.1
Remoteness	11.5	6.2
Natural region	15.8	8.3
Indigenous language	9.5	5.4
Population	—	49.1

Notes: Baseline uses six circumstances. "With population" adds log population as a seventh. Shares in %.

Table A5: D-index: full model vs. geography-only, 2025

Specification	D-index	Share of full (%)
Full model (6 circumstances)	0.300	100.0
Geography only (3 circumstances)	0.234	77.8

Notes: Geography-only uses altitude, remoteness, and natural region. Full model adds poverty, density, and indigenous language.

Table A6: Average marginal effects from year-specific logit models

Variable	2002	2012	2025
Poverty	-0.0536	-0.0596	-0.0668
Altitude	-0.0034	-0.0293	-0.0274
Density	0.0408	0.1173	0.1272
Remoteness	-0.0064	0.0151	0.0037
Coast	0.0220	0.0323	0.0442
Amazon	0.0434	0.0599	0.0808
Indigenous	-0.0043	-0.0172	-0.0066
Observations	1,873	1,873	1,873
Coverage	0.080	0.251	0.276
Pseudo-R^2	0.406	0.181	0.190
AUC	0.922	0.782	0.782

Notes: AMEs computed as $\overline{\beta_k \hat{p}(1-\hat{p})}$, averaged over all districts. AUC is area under the ROC curve. Pseudo-R^2 declines mechanically as coverage rises from 8% to 28%.

Table A7: D-index inference: iid vs. cluster bootstrap, 2025

Method	SE	95% CI
District iid (999 reps)	0.014	[0.275, 0.330]
Department cluster (999 reps)	0.023	[0.261, 0.349]

Notes: Point estimate D = 0.300. Seeds: 42 (iid), 43 (cluster). Zero failures in both.

Table A8: Circumstance variables: sources and construction

Variable	Source	Year	Notes
Poverty rate	INEI district poverty map	2018	Complete for estimation sample
Altitude	Census coordinates	2017	Log-transformed, standardized
Density	Census pop. / area	2017	Log-transformed, standardized
Remoteness	Haversine to dept. hub	2017	Log-transformed, standardized
Natural region	SBS panel classification	—	Coast, Sierra (ref.), Amazon
Indigenous lang.	ENAHO pooled 2012–2023	Pooled	1,633 direct; 240 province-imputed
Population	Census 2017	2017	Used in robustness only

Notes: Indigenous-language share is the proportion of individuals whose mother tongue is Quechua, Aymara, or another native language. All continuous variables standardized to mean zero, unit variance across the pooled panel.