

Property with Different Rights: The Long-Run Effects of Ecuador's Agrarian Reform

Nicholas Gachet*

2026/10/16

Abstract

In its coastal region, Ecuador's agrarian reform transferred land with different rights. Public land transfers (PLT) gave beneficiaries broad control over crop choice. Outside PLT, the state relied more on expropriation, and Decree 1001 required recipients of expropriated land in rice-growing zones to continue producing rice. A geographic regression discontinuity at the PLT boundary estimates the relative long-run effects of these reform environments. Three decades later, PLT-side areas cultivated less rice and shifted land and revenue toward sugar cane, oil palm, and plantain. These crops widened production beyond rice and bananas, the region's established high-productivity crops. Adoption, rather than higher output per hectare among incumbent producers, drove this shift. Banana yields remained higher outside PLT. Historical crop knowledge, producer selection, and greater banana-land concentration offer plausible explanations, although the design cannot distinguish among them. Aggregate productivity can therefore obscure diversification gains when incumbent-crop performance moves in the opposite direction. I explore several development-related implications of these results.

JEL classification: N56, O17, O43, P14, Q15.

*Department of Economics, Tecnológico de Monterrey (ngachet@tec.mx). I received excellent comments and suggestions from Victoria Alomía, Leticia Arroyo-Abad, James Banks, Daniel Baquero, Guillaume Blanc, Felipe Bragues, Jonathan Chapman, David Delacrétaz, James Fenske, Santiago José Gangotena, Rachel Griffith, Sonya Krutikova, Rocco Macchiavello, Luciano Martínez, Peter Murrell, Nuno Palma, Samuel Pienknagura, James A. Robinson, Miriam Venturini, Chris Wallace, and Mazhar Wasseem. Nicolás Chuquimarca and Jorge Paredes provided outstanding research assistance.

1. Introduction

Property rights shape both whether farmers invest and what they produce. Productivity in one established crop does not guarantee broad development. A rights regime can preserve high output in incumbent crops while limiting movement into new activities. In agrarian settings, the key long-run question is whether institutions let productive activity spread beyond a narrow crop structure. This distinction helps explain why aggregate measures can make reform look disappointing even when crop choice and labor allocation change. Latin American debates on agrarian reform have often emphasized disappointing aggregate results, distortions, and weak incentives under expropriation-based settings (de Janvry, 1981; Kay, 2002, 2019). Much less is known about public land settlement, internal colonization, or public land transfers¹, even though these policies were also common in several contemporary agrarian reforms in Latin America (Albertus, 2015). Few studies examine that margin directly. Almost no quantitative work compares expropriation and public land transfers within the same reform episode, even though both strategies often coexisted.

This paper studies the long-run effects of transferring property with different rights. I compare adjacent coastal areas assigned to public land transfers (PLT) with areas outside the realized PLT footprint. PLT gave beneficiaries more freedom over crop choice. Outside PLT, the state relied more on expropriation, and Decree 1001 required recipients of expropriated land in rice-growing zones to continue producing rice. A geographic regression discontinuity compares assignment to these neighboring reform environments (Keele and Titiunik, 2015; Cattaneo et al., 2019). It identifies their relative long-run effects, not the absolute effect of either policy or any single right. Three decades later, PLT-side areas cultivated less rice and shifted land and revenue toward sugar cane, oil palm, and plantain. These crops widened production beyond rice and bananas, the region's established high-productivity crops. Adoption, rather than higher output per hectare among incumbent producers, drove this shift. Banana output per hectare remained higher outside PLT. Historical crop knowledge, producer selection, local conditions, and greater banana-land concentration offer plausible explanations, although the design cannot distinguish among them. Aggregate productivity can therefore obscure diversification gains when incumbent-crop performance moves in the opposite direction. I also find suggestive evidence of movement out of agricultural employment on the PLT side. Although PLT often granted larger plots elsewhere, agricultural surface is locally similar across the border. The design therefore

¹ Throughout the paper, I use the term public land transfers to refer to this land-allocation policy.

compares different reform regimes rather than large and small farms.

Ecuador's agrarian reform placed both allocation strategies in a common agrarian landscape. The same history shaped both sides of the border. Large haciendas first organized the cocoa economy. External shocks then brought its collapse. Rice spread across abandoned or weakly controlled estates, and the banana boom created new opportunities for accumulation (Chiriboga, 1980; Arosemena, 1991; Guerrero, 1994; Uggen, 1975; Espinosa, 2014; Larrea, 1987). The Ecuadorian Institute of Agrarian Reform and Colonization (IERAC, its Spanish acronym) later assigned neighboring areas to two institutional environments. PLT gave beneficiaries broad freedom over crop choice. Outside PLT, the state used expropriation more intensively, and Decree 1001 tied transfers of expropriated land in rice-growing zones to continued rice production.

The empirical strategy compares locations around the digitized outer edge of the PLT footprint in IERAC's 1978 map. The baseline uses a sharp geographic regression discontinuity with robust bias correction, a 40 km bandwidth, and geographic controls. It excludes the two parishes closest to the threshold to reduce uncertainty about the historical border. A digitized Decree 1001 registry and mapped policy footprints clarify the treatment contrast. In the preferred corridor, Decree 1001 awards appear in nine of twelve parishes outside the PLT footprint and none of the fourteen parishes inside it. The 2000 agricultural census provides the main outcomes, while earlier censuses establish the pre-reform context. Evidence on credit, crop structure, and structural transformation helps interpret the results. A broad set of robustness checks tests the design. For interpretation only, I also report a brief back-of-the-envelope adaptation of Kim and Wang (2025), using the reallocation logic in McCaig and Pavcnik (2018). It shows how aggregate productivity measures can obscure diversification and labor-reallocation gains within the same reform.

The crop-specific results reveal opposing margins that aggregate productivity can hide. Outside PLT, Decree 1001 required recipients of expropriated land in rice-growing zones to continue producing rice. The informal-credit pattern is consistent with persistence along that direct route. Banana output per hectare also remained higher outside the PLT footprint. Accumulated expertise, producer selection, local conditions, and a more concentrated banana sector offer plausible explanations for this incumbent-crop advantage. PLT instead gave producers more freedom to adopt sugar cane, oil palm, and plantain. Adoption, rather than uniformly higher output per hectare among incumbent producers, drove this broader portfolio. Property rights can therefore change the direction of adaptation even when aggregate agricultural outcomes remain weak.

The paper contributes to three literatures. First, it shifts the Latin American agrarian-reform


literature from average reform effects toward the allocation strategies within reforms. Most empirical studies examine expropriation and its effects on productivity, inequality, and long-run development (Besley and Burgess, 2000; Do and Iyer, 2008; Galán, 2020; Montero, 2022, 2023; Edwards et al., 2022). Quantitative work on the Philippines shows that government land assignment and restrictions on transfer can reduce agricultural productivity by misallocating land and distorting occupational and technology choices (Adamopoulos and Restuccia, 2020). I study a complementary margin: how the rights attached to redistributed land shape crop choice within one reform. A smaller literature studies public land distribution and PLT, often through their political economy or distributive effects (Albertus, 2015; Faguet et al., 2020). I connect these strands by comparing assignment to neighboring reform environments within one episode. Work on Taiwan likewise shows that distinct phases of one agrarian reform can produce different effects on productivity and structural change (Kim and Wang, 2025). Agrarian reform is not a single policy when its components assign different rights and constraints. Second, the paper shows that agricultural development depends on the bundle of property rights attached to land allocation, not only on whether rights exist (Alchian, 1965; Hornbeck, 2010; Galiani and Schargrodsky, 2010; Bühler, 2021; Smith, 2020). The design compares bundles and does not separately identify each right. Third, it links local reform design to persistent differences in crop structure, misallocation, and structural transformation (Adamopoulos, 2008; Herrendorf et al., 2014; Kitamura, 2022; Blankespoor et al., 2022).

Section 2 presents the historical and institutional context. Sections 3 and 4 describe the data and empirical strategy. Section 5 reports the results and robustness exercises, Section 6 discusses mechanisms, and Section 7 concludes.

2. Historical and Institutional Context

Latin American agrarian reforms transferred property with different rights. Some governments relied mainly on expropriation. Others, including Ecuador, Bolivia, and Colombia, combined expropriation with large public land transfers (PLT) (Albertus, 2015). Figure 1 shows the importance of both channels across the region. The channels assigned different rights even within the same national reform. Ecuador therefore permits a comparison of two rights bundles within one reform episode and regional context.

Figure 1: Public land transfers and expropriation in Latin American agrarian reforms



Note: Public land transfers were quantitatively important in several Latin American agrarian reforms, including Bolivia, Colombia, and Ecuador. Source: [Albertus \(2015\)](#).

2.1. From cocoa to rice and bananas

The cocoa economy created the agrarian landscape that IERAC later divided. During the nineteenth and early twentieth centuries, Ecuador became a major cocoa exporter, and the coastal provinces of Guayas and Los Ríos organized production around large haciendas. A small landed elite controlled much of the land, often from abroad, while workers cultivated it under precarious labor and tenancy arrangements ([Chiriboga, 1980](#); [Arosemena, 1991](#); [Guerrero, 1994](#)).

The cocoa system collapsed in the 1920s. Falling world prices and witches' broom disease made many estates unprofitable. Absentee owners abandoned them or lost effective control. Workers remained on the land, divided fields, and entered informal arrangements that the state did not fully recognize. The collapse changed crops and property relations at the same time. Productive communities remained, but estate boundaries and legitimate claims became difficult to establish ([Uggen, 1975](#); [Espinosa, 2014](#); [Carrasco, 1994](#)).

Rice spread across parts of this post-cocoa landscape. Unlike cocoa, rice served the domestic market and became increasingly important to food security. Rising imports and balance-of-payments pressure made the crop politically salient, and the government later treated local rice production as a national priority. Small rice plots could coexist with unsettled ownership because former workers continued to cultivate land whose legal status remained disputed ([Espinosa, 2014](#); [Redclift, 1976](#)).

The banana boom of the 1950s created a second export economy before agrarian reform. Pro-

ducers in central and southern Los Ríos and Guayas converted some former cocoa land to bananas, and both large and small landholders entered the sector. Bananas brought commercial opportunities and valuable crop-specific experience, but they did not resolve the region's unequal or uncertain land relations (Larrea, 1987). By the 1960s, the coastal economy therefore combined two established crops: rice as the principal annual staple and bananas as the principal perennial export. Agrarian reform entered this settled, commercially active landscape rather than a blank frontier.

2.2. Agrarian Reform, PLT, and Decree 1001

The 1964 Agrarian Reform and Colonization Law created the Ecuadorian Institute of Agrarian Reform and Colonization (IERAC) and charged it with redistributing land and expanding cultivation. IERAC used two main channels. Through PLT, called colonization in the legislation, the state adjudicated public, abandoned, or underused land to farmers and sought to expand the agricultural frontier (Canelos, 1980; Barsky, 1984). Through expropriation, the state intervened in large estates and transferred land, often to people who already worked it. After an initial slowdown, the government renewed coastal reform in 1970, but IERAC continued to administer the two channels under different legal rules (Barsky, 1984; Redclift, 1976).

The two channels shared some financial features. Beneficiaries generally repaid the state over 20 to 25 years, after a grace period of three to five years, on terms more favorable than private credit (Barsky, 1984). Yet the contracts did not grant the same control over land use. PLT followed a relatively uniform policy and gave beneficiaries more flexible allocation terms, often larger adjudicated plots, and greater freedom over production (Barsky, 1984; Carrasco, 1994). Pilot settlement projects began in Santo Domingo in the 1950s, but the 1964 law turned PLT into a central allocation program.

PLT did more than settle untouched land. Farther north, it extended the agricultural frontier. In Guayas and Los Ríos, however, IERAC also applied PLT to abandoned cocoa haciendas whose former workers already occupied and cultivated the land. The state used the same legal channel to formalize claims in an old agricultural region. Near the policy border, IERAC therefore did not reserve PLT for empty land or places without a crop history (Carrasco, 1994).

Expropriation took a more restrictive form in southern Guayas and Los Ríos. President José María Velasco Ibarra issued Decree 1001 amid concern about precarious labor arrangements, rising rice imports, and falling domestic production. The decree accelerated intervention in rice-

growing estates and tied access to expropriated land to continued rice cultivation. Its terms effectively committed beneficiaries to rice for ten years. Beneficiaries received land and favorable repayment terms, but the state directed how they should use it. Decree 1001 therefore supplied a direct route from expropriation to crop specialization: it made rice production part of the bargain rather than leaving crop choice entirely to the farmer (Barsky, 1984; Redclift, 1976).

Bananas reveal an apparent puzzle. Decree 1001 targeted rice, yet the later data show higher banana output per hectare outside the PLT footprint. The decree did not require farmers to grow bananas. Established expertise, producer selection, land quality, or reluctance to leave a familiar crop could all generate this pattern. A later agricultural cadastre also shows greater concentration of banana land outside PLT. These facts make persistence in an incumbent crop plausible, but the design cannot distinguish among its causes. The banana result therefore documents a contrasting incumbent-crop margin rather than a second statutory route.

The policies differed in control over land use, not simply ownership. In the sense of Alchian (1965), property consists of a bundle of rights. PLT beneficiaries gained broader authority to redirect land toward another crop. Decree 1001 beneficiaries received land under a rule that reserved more control for the state and explicitly required rice production. Cooperative organization and the selection of land and beneficiaries also differed across the reform environments. The design estimates the combined relative effect of these rights bundles.

2.3. IERAC Policy Border

The 1978 IERAC map records the reform's realized spatial footprint. Its green polygon marks the area reached by PLT. I digitize the polygon's outer edge and use it as the policy border. The cutoff therefore comes from the realized expansion of PLT, not from a line jointly drawn to separate PLT from expropriation. Beyond the polygon, especially to the south, IERAC relied more heavily on expropriation after Decree 1001 intensified intervention in the rice economy. Figure 2 shows the reconstructed PLT footprint, the nearby expropriation areas, and the boundary used in the analysis.


The post-cocoa land problem helps explain why PLT could expand through otherwise similar places and stop at this edge. Abandoned haciendas stretched across both sides. Former workers occupied and cultivated land even when estate limits and ownership claims remained unclear. IERAC had discretion to classify these properties as abandoned land eligible for PLT or as estates subject to intervention and expropriation. Near the polygon's edge, the agency applied different

legal solutions to communities that shared the same hacienda history (Carrasco, 1994; Guerrero, 1994).

Quevedo, Mocache, Zapotal, and Balzar make the contrast especially clear. These places lay near the edge of the realized PLT footprint and contained land from the former cocoa economy. Farther north, PLT increasingly resembled frontier settlement; close to the border, it instead reached established farms and resident workers. The local comparison therefore does not pit frontier settlers against an old plantation economy. It compares neighboring parts of the same post-cocoa agricultural region.

Historical evidence also ties the northern side to the same cocoa collapse. Contemporary accounts describe the abandonment of Pichilingue, near Quevedo, after crop disease damaged the cocoa economy; agricultural programs later used the area as an important center of experimentation (Leonard, 1947; Cuvi, 2015). The international price decline and plant disease preceded the reform and affected haciendas on both sides of the later IERAC line (Chiriboga, 1980; Guerrero, 1994). This common shock left IERAC to resolve similar land claims under different statutes.

Figure 2: GIS reconstruction of the IERAC policy border and reform regimes



Note: Green polygons show the realized PLT footprint recorded in IERAC's 1978 map, and red polygons mark expropriations associated with Decree 1001. The black line traces the digitized outer edge of the PLT polygon used as the RD cutoff; gray lines show administrative boundaries.

The polygon edge separates assignment to the PLT footprint more clearly than initial agricultural potential. It follows neither the limit of cultivated land nor a sharp break in the pre-reform

crop economy. Outside the footprint, IERAC relied more on expropriation, and Decree 1001 governed many transfers of expropriated land in rice-growing areas. Neither policy covered every hectare uniformly. The design therefore estimates the relative effect of the PLT-side reform environment within a common historical landscape. The balance and continuity tests below assess the required local-continuity assumption (Keele and Titiunik, 2015).

The reform environments created different incentives for crop choice. Decree 1001 required recipients of expropriated land in rice-growing areas to continue producing rice. PLT gave producers more room to change course. Banana output per hectare remained higher outside the PLT footprint. Historical expertise, producer selection, local conditions, and later banana-land concentration offer plausible explanations, although the design cannot separate their contributions. The core institutional contrast concerns mandated rice production and the freedom to adopt sugar cane, oil palm, and plantain. Bananas show that an established crop could remain productive without producing a broader portfolio.

3. Data

3.1. Agricultural census and historical sources

Ecuador's National Institute of Statistics and Censuses (INEC) collected the 2000 National Agricultural Census. It supplies the core plot-level data (Instituto Nacional de Estadística y Censos, 2002b,a). Its crop questionnaires report production, area, crop age, and land use. The census also records agricultural surface, tenure, workers, and farm credit. I use crop records for the main outcomes and tenure as a baseline control. Credit provides suggestive evidence on rice persistence. Workers provide a post-treatment diagnostic of farm labor scale. Surface enters the landholding checks. The main sample covers the coastal reform corridor. The appendix extends the comparison to Manabi. The census places annual and perennial crops in separate questionnaires. This feature determines the outcome samples described below.

Two earlier national agricultural censuses establish the crop history around the reform. I digitized canton-level records from the 1954 census and parish-level records from the 1974 census. The 1954 census describes agriculture before the 1964 reform and reports production and cultivated area for rice and bananas. The 1974 census records land in annual and perennial crops during the reform's first decade. Together with the 2000 census, these sources trace crop composition from the pre-reform economy to the long run. Their geographic units differ from the

plot-level data, so I use them as descriptive evidence rather than as part of the local RD.

IERAC's archival records identify where the state applied each reform channel. Statistical reports record the amount of land assigned through PLT and expropriation by local administrative unit. A digitized Decree 1001 registry records adjudications, hectares, and beneficiary households. IERAC's 1978 map records the realized PLT footprint. I digitized and georeferenced its outer edge, overlaid it with current parish boundaries, and calculated each parish's signed distance to that edge. I also overlay the mapped PLT and expropriation footprints with parish boundaries. The outer PLT edge defines treatment status and the running variable in the geographic RD. The registry and footprint overlays instead describe the composition of the two sides. Appendix Figure A1 reproduces the original 1978 map, Figure 2 shows the georeferenced reconstruction, and Appendix Table A6 reports the treatment-composition evidence.

3.2. Outcome construction

I construct perennial revenue per hectare from banana, cacao, sugar cane, oil palm, and plantain. I construct annual revenue per hectare from rice, maize, beans, and potatoes. For these aggregates, I value each crop with the fixed 2013 INEC producer-price schedule implemented in the analysis code. I divide revenue by the corresponding area and use $\log(1 + \text{revenue per hectare})$.² These coefficients describe changes in a zero-inclusive transformed outcome. They are not percentage changes in untransformed revenue, profits, or total factor productivity. I also study physical output per hectare for rice, cacao, and banana. A narrower diversification basket contains sugar cane, oil palm, and plantain. Tables refer to these three crops as the main other-perennial basket.³ This basket measures movement beyond the established cacao-banana structure. The crop groups follow INEC's distinction between permanent (perennial) and transitory (annual) crops and the principal coastal crop patterns in the 2000 census ([Instituto Nacional de Estadística y Censos, 2002b,a](#)).

The 2000 census records perennial and annual crops in separate modules. I therefore use the relevant questionnaire for each outcome rather than force every regression onto the overlap sample. Perennial outcomes use the perennial module, and annual outcomes use the annual module. I define perennial-questionnaire membership before applying the crop-age restriction. I

² The implemented bean price averages the monthly 2013 tender red-bean series. A broader bean average recorded in the source workbook changes the annual-crop coefficient from -0.524 to -0.511 . It changes the aggregate residual-productivity coefficient from -0.573 to -0.570 . The conclusions remain unchanged.

³ Appendix robustness exercises expand this narrower perennial basket to include additional crops such as coffee, passion fruit, lemon, palmito, and pineapple. The positive diversification result remains.

retain records with missing crop age and exclude only reported ages of at least 37 years, measured from the reform's start in 1964. I assign zero only when a plot appears in the relevant module but reports no eligible production for the outcome. The overlap sample provides a robustness check.

I construct the perennial revenue denominator conservatively: harvested area first, productive area when harvested area is missing, and planted area as the final fallback. Observations without usable area after these steps remain missing. For cacao and banana yields, I divide total plot production by total usable crop area rather than average crop-record yields. A fixed cross-crop schedule of Ecuadorian 2013 INEC producer prices keeps local price variation from driving the comparisons.

3.3. Additional data

Spatial data describe the geography around the IERAC border. WorldClim supplies precipitation and temperature. The Center for Climatic Research at the University of Wisconsin–Madison supplies elevation and general agricultural suitability, and [Nunn and Puga \(2012\)](#) supplies terrain ruggedness. The Food and Agriculture Organization's Global Agro-Ecological Zones database (FAO-GAEZ) supplies agro-climatic potential yields for banana, rice, cacao, and oil palm. I merge these gridded measures with the agricultural and parish data. The preferred RD includes banana potential yield because bananas predated the reform and could otherwise confound the comparison. The balance tests and auxiliary specifications use the broader geographic set.

Harmonized IPUMS International population censuses provide the labor-allocation outcomes ([Center, 2020](#)). I merge person and household variables one-to-one by census sample, household serial number, and person number. I use IPUMS person weights. The analysis covers the 1962, 1982, 1990, 2001, and 2010 censuses. It measures agricultural employment, manufacturing employment, and urban residence among the active labor force. I exclude the harmonized 1974 IPUMS file because of data-quality problems. This exclusion does not affect the separately digitized 1974 agricultural census described above. The urban indicator is unavailable in 1982. IPUMS identifies cantons rather than parishes. I therefore aggregate reform exposure to cantons and treat these results as supporting evidence. The productivity-reallocation calibration uses a separate national source. Ecuador's 2000 observations in the GGDC/UNU-WIDER Economic Transformation Database provide the agricultural employment share and the productivity gap between agriculture and non-agriculture ([Kruse et al., 2023](#)).

The Ministerio de Agricultura y Ganaderia (MAGAP) supplies two georeferenced sources on later agricultural structure. Its 2021 land-use data classify grid cells by crop duration and by whether observed land use matches local productive potential. I use these data to compare land-use alignment for perennial and other agricultural cells. A later MAGAP banana cadastre records banana producers and cultivated area. I use it descriptively to compare producer size and concentration within 40 and 60 km of the border; its uneven coverage does not support an RD estimate.

Finally, Defense Meteorological Satellite Program night-light rasters distributed by the National Oceanic and Atmospheric Administration’s National Centers for Environmental Information provide a proxy for local economic activity. I construct parish-level light measures for 1992, 2001, and 2010 from the raster data and compare them around the reform border with heteroskedasticity-robust inference. The IPUMS, MAGAP, night-light, and ETD exercises help interpret the crop results, but none forms part of the paper’s core identification strategy.

4. Empirical Strategy

4.1. Baseline geographic RD

The main design uses a sharp geographic regression discontinuity around the digitized outer edge of the PLT polygon in IERAC’s 1978 map. Let d_p denote parish p ’s signed distance from that edge, with positive values inside the PLT polygon.⁴ I estimate:

$$Y_{ipb} = \alpha + \tau PLT_p + \beta d_p + X_p' \Gamma + f(\text{lat}, \text{lon})_p + \phi_b + \delta_t + \gamma_{\text{prov}} + \varepsilon_{ipb} \quad (1)$$

Here, Y_{ipb} is an agricultural outcome for plot i in border segment b and parish p . The indicator PLT_p equals one inside the realized PLT footprint. The variable d_p measures signed distance to its outer edge. The vector X_p contains parish geography and crop-suitability controls. The function $f(\text{lat}, \text{lon})_p$ captures local geography. The terms ϕ_b , δ_t , and γ_{prov} absorb average differences across border segments, tenure categories, and provinces. The coefficient τ measures the relative jump on the PLT side. It combines the features of the neighboring reform environments and does not isolate one contractual right. I estimate it with robust bias-corrected local-linear RD methods following Calonico et al. (2014, 2017, 2021) and the practical guidance in Cattaneo et al. (2019).

The baseline uses a 40 km bandwidth and clusters standard errors by parish. It excludes the

⁴ This specification follows the canonical model in Dell (2010).

two parishes closest to the threshold. I call this exclusion the doughnut. The historical map does not locate its edge with modern surveying precision. The doughnut therefore reduces boundary ambiguity and treatment misclassification (Barreca et al., 2011). Controls include latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield. Full tenure occurs more often on the PLT side in the broader sample but shows no equally sharp local break at the cutoff. The controls absorb remaining observed geographic differences. They do not create identification.⁵

4.2. Identification

The source of the cutoff anchors identification. The 1978 IERAC map records the realized PLT footprint. I use the outer edge of its green polygon as the RD threshold. The cutoff is a PLT boundary, not a line jointly drawn to separate PLT from expropriation. PLT expanded across an already settled post-cocoa region and stopped at this edge. IERAC relied more heavily on expropriation outside the polygon, especially to the south. The sharp RD estimates the relative effect of assignment to the historical PLT side. It does not estimate an absolute PLT effect, an absolute expropriation effect, or the effect of an additional hectare treated through either program.

IERAC's administrative reports provide a supporting check on the mapped treatment geography. I match the reported PLT hectares to current parish boundaries. When parish borders changed, I use the historical parent jurisdiction. The reports, the 1978 map, and the agricultural census use different geographic units. The reconstruction therefore cannot produce a perfect parish-by-parish match.⁶ The administrative share remains noisy, but the fitted probability of lying inside the PLT polygon rises at its edge. Appendix Figure A5 reports both measures.

The Decree 1001 registry and polygon overlays show what differs across the corridor. In the full 40 km sample, ten of thirteen parishes outside the PLT footprint contain a recorded Decree 1001 award, compared with one of fifteen parishes inside it. The sole inside parish with an award is Palenque, the parish closest to the cutoff. The preferred doughnut removes Palenque and Zapotal. Nine of twelve remaining outside parishes then contain an award, while none of the fourteen PLT-side parishes does. The mapped footprints show the same asymmetry. In the doughnut sample, mapped expropriation covers an average of 16.1 percent of parish area outside the foot-

⁵ Where I report standard errors robust to unequal residual variance, I use the leverage-adjusted HC3 correction.

⁶ The crosswalk treats San Carlos in Los Rios as a PLT parish. The digitized 1978 polygon covers 94.9 percent of its current area. Manga del Cura was a nondelimited census territory that joined Manabi in 2015. Its 87 historical census sectors map to El Carmen, Santa Maria, and El Paraiso La 14. Their combined current territory is 98.8 percent covered by the mapped PLT footprint and contains no recorded Decree 1001 award. I therefore retain Manga del Cura on the PLT side.

print and 0.5 percent inside it. Mapped PLT covers 90.0 percent inside and less than 0.1 percent outside. Appendix Table A6 reports these comparisons. The award registry and mapped polygons measure different objects. They clarify treatment composition but do not create a second assignment rule or prove local exogeneity.


Appendix Tables A7 and A8 use the administrative PLT share as treatment intensity. This is the fuzzy RD exercise. The share varies only at the parish level. The full sample contains 28 parishes, and the doughnut sample contains 26. The boundary separates two reform environments rather than treatment from no treatment, and the administrative reconstruction adds measurement error. I therefore treat the fuzzy estimates as robustness evidence. The sharp RD remains the preferred assignment comparison.

The map records realized treatment geography. It does not prove random placement. Identification requires potential outcomes to vary smoothly at the footprint's edge without the change in reform environment. Near the cutoff, the polygon divided neighboring parts of the same post-cocoa economy. It did not divide cultivated land from an empty frontier. Nor did it follow one natural boundary or a sharp ecological break. The 1954 crop comparisons and the balance tests support this continuity assumption. A separate test finds no detectable break in agricultural surface in 2000. The doughnut removes the two parishes where the historical border and policy composition are most ambiguous. It loses the closest observations but reduces treatment misclassification.

4.2.1. Balance and pre-reform evidence

Predetermined geography shows no statistically detectable discontinuity at the cutoff. Figure 3 reports sharp parish-level local-linear estimates for climate, terrain, and crop suitability using the preferred 40 km doughnut sample. The test has limited power with 26 parishes, so the absence of significant jumps supports local continuity without proving exact equivalence. Full tenure occurs more often on the PLT side in the broader sample but shows no comparable local break at the cutoff. The figure also reports mean agricultural surface from the 2000 census as a separate comparability check. Surface is post-treatment rather than predetermined, but the estimate reveals no statistically detectable difference in the scale of agricultural units at the border. The design therefore compares neighboring places with a common agrarian background and no evident local break in observed farm scale, but different legal regimes.


Figure 3: Continuity in geography and agricultural surface



Note: The figure reports robust bias-corrected local-linear sharp RD estimates in standard-deviation units. The sample contains 26 parishes within 40 km of the IERAC border after excluding the two parishes closest to the cutoff. Bars show 95 percent confidence intervals based on standard errors robust to unequal residual variance. Climate, terrain, and crop-suitability measures describe predetermined geography. Agricultural surface is mean reported agricultural-unit area in the 2000 agricultural census; it is post-treatment and serves only as a comparability diagnostic.


Historical censuses also place both sides on a common pre-reform trajectory. Figure 4 compares Quevedo with the 1954 control mean for rice and bananas, the coast's central crops. The areas need not match on every margin; the relevant fact is that the later boundary did not divide clearly different crop economies. Both belonged to the same post-cocoa landscape. Figure 5 tracks Quevedo-minus-control gaps in annual and perennial land shares across available censuses. Quevedo shows no clear perennial advantage in 1954; the gap widens in 1974 and narrows again by 2000. Because different censuses and mixed control groups produce these canton-level aggregates, I treat them as descriptive context rather than a preview of the local RD estimates.

Figure 4: Pre-reform rice and banana yields in 1954



Note: Descriptive evidence from the 1954 agricultural census. The figure compares crop yields in Quevedo with the corresponding control mean for rice and bananas, the two historically central coastal crops highlighted in the paper's background discussion.

Figure 5: Historical annual and perennial land-share gaps over time



Note: Descriptive comparison across the historical agricultural censuses available for this project. The figure tracks Quevedo-minus-control gaps in the annual and perennial land shares over time for alternative control groups. These are canton-level historical aggregates rather than local RD objects.

5. Results

5.1. Main results

PLT changed the composition of agriculture more clearly than aggregate perennial revenue per hectare. The aggregate perennial coefficient is -0.172 , only 4 percent of the estimation-sample mean of 3.910 on the transformed scale, and is statistically imprecise. The crop-specific estimates reveal much larger shifts in opposite directions.

Rice output falls by 0.661 tonnes per hectare on the PLT side. This decline equals 30 percent of the estimation-sample mean of 2.238 tonnes. Revenue from sugar cane, oil palm, and plantain moves in the opposite direction. I group these crops in the tables as the main other-perennial basket. Its coefficient is 1.239 in $\log(1 + \text{fixed-price revenue per hectare})$, or 54 percent of the outcome's estimation-sample mean of 2.316. The discontinuity is therefore large relative to the typical value of the index. Because the index includes zero revenue and then applies the logarithmic transformation, the 54 percent comparison describes its scale rather than a percentage change in raw revenue.

The remaining outcomes reinforce the contrast. Annual-crop revenue falls by 0.524 on the transformed scale, about 8 percent of its sample mean. Banana output falls by 1.645 tonnes per hectare, or 78 percent of its zero-inclusive sample mean. Cacao output falls by 0.047 tonnes per hectare, or 73 percent of its much smaller sample mean. These comparisons include non-producers observed in the relevant crop questionnaire. They therefore capture both crop adoption and output among producers.⁷

The crop pattern matches the institutional contrast. Decree 1001 required recipients of expropriated land in rice-growing zones to continue producing rice. PLT gave producers more room to move beyond the established rice, cacao, and banana economy. Work on Taiwan also shows why higher rice yields need not capture the broader development effects of land reform. Some reform phases raised rice yields through narrow channels rather than through the forces associated with the classic miracle narrative (Kim and Wang, 2025). In Ecuador, output per hectare remained higher in some established crops outside PLT, while PLT shifted land and revenue to-

⁷ The perennial-crop census records crop age incompletely. Many raw banana observations lack age, plausibly because respondents struggle to date ratoon and continuously renewed plants. I define perennial-questionnaire membership before applying the age restriction. I retain records with missing age and exclude only reported ages of at least 37 years. Plots therefore remain in the perennial-questionnaire sample even when no crop record survives the age screen. The remaining age evidence is suggestive. Among banana observations with reported age in the final local sample, bananas are slightly older on average outside the PLT footprint. The local RD points in the same direction. In the raw perennial file, older banana observations linked to the outside parishes also tilt in that direction (34 versus 12 under PLT). I treat banana age as supportive rather than decisive evidence of persistence.

ward a broader perennial portfolio. These opposing margins explain why aggregate perennial revenue changes little even though crop allocation changes sharply.

Table 1: Main RD results

	Perennial log revenue/ha	Annual log revenue/ha	Rice yield	Cacao yield	Banana yield	Main other perennials log revenue/ha
PLT-side jump	-0.172 (0.299) [0.317]	-0.524*** (0.165) [0.168]	-0.661*** (0.210) [0.226]	-0.047*** (0.012) [0.012]	-1.645*** (0.539) [1.310]	1.239*** (0.189) [0.316]
Observations	5590	4364	4364	5590	5590	5590
Parishes	26	26	26	26	26	26
Bandwidth (km)	40	40	40	40	40	40
Mean dep. var.	3.910	6.432	2.238	0.064	2.106	2.316

Every specification controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield.


The sample follows the questionnaire used to construct each outcome: perennial outcomes use plots observed in the perennial-crop questionnaire, while annual and rice outcomes use plots observed in the annual-crop questionnaire. Zeros are assigned only within the relevant questionnaire. Perennial-questionnaire membership is defined before the crop-age restriction; missing ages remain, and only reported ages of 37 years or more are excluded. Revenue outcomes use fixed Ecuador 2013 producer prices. Cacao and banana yields divide plot-level total production by total usable crop area.

The perennial, annual, and main-other-perennial outcomes are $\log(1 + \text{fixed-price revenue per hectare})$. The main other-perennial basket contains sugar cane, oil palm, and plantain. These coefficients are changes in transformed outcomes, not percentage changes in untransformed revenue. Rice, cacao, and banana are physical output per hectare, with zeros for non-producers observed in the relevant crop questionnaire.

Robust bias-corrected estimates following [Calonico et al. \(2014\)](#). Parentheses report parish-clustered standard errors; brackets report Huber–White standard errors. Stars use parish-clustered inference.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 6: Main RD evidence



Note: The dots show binned plot-level means within 100 km of the IERAC border. Positive distances denote the PLT side. The lines come from separate linear fits on each side. The outer gray lines mark 95 percent confidence intervals based on standard errors clustered by parish. These visual fits are unadjusted. The preferred estimates in Table 1 use a 40 km bandwidth, exclude Zapotal and Palenque, and include latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield. The annual-crop and diversification outcomes use $\log(1 + \text{fixed-price revenue per hectare})$. Rice reports physical output per hectare. Each outcome assigns zero only to non-producers observed in the relevant crop questionnaire.

5.2. Portfolio shift and alternative measures of diversification

PLT farms devote substantially more land and revenue to sugar cane, oil palm, and plantain. The land-share coefficient is 0.143, a rise of 14.3 percentage points against a sample mean of 19.0 percent. The increase equals about three-quarters of that mean. The revenue-share coefficient is 0.224, or 22.4 percentage points against a mean of 40.7 percent. It equals 55 percent of the mean. These estimates show a large reweighting of the perennial portfolio.

The number of distinct perennial crops rises by 0.376 against a sample mean of 4.054, an increase of about 9 percent. This estimate is imprecise. The evidence therefore shows that farms placed more land and revenue in the diversification basket, but it does not establish that they cultivated more crop names. A farm can reallocate a large area without adding many crops, while a crop count gives the same weight to a minor activity and a major one. Appendix Figures [A6](#), [A7](#), and [A8](#) show the corresponding RD graphs.

Table 2: Alternative measures of the portfolio shift

	Main other perennial land	Main other perennial rev. share	No. perm. crops
PLT-side jump	0.143*** (0.033)	0.224*** (0.027)	0.376 (0.334)
Observations	5590	5590	5590
Parishes	26	26	26
Bandwidth (km)	40	40	40
Mean dep. var.	0.190	0.407	4.054

The first two columns measure the land and revenue shares devoted to sugar cane, oil palm, and plantain. The sharp geographic RD uses a 40 km window and excludes the two parishes closest to the cutoff. It uses plots observed in the perennial-crop questionnaire. Questionnaire membership is defined before the crop-age restriction; records with missing ages remain, and only reported ages of 37 years or more are excluded. The revenue-share outcome uses fixed Ecuador 2013 producer prices. Every specification controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield, and clusters standard errors by parish.

The first two outcomes are shares, so multiplying their coefficients by 100 gives percentage-point changes. The crop-count coefficient measures the change in the number of distinct perennial crops.

Robust bias-corrected estimates following [Calonico et al. \(2014\)](#). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5.3. Intensive and extensive margins

The combined outcomes mix two decisions: whether a farm produces a crop and how much it produces after entry. Table [3](#) separates them. The extensive-margin coefficients measure changes in the probability of positive production. The intensive-margin coefficients compare revenue or yield among positive producers. These margins describe where the discontinuity arises, but they do not add mechanically to the combined coefficient.

Entry drives the rice and diversification results. PLT raises the probability of producing sugar cane, oil palm, or plantain by 20.4 percentage points. Conditional revenue among producers

changes by only 0.069 on the transformed scale and is imprecise. PLT lowers the probability of producing rice by 14.6 percentage points and the probability of producing any annual crop by 7.8 percentage points. Conditional rice output differs little across the border. The reform environments therefore changed which crops farms entered more clearly than how much incumbent producers obtained per hectare.

Bananas follow a different pattern. PLT raises the probability of positive banana production by 5.4 percentage points, but banana output among positive producers falls by 28.343 tonnes per hectare. This intensive-margin difference drives the combined decline of 1.645 tonnes per hectare once non-producers enter the outcome as zeros. Cacao declines on both margins: the probability of production falls by 25.6 percentage points, and output among producers falls by 0.044 tonnes per hectare.

The margin split connects the estimates to the institutional setting. Decree 1001 directly affected entry into rice, while flexible PLT rights expanded entry into sugar cane, oil palm, and plantain. The banana result concerns output among a selected group of producers. Historical expertise, producer selection, local conditions, and land concentration could all contribute to that difference. The design does not rank those explanations.

Table 3: Main RD results by combined, extensive, and intensive margins

Outcome	Combined	Extensive	Intensive
Perennial revenue per hectare	−0.172 (0.299)	−0.018 (0.040)	−0.289* (0.156)
Annual revenue per hectare	−0.524*** (0.165)	−0.078*** (0.018)	−0.046 (0.080)
Rice yield	−0.661*** (0.210)	−0.146*** (0.038)	−0.124 (0.212)
Cacao yield	−0.047*** (0.012)	−0.256*** (0.073)	−0.044*** (0.015)
Banana yield	−1.645*** (0.539)	0.054*** (0.020)	−28.343*** (2.622)
Sugar cane, oil palm, and plantain	1.239*** (0.189)	0.204*** (0.029)	0.069 (0.118)

Note: The combined column reports the baseline outcome construction, which assigns zeros only to plots observed in the relevant crop questionnaire. I define perennial-questionnaire membership before applying the crop-age restriction, retain missing ages, and exclude only reported ages of 37 years or more. The extensive column indicates positive outcomes within that questionnaire. The intensive column reports the original outcome among positive producers. Aggregate perennial and annual outcomes use price-weighted revenue per hectare; rice, cacao, and banana use physical yields. Cacao and banana yields divide plot-level total production by total usable crop area. Parentheses report parish-clustered standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.4. Further evidence

The main estimates show a broader crop portfolio on the PLT side. Three supporting exercises ask whether this shift coincides with movement out of agriculture, better alignment between crops and productive potential, and stronger local economic activity. The data require a different model for each outcome, so I use these results to interpret the core RD rather than as additional estimates of the same effect.

The labor exercise compares cantons near the border across five censuses. The land-use exercise compares perennial and other agricultural land within small grid cells. The growth exercise returns to the geographic RD and studies parish-level night lights at several bandwidths. Together, they trace three possible development margins without imposing one causal chain from diversification to local growth.

5.4.1. Structural transformation

IPUMS identifies location at the canton level, so I aggregate reform exposure from parishes to cantons. I pool the 1962, 1982, 1990, 2001, and 2010 censuses, compare cantons within 50 km of the border, and include census-year fixed effects. The comparison uses 21 cantons and clusters uncertainty at that level.

The agricultural employment share is 13.6 percentage points lower on the PLT side. This difference equals 23 percent of the sample mean of 59.3 percent. The canton-clustered standard error is 7.2 percentage points, and a bootstrap designed for 21 clusters gives a p -value of 0.098. Manufacturing employment rises by 2.6 percentage points, about half of its 5.3 percent sample mean, but the estimate is imprecise. The urban share falls by 2.7 percentage points, about 6 percent of its 44.2 percent sample mean, and is also imprecise. Table 4 and Figure 7 therefore show movement out of agriculture without identifying the destination sector. The canton-level evidence supports labor reallocation as a possible companion to crop diversification, but it does not establish that diversification caused that movement.

Table 4: Supporting labor-allocation indicators


	Ag. share	Mfg. share	Urban
PLT indicator	-0.136* (0.072)	0.026 (0.018)	-0.027 (0.095)
Observations	64	64	53
Cantons	21	21	21
Mean dep. var.	0.593	0.053	0.442
Bandwidth (km)		50	

Outcome columns are the agricultural employment share, manufacturing employment share, and urban-residence share among the active labor force. The coefficient is the PLT indicator from pooled canton-level ordinary least squares. I construct the canton-year data by merging harmonized IPUMS person and household records one-to-one by census sample, household serial number, and person number. The construction uses person weights.

The outcomes are shares, so multiplying a coefficient by 100 gives the change in percentage points.

Every specification includes census-year fixed effects, restricts cantons to 50 km from the border, and clusters standard errors by canton. The sample uses the 1962, 1982, 1990, 2001, and 2010 censuses and excludes 1974. The urban indicator is unavailable in 1982. A null-imposed wild-cluster bootstrap with 9,999 Rademacher replications gives $p=0.098$ for the agricultural-employment coefficient. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 7: Agricultural employment near the IERAC border



Note: The dots show binned canton-year means for the agricultural employment share within 100 km of the IERAC border. Positive distances denote the PLT side. The lines come from separate linear fits on each side. The outer gray lines mark 95 percent confidence intervals based on standard errors clustered by canton. The graph pools the 1962, 1982, 1990, 2001, and 2010 censuses. I treat this evidence as supportive because treatment varies across only 21 cantons. Table 4 reports the 50 km pooled estimates with year fixed effects and canton-clustered inference.

5.4.2. Pixel-level land-use conflict and alignment

The 2021 Ministry of Agriculture and Livestock data indicate whether observed land use fits local productive potential. I compare perennial land with other agricultural land inside the same grid cell, which holds fixed the cell's permanent characteristics. The interaction between the PLT side and perennial status measures whether perennial land gains an additional alignment advantage under PLT. I cluster uncertainty by parish. The first estimate uses all agricultural records outside the two excluded parishes. The second restricts the sample to 50 km from the

border.

In the full sample, the alignment advantage of perennial over other agricultural land is 20.3 percentage points larger on the PLT side. This estimate equals 83 percent of the outcome mean of 24.5 percent. Within 50 km, the corresponding advantage is 14.1 percentage points, or 67 percent of the mean of 21.1 percent. Both estimates differ statistically from zero. Table 5 therefore links the crop shift to economically large differences in land-use alignment. Because the 2021 model compares crop types within cells rather than locations at the policy cutoff, it serves as supporting evidence rather than a second estimate of the main RD effect.

Table 5: Pixel-level land-use alignment in MAGAP 2021

	Whole sample	Within 50 km
PLT side $\times$ perennial crop	0.203*** (0.050)	0.141** (0.060)
Observations	74431	38417
Parishes	116	32
Grid cells	421	151
Mean dep. var.	0.245	0.211
Grid-cell fixed effects	Yes	Yes

The outcome indicates good land use in the 2021 Ministry of Agriculture and Livestock data. Each column reports the interaction between the PLT side and a perennial-crop cell.

The outcome is an indicator, so multiplying a coefficient by 100 gives the change in percentage points. The linear probability model includes grid-cell fixed effects, excludes the two parishes closest to the cutoff, and clusters standard errors by parish.


* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.4.3. Economic growth

NOAA night-light rasters provide a parish-level measure of local economic activity in 1992, 2001, and 2010. I estimate the same border-distance design at 40, 75, and 100 km and exclude the two parishes closest to the cutoff. The regressions control for latitude, longitude, agricultural suitability, elevation, and precipitation. The wider windows show whether the direction depends on the narrowest local sample.

All nine estimates are positive, so each comparison places more night-light activity on the PLT side. Precision varies. The 90 percent intervals exclude zero for 1992 at 40 km and for 2001 and 2010 at 100 km. The other six intervals include zero. Figure 8 therefore provides directionally consistent but imprecise evidence of stronger local activity. Night lights combine many economic channels, so they cannot attribute this difference to crop diversification alone.

Figure 8: Night-light evidence



Note: The figure reports bias-corrected local-linear RD estimates for parish-level night-light measures in 1992, 2001, and 2010, using bandwidths of 40, 75, and 100 km and excluding the two parishes closest to the cutoff. Every estimate controls for latitude, longitude, general agricultural suitability, elevation, and precipitation. The intervals cover 90 percent and use heteroskedasticity-robust standard errors.

5.4.4. A simple productivity-reallocation calibration

The aggregate agricultural estimate combines gains and losses across crops and omits productivity gains from movement out of agriculture. A weak aggregate result can therefore coexist with diversification into sugar cane, oil palm, and plantain and with labor reallocation toward more productive sectors.

I illustrate this possibility with a back-of-the-envelope adaptation of [Kim and Wang \(2025\)](#). I first scale the within-agriculture productivity proxy by agriculture's employment share. I then add the potential gain implied by the lower agricultural employment share on the PLT side, following the reallocation logic in [McCaig and Pavcnik \(2018\)](#). Ecuador's 2000 GGDC/UNU-WIDER Economic Transformation Database supplies the national employment share and the productivity gap between agriculture and the rest of the economy ([Kruse et al., 2023](#)). The residual-productivity proxy follows [Montero \(2022\)](#) and [Aragon et al. \(2020\)](#).

The scaled within-agriculture term is -0.181 . Labor reallocation contributes 0.099 and offsets 55 percent of that negative term. The combined index remains negative at -0.082 . Aggregate productivity therefore remains weak, but the headline masks a sizable positive reallocation margin.

The crop scenarios expose the same offset within agriculture. Inserting the sugar cane, oil palm, and plantain coefficient produces a positive combined index of 0.490 . Inserting the rice

coefficient gives -0.110 , while the banana coefficient gives -0.420 . These values are accounting scenarios, not estimated percentage changes in aggregate labor productivity. They show how diversification and labor reallocation can move favorably while established crops pull the aggregate in the opposite direction. Appendix C.7 reports the equations and full arithmetic.

5.5. Robustness

The main crop pattern survives the robustness checks. Rice and banana remain negative when I use the common questionnaire sample, change the bandwidth or polynomial, or include the two parishes nearest the border. Sugar cane, oil palm, and plantain remain positive, including under broader definitions of the basket. Statistical precision varies, but the signs that define the paper’s comparison remain stable. Aggregate perennial revenue changes sign and remains imprecise because it combines crop margins that move in opposite directions.

Nearby farms may share weather, market access, or other unobserved shocks. I therefore recalculate uncertainty while allowing regression errors to remain correlated within 25 km. Because the main bias-corrected estimator cannot incorporate this correction, I approximate the preferred design with the same 40 km window, controls, and triangular weights. The central crop signs remain unchanged. The boundary method of Cattaneo et al. (2026) also preserves those signs when it measures location along the policy border. This holds both with and without the two nearest parishes, as Table A10 shows.

Randomization inference asks whether direct reassignment tests preserve the crop pattern. At the plot level, rice and banana remain negative and sugar cane, oil palm, and plantain remain positive. I then aggregate to the parish level so that each parish contributes one observation. Rice stays negative and the diversification basket stays positive in every fixed window. Banana remains negative but becomes imprecise. Annual crops and cacao vary more. These exercises compare covariate-adjusted means and report finite-sample p -values, so their coefficients need not reproduce the local-linear RD exactly.⁸

The robustness evidence preserves the paper’s qualitative conclusion: PLT-side areas specialize less in rice, produce less banana output per hectare, and allocate more production to sugar cane, oil palm, and plantain. Magnitudes and significance levels change across procedures. Aggregate perennials, annual crops, and cacao show less stability. The evidence therefore supports the comparative crop pattern rather than a claim that every agricultural outcome moved uni-

⁸ The randomization-inference tables compare adjusted differences in means within fixed windows and then calculate finite-sample randomization p -values.

formly.

The appendix reports the overlap sample in Table A1 and Figure A2. Tables A2 and A3 vary the bandwidth and polynomial. Figure A5 compares the mapped polygon with administrative PLT exposure, while Table A6 describes Decree 1001 and mapped reform footprints on each side. Tables A7 and A8 report the fuzzy-RD robustness exercises. Figure A3 and Table A4 report the density test. Figure A4 and Table A5 report the worker diagnostic. Tables A9 and A10 address spatial correlation and distance along the border.

The appendix also reports the calibration in Table A18 and Appendix C.7. Table A17 contains the Manabi benchmark. Tables A11 and A12 report randomization inference. Tables A13, A14, and A15 report the annual-crop credit check, estimates that retain the two closest parishes, and the diversification decomposition. Figures A6–A8 show the portfolio outcomes. The calibration and Manabi comparison support interpretation rather than identification.

6. Mechanisms

The two reform environments directed adaptation toward different crops. PLT-side areas shifted toward sugar cane, oil palm, and plantain. Outside PLT, Decree 1001 tied access to expropriated land in rice-growing zones to continued rice production. The informal-credit pattern is consistent with persistence along that direct route. Adoption drives both the rice result and the diversification result. Output per hectare among positive producers drives the banana result. The historical record documents accumulated banana expertise. The conditional-yield estimates allow for producer selection, and the cadastre shows greater banana-land concentration outside PLT. Together, these facts make persistence in the established banana sector plausible, but the design cannot separate their contributions.

6.1. Mandated rice specialization

Decree 1001 created the direct route to specialization by tying access to expropriated land to rice production. Credit markets may then have helped producers sustain that choice. Whereas much of the property-rights literature asks whether stronger rights expand formal finance and investment, producers under expropriation report more informal credit use.

Informal credit favors short production horizons. It often requires faster repayment and offers fewer opportunities to renegotiate than formal bank credit. Those features suit rice better than experimentation with a perennial portfolio. In the full sample of agricultural units, reported bor-

rowing from informal moneylenders is 5.6 percentage points lower on the PLT side, while formal borrowing does not change. Parish-clustered inference rejects zero at the 1 percent level. Standard errors robust to unequal residual variance reject zero at the 10 percent level. The estimate attenuates when I restore the two boundary-adjacent parishes, so I treat it as suggestive.

Annual-crop producers sharpen the pattern. Among agricultural units observed in the annual-crop questionnaire, informal credit use is 5.3 percentage points lower on the PLT side, and parish-clustered inference rejects zero at the 5 percent level. Annual crops turn over quickly, so short-horizon informal credit fits them more readily than perennials. The estimate strengthens the credit interpretation, but it does not show that informal credit caused rice persistence. Table 6 reports the main credit estimates, and Figure 9 plots the annual-revenue and informal-credit discontinuities. Appendix Table A13 reports the annual-crop credit estimates. The extensive-margin results more clearly show sustained participation in rice outside the PLT footprint than higher output per hectare among rice producers.

Table 6: Credit-market evidence


	Informal credit	Formal credit
PLT-side jump	-0.056*** (0.017) [0.031]	0.002 (0.013) [0.034]
Observations	6560	6560
Parishes	26	26
Mean dep. var.	0.061	0.080
Bandwidth (km)	40	40

The sample includes all agricultural units. Informal credit identifies loans from informal moneylenders (*chulqueros*); formal credit includes the Banco Nacional de Fomento, private banks, and credit cooperatives. Entries report the PLT-side jump from the preferred robust bias-corrected sharp RD following Calonico et al. (2014), with the two closest parishes excluded and a 40 km bandwidth.

The outcomes are indicators, so multiplying a coefficient by 100 gives the change in percentage points. The informal-credit estimate remains statistically distinguishable from zero at the 10 percent level with Huber–White standard errors but attenuates when the two closest parishes return. I therefore interpret it as suggestive.

Every specification controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield. Parentheses show parish-clustered standard errors; brackets show Huber–White standard errors. Stars use parish-clustered inference. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 9: Annual-crop revenue and informal credit near the policy border



Note: The dots show binned agricultural-unit means within 100 km of the IERAC border. Positive distances denote the PLT side. The lines come from separate linear fits on each side. The outer gray lines mark 95 percent confidence intervals based on standard errors clustered by parish. Panel (a) plots $\log(1 + \text{fixed-price annual-crop revenue per hectare})$. Panel (b) plots the share of agricultural units that report borrowing from informal moneylenders. Both panels exclude Zapotal and Palenque. The corresponding estimates use the preferred 40 km bandwidth. The graphs describe the discontinuities; they do not establish that informal credit caused rice specialization.

6.2. Banana yield and concentration

Bananas move differently from rice and the sugar cane, oil palm, and plantain basket. The extensive-margin estimate is slightly positive on the PLT side, but output per hectare among positive producers is substantially lower. Fewer producers do not drive the combined result. The difference comes from conditional yields. Because that regression selects on positive banana production, it may also reflect changes in producer composition.

Table 7 adds a descriptive fact. PLT areas contain larger banana producers on average, while banana land is more concentrated outside the PLT footprint in both the 40 km and 60 km windows. Higher yields outside PLT therefore do not reflect a uniformly larger banana sector. They coexist with a more concentrated one.

The data admit several interpretations. Established producers may have accumulated crop-specific knowledge. Land quality or selection into banana production may differ in ways that the controls do not capture. Tighter control over crop choice may also have discouraged experimentation with new perennials. The historical sequence and cadastre make this account plausible, but the design cannot separate it from the alternatives. The banana evidence therefore shows that strong performance in an established crop can coexist with weak portfolio diversification. It does not identify a separate property-rights mechanism.

Strong performance in incumbent crops can coexist with weak diversification. Decree 1001

Table 7: Banana-sector concentration in the banana cadastre

	40 km		60 km	
	Expro	PLT	Expro	PLT
Banana producers	94	383	118	383
Mean banana area	26.588	50.190	34.269	50.190
Median banana area	10.564	32.648	14.155	32.648
Banana-area Gini	0.644	0.551	0.616	0.551

The rows summarize the later banana cadastre from Ecuador's Ministry of Agriculture and Livestock: number of producers, mean banana area, median banana area, and the Gini coefficient of banana area. Columns compare the expropriation and PLT sides within 40 km and 60 km of the policy boundary. The table is descriptive rather than RD-based because the cadastre sample is highly imbalanced near the cutoff.

mandated rice, and the suggestive credit evidence is consistent with continued rice specialization among constrained producers. Higher banana yields and greater banana-land concentration, together with limited crop-age evidence, point to persistence outside PLT. Historical expertise, producer selection, and weaker incentives to change crops offer plausible channels, although the design cannot rank them. The PLT-side environment opened room for a broader perennial portfolio. The rights attached to redistributed land shaped crop allocation without raising every aggregate or incumbent-crop outcome.

7. Conclusion

Ecuador's agrarian reform transferred property with different rights, and those rights shaped crop choice for decades. In the coastal region, the state assigned PLT to one side of a historical policy boundary. Outside that footprint, the state relied more on expropriation, and Decree 1001 required recipients of expropriated land in rice-growing zones to continue producing rice. The resulting reform environments supported different patterns of agricultural specialization.

Decree 1001 sustained narrow specialization through a direct rice mandate. PLT gave producers more freedom to adopt sugar cane, oil palm, and plantain. These crops widened production beyond the established rice and banana economy. Adoption, rather than higher output per hectare within every incumbent crop, drives this result. Higher banana yields and greater banana-land concentration, together with limited crop-age evidence, point to persistence outside PLT. Historical expertise, producer selection, local conditions, and tighter control over crop choice offer plausible explanations, although the design cannot rank them. Supporting evidence connects the broader crop contrast to labor reallocation, land-use alignment, and local economic activity.

Land reform assigns rights and constraints as well as land. The reform environments differed

more clearly in crop choice than in output per hectare within every inherited crop. The PLT-side environment supported a broader rural portfolio, while high rice or banana productivity outside PLT coexisted with a narrower crop structure. Aggregate productivity can therefore obscure diversification gains when established crops move in the opposite direction. Suggestive evidence also points to better land-use alignment, stronger local activity, and movement out of agriculture on the PLT side.

The local design estimates this relative contrast. It does not isolate the absolute effect of either reform arm or each component of its rights bundle. Public land transfers often granted larger plots in other settings, yet plots have similar agricultural surface across the IERAC boundary. The comparison therefore studies neighboring reform environments within a common agrarian landscape rather than large and small farms. Once we distinguish reform designs, broad claims about agrarian reform become harder to sustain. Some findings attributed to reform as a whole may instead reflect the rights and production conditions attached to particular allocation strategies.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the author(s) used ChatGPT and Claude in order to correct spelling and grammar. Claude Code assisted in preparing a replication package for this paper. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the published article.

References

- Adamopoulos, T. (2008). Land inequality and the transition to modern growth. *Review of Economic Dynamics*, 11(2):257–282.
- Adamopoulos, T. and Restuccia, D. (2020). Land reform and productivity: A quantitative analysis with micro data. *American Economic Journal: Macroeconomics*, 12(3):1–39.
- Albertus, M. (2015). *Autocracy and Redistribution*. Cambridge University Press, Cambridge.
- Alchian, A. A. (1965). Some economics of property rights. *Il Politico*, 30(4):816–829.
- Aragon, F., Restuccia, D., and Rud, J. P. (2020). Are Small Farms Really more Productive than Large Farms ?
- Arosemena, G. (1991). *El fruto de los Dioses*, volume Volumen 1. Editorial Graba.
- Barreca, A. I., Guldi, M., Lindo, J. M., and Waddell, G. R. (2011). Saving babies? Revisiting the effect of very low birth weight classification. *Quarterly Journal of Economics*, 126(4):2117–2123.
- Barsky, O. (1984). *La Reforma Agraria Ecuatoriana*. Corporacion editora nacional, Quito.
- Besley, T. and Burgess, R. (2000). Land reform, poverty reduction, and growth: Evidence from India. *Quarterly Journal of Economics*, 115(2):389–430.
- Blankespoor, B., Ru, Y., Wood-sichra, U., Thomas, T. S., You, L., and Kalvelagen, E. (2022). Estimating Local Agricultural GDP across the World.
- Bühler, M. (2021). On the other side of the fence: Property rights and productivity in the United States.
- Calonico, S., Cattaneo, M. D., and Farrell, M. H. (2021). Optimal bandwidth choice for robust bias-corrected inference in regression discontinuity designs. *Econometrics Journal*, 23:192–210.
- Calonico, S., Cattaneo, M. D., Farrell, M. H., and Titiunik, R. (2017). Rdrobust: Software for regression-discontinuity designs. *Stata Journal*, 17(2):372–404.
- Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82:2295–2326.
- Canelos, F. (1980). *COLONIZACION y VIAS DE DESARROLLO, EL CASO DE LA PARROQUIA PEDRO VICENTE MALDONADO*. PhD thesis, FLACSO.

- Carrasco, H. (1994). *Campesinos y mercado de tierras en la costa ecuatoriana*. FAO, Rome.
- Cattaneo, M. D., Frandsen, B. R., and Titiunik, R. (2014). Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate. *Journal of Causal Inference*, 3(1):1–24.
- Cattaneo, M. D., Idrobo, N., and Titiunik, R. (2019). *A Practical Introduction to Regression Discontinuity Designs*.
- Cattaneo, M. D., Titiunik, R., and Yu, R. R. (2026). Estimation and Inference in Boundary Discontinuity Designs: Location-Based Methods.
- Center, M. P. (2020). Integrated Public Use Microdata Series, International: Version 7.3 [dataset].
- Chiriboga, M. (1980). *Jornaleros y gran propietarios en 135 años de Exportación cacaotera (1790-1925)*, volume volumen1. Consejo provincial de Pichincha.
- Cuvi, N. (2015). Las semillas del imperialismo agr  cola estadounidense en el Ecuador. *Procesos. Revista ecuatoriana de historia*, 1:69.
- de Janvry, A. (1981). *The Agrarian Question and Reformism in Latin America*. The Johns Hopkins Studies in Development. Johns Hopkins University Press.
- Dell, M. (2010). The Persistent Effects of Peru’s Mining Mita. *Econometrica*, 78(6):1863–1903.
- Do, Q. T. and Iyer, L. (2008). Land titling and rural transition in Vietnam. *Economic Development and Cultural Change*, 56:531–579.
- Edwards, E. C., Fiszbein, M., and Libecap, G. D. (2022). Property Rights to Land and Agricultural Organization: An Argentina United States Comparison. *Journal of Law and Economics*, 65(S1):S1–S33.
- Espinosa, R. (2014). *Desmemoria y olvido: la econom  a arrocer  a en la cuenca del Guayas, 1900-1950*. Universidad Andina Sim  n Bol  var / Corporaci  n Editora Nacional, Quito.
- Faguet, J. P., S  nchez, F., and Villaveces, M. J. (2020). The perversion of public land distribution by landed elites: Power, inequality and development in Colombia. *World Development*, 136:105036.
- Gal  n, J. S. (2020). Tied to the Land? Intergenerational Mobility and Agrarian Reform in Colombia.

- Galiani, S. and Schargrodsky, E. (2010). Property rights for the poor: Effects of land titling. *Journal of Public Economics*, 94(9-10):700–729.
- Guerrero, A. (1994). *Los oligarcas del cacao*. Editorial El Conejo, Quito.
- Herrendorf, B., Rogerson, R., and Valentinyi, Á. (2014). Growth and Structural Transformation. In *Handbook of Economic Growth*, volume 2, pages 855–941. Elsevier B.V.
- Hornbeck, R. (2010). Barbed Wire: Property Rights and Agricultural Development *. *Quarterly Journal of Economics*, 125(2):767–810.
- Instituto Nacional de Estadística y Censos (2002a). *III Censo Nacional Agropecuario: Resultados nacionales, incluye resúmenes provinciales. Volumen 1*. INEC, Quito.
- Instituto Nacional de Estadística y Censos (2002b). Síntesis metodológica del III Censo Nacional Agropecuario. Technical report, INEC, Quito.
- Kay, C. (2002). Why east asia overtook latin america: Agrarian reform, industrialisation and development. *Third World Quarterly*, 23:1073–1102.
- Kay, C. (2019). Land reform in latin america: Past, present, and future. *Latin American Research Review*, 54:747–755.
- Keele, L. J. and Titiunik, R. (2015). Geographic boundaries as regression discontinuities. *Political Analysis*, 23(1):127–155.
- Kim, O. and Wang, J.-K. (2025). Roots of the Taiwanese Miracle? Reassessing Land Reform, 1950-1961.
- Kitamura, S. (2022). Tillers of Prosperity: Land Ownership, Reallocation, and Structural Transformation. (November):42.
- Kruse, H., Mensah, E., Sen, K., and de Vries, G. (2023). A Manufacturing (Re)Naissance? Industrialization in the Developing World. *IMF Economic Review*, 71(2):439–473.
- Larrea, C. (1987). *El banano en el Ecuador: Transnacionales, modernización y Subdesarrollo*. FLACSO, Quito.
- Leonard, O. (1947). Pichilingue: A study of rural life in coastal ecuador.

- McCaig, B. and Pavcnik, N. (2018). Export Markets and Labor Allocation in a Low-Income Country. *American Economic Review*, 108(7):1899–1941.
- Montero, E. (2022). Cooperative Property Rights and Development: Evidence from Land Reform in El Salvador. *Journal of Political Economy*, 130(1):48–93.
- Montero, E. (2023). Erratum: Cooperative Property Rights and Development: Evidence from Land Reform in El Salvador. *Journal of Political Economy*, 131(8):2286–2287.
- Nunn, N. and Puga, D. (2012). Ruggedness: The Blessing of Bad Geography in Africa. *Review of Economics and Statistics*, 94(1):20–36.
- Redclift, M. R. (1976). Agrarian class structure and the State : the case of coastal Ecuador. *Boletín de Estudios Latinoamericanos y del Caribe*, 21(21):16–31.
- Smith, C. (2020). Land Concentration and Long-Run Development in the Frontier United States. *Working Paper*.
- Uggen, J. (1975). *PEASANT MOBILIZATION IN ECUADOR: A CASE STUDY OF GUAYAS PROVINCE*. PhD thesis, University of Miami.


A. Appendix Roadmap

The appendix presents the shared pre-reform history, outcome-construction details, robustness tests, and auxiliary identification and interpretation checks.

B. Historical Appendix

The historical figures and tables document the shared post-cocoa context across the border and provide the chronology, map evidence, and archival detail behind the shorter main-text discussion.

Figure A1: Original 1978 IERAC archival map



Note: The 1978 IERAC map records the realized spatial footprint of agrarian reform. I digitize the outer edge of its PLT polygon to construct the RD cutoff used in the main text.

C. Empirical Appendix

The empirical appendix reports alternative samples and outcome rules, bandwidth and polynomial checks, parish-level density tests, the no-doughnut RD, segment fixed effects, and diversification decompositions. In its tables, “main other perennials” refers to sugar cane, oil palm, and plantain. Auxiliary within-expropriation exercises help interpret the rice and banana margins but do not belong to the core design.

C.1. Overlap-sample robustness

Table A1: Overlap-sample robustness estimates

	Perennial log revenue/ha	Annual log revenue/ha	Rice yield	Cacao yield	Banana yield	Main other perennials log revenue/ha
PLT-side jump	0.332 (0.296) [0.371]	-0.328* (0.168) [0.174]	-0.611*** (0.205) [0.242]	-0.045*** (0.014) [0.015]	-1.126** (0.538) [1.277]	1.768*** (0.132) [0.388]
Observations	3632	3632	3632	3632	3632	3632
Parishes	26	26	26	26	26	26
Bandwidth (km)	40	40	40	40	40	40
Mean dep. var.	3.655	6.388	2.178	0.073	1.510	2.123

Every specification controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield.


This robustness check restricts every outcome to plots observed in both the annual- and perennial-crop questionnaires and assigns zeros only within that common sample. Perennial records retain missing crop ages and exclude only reported ages of 37 years or more. Revenue outcomes use fixed Ecuador 2013 producer prices; cacao and banana yields use plot-level production divided by usable crop area.

The perennial, annual, and main-other-perennial outcomes are $\log(1 + \text{fixed-price revenue per hectare})$. The main other-perennial basket contains sugar cane, oil palm, and plantain. These coefficients are changes in transformed outcomes, not percentage changes in untransformed revenue. Rice, cacao, and banana are physical output per hectare, with zeros for non-producers observed in the relevant crop questionnaire.

Robust bias-corrected estimates following [Calonico et al. \(2014\)](#). Parentheses report parish-clustered standard errors; brackets report Huber–White standard errors. Stars use parish-clustered inference.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure A2: Overlap-sample robustness summary



Note: These estimates restrict every outcome to plots observed in both the annual- and perennial-crop questionnaires. The main analysis instead uses the questionnaire relevant to each outcome.

C.2. Specification robustness

Table A2: Bandwidth robustness

	Perennial	Annual	Rice	Cacao	Banana	Main other perennials
<i>(1) $h = 35$ km, parish-clustered standard errors</i>						
PLT-side jump	0.527** (0.215)	-0.283** (0.130)	-0.627*** (0.226)	-0.030*** (0.010)	-1.212** (0.590)	1.727*** (0.109)
Observations	4760	3863	3863	4760	4760	4760
<i>(2) $h = 40$ km, parish-clustered standard errors</i>						
PLT-side jump	-0.172 (0.299)	-0.524*** (0.165)	-0.661*** (0.210)	-0.047*** (0.012)	-1.645*** (0.539)	1.239*** (0.189)
Observations	5590	4364	4364	5590	5590	5590
<i>(3) $h = 50$ km, parish-clustered standard errors</i>						
PLT-side jump	0.138 (0.390)	-0.455** (0.208)	-0.607** (0.268)	-0.039*** (0.012)	-1.658*** (0.630)	1.373*** (0.229)
Observations	6967	5675	5675	6967	6967	6967

Columns report estimates at fixed 35, 40, and 50 km bandwidths. Each outcome uses its relevant crop questionnaire, and zeros are assigned only to plots observed in that questionnaire. Perennial-questionnaire membership is defined before the crop-age restriction; missing ages remain, and only reported ages of 37 years or more are excluded. Cacao and banana yields divide plot-level total production by total usable crop area.

Every sharp geographic RD uses robust bias-corrected inference following Calonico et al. (2014), excludes the two parishes closest to the cutoff, and controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield. Parentheses report parish-clustered standard errors. Fixed bandwidths provide a transparent comparison because automatic selection can be unstable with parish-level mass points. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: Local-quadratic robustness

	Perennial	Annual	Rice	Cacao	Banana	Main other perennials
PLT-side jump	-0.214 (0.504) [0.433]	-0.752*** (0.187) [0.222]	-1.156*** (0.346) [0.300]	-0.060*** (0.015) [0.016]	-1.001 (1.014) [1.750]	1.125*** (0.291) [0.424]
Observations	5590	4364	4364	5590	5590	5590
Parishes	26	26	26	26	26	26
Bandwidth (km)	40	40	40	40	40	40
Mean dep. var.	3.910	6.432	2.238	0.064	2.106	2.316

Every specification controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield.

This check uses the main outcome-specific questionnaire samples, the doughnut exclusion, and a 40 km bandwidth, but fits a local quadratic polynomial instead of the baseline local linear specification. Zeros are assigned only to plots observed in the questionnaire used to construct each outcome.

Perennial, annual, and main-other-perennial revenue use $\log(1 + \text{fixed-price revenue per hectare})$, so their coefficients are not percentage changes in untransformed revenue. Rice, cocoa, and banana report physical output per hectare.

Robust bias-corrected estimates follow [Calonico et al. \(2014\)](#) and the bandwidth refinements in Calonico et al. (2020). Parentheses report parish-clustered standard errors; brackets report Huber–White standard errors. Stars use parish-clustered inference.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure A3: Parish-level running-variable density around the cutoff





Table A4: Parish-level running-variable density around the cutoff statistics

Statistic	Conventional	Robust
Test statistic	-0.432	1.108
<i>p</i> -value	0.665	0.268
Obs. left / right		26 / 18
Eff. obs. left / right		25 / 18
Bandwidth		59.657 km

The density test uses one observation per parish centroid within 60 km of the IERAC boundary. Figure A3 plots the same parish-level running variable.

The conventional and robust columns report the uncorrected and bias-corrected versions of the local-polynomial density-discontinuity test. This sorting diagnostic complements the paper's robust bias-corrected outcome RD following Calonico et al. (2014).

Figure A4: Post-treatment farm-labor-scale diagnostic




Note: The dots show binned agricultural-unit means within 100 km of the IERAC border. The lines come from separate linear fits on each side. The outer gray lines mark 95 percent confidence intervals based on standard errors clustered by parish. The outcome is the number of workers reported per agricultural unit in the 2000 agricultural census. Because the outcome is post-treatment, it cannot test pre-reform sorting. It only shows whether observed farm-labor scale changes at the policy border in 2000. The corresponding estimate uses a 40 km bandwidth.

Table A5: Post-treatment agricultural-worker diagnostic

	Number of workers
PLT-side jump	0.245 (0.474)
Observations	5505
Parishes	28
Bandwidth (km)	40
Mean dep. var.	4.133

The sharp geographic RD uses robust bias-corrected inference following [Calonico et al. \(2014\)](#) for the number of workers recorded in the 2000 agricultural census. It estimates the discontinuity within a 40 km bandwidth from observations located within 100 km of the IERAC boundary and clusters standard errors by parish. Because workers are post-treatment, this outcome cannot test pre-reform sorting; it only diagnoses whether observed farm-labor scale differs locally at the boundary in 2000. The specification does not add covariates or apply the doughnut exclusion. Figure companion: Figure [A4](#). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A5: Administrative PLT correspondence at the polygon edge



Note: Panel (a) plots the share of parish land assigned through PLT in the reconstructed IERAC records. Panel (b) uses that share in a parish-level probit model and plots the fitted probability that a parish lies inside the digitized PLT polygon. Each dot represents a binned parish mean. The lines come from separate linear fits on each side. The outer gray lines mark 95 percent confidence intervals based on heteroskedasticity-robust standard errors. Each underlying observation is one parish, so no repeated observations remain within parish. Both panels use parish-centroid distance to the polygon's outer edge and display the full sample within 100 km. They retain the two parishes closest to the border. Positive distances lie inside the polygon. The reconstruction matches historical reports to current parishes and assigns later parishes the exposure of their historical parent jurisdiction when required. The panels show correspondence between administrative reports and mapped geography. They do not provide a separate treatment assignment.

Table A6: Treatment composition around the PLT footprint

	Full 40 km corridor		Doughnut corridor	
	Outside PLT	PLT side	Outside PLT	PLT side
Parishes	13	15	12	14
Parishes with a Decree 1001 award	10	1	9	0
Decree 1001 adjudications	57	1	56	0
Decree 1001 hectares	9,280	1,064	9,192	0
Decree 1001 beneficiary households	1328	20	1328	0
Parishes overlapping mapped expropriation	13	5	12	4
Mean expropriation footprint share	0.168	0.019	0.161	0.005
Parishes overlapping mapped PLT	3	15	2	14
Mean PLT footprint share	0.010	0.858	0.001	0.900

The table describes the reform evidence within 40 km of the digitized PLT boundary. Positive distance defines the PLT side. The doughnut corridor excludes Zapotal and Palenque, the two parishes closest to the historical edge. Decree 1001 entries come from the digitized IERAC award registry. Footprint measures come from overlays of the mapped expropriation and PLT polygons with parish boundaries. The registry and polygons measure different objects, so the table describes treatment composition rather than a first stage or a new causal design. Manga del Cura was a nondelimited census territory that joined Manabi in 2015. I reconstruct it from the three current El Carmen units that contain its historical census sectors: El Carmen, Santa Maria, and El Paraiso La 14. Their combined territory is 98.8 percent covered by the mapped PLT footprint and contains no recorded Decree 1001 award.

Table A7: Fuzzy RD robustness: Full sample

	Perennial	Annual	Rice	Cacao	Banana	Main other perennials
First stage	0.410*** (0.105)	0.432*** (0.097)	0.432*** (0.097)	0.410*** (0.105)	0.410*** (0.105)	0.410*** (0.105)
Fuzzy RD estimate	1.190 (1.219)	-0.678* (0.395)	-1.651*** (0.383)	-0.035 (0.043)	-3.551** (1.425)	3.143*** (1.025)
Observations	6020	4799	4799	6020	6020	6020
Parishes	28	28	28	28	28	28
Bandwidth (km)	40	40	40	40	40	40
Mean dep. var.	3.843	6.475	2.248	0.063	2.061	2.236

The table reports a fuzzy geographic RD that instruments the reconstructed administrative PLT land share with location inside the digitized PLT polygon. The dependent variables are $\log(1 + \text{revenue per hectare})$ for the aggregate crop groups and the paper's zero-inclusive transformed yield measures for individual crops. Each column uses the relevant annual- or perennial-crop questionnaire and assigns zeros only to plots observed in that questionnaire. Coefficients are changes in transformed outcomes, not percentage changes in unlogged productivity.

The full sample does not exclude the two parishes closest to the cutoff. Every specification uses a 40 km bandwidth and controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield. Parentheses report parish-clustered standard errors. The first-stage and fuzzy coefficients use robust bias-corrected inference following [Calonico et al. \(2014\)](#). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A8: Fuzzy RD robustness: Doughnut sample

	Perennial	Annual	Rice	Cacao	Banana	Main other perennials
First stage	0.506*** (0.196)	0.573** (0.242)	0.573** (0.242)	0.506*** (0.196)	0.506*** (0.196)	0.506*** (0.196)
Fuzzy RD estimate	-0.619 (0.976)	-0.966** (0.465)	-1.215* (0.694)	-0.102*** (0.033)	-3.234* (1.752)	2.250 (1.735)
Observations	5590	4364	4364	5590	5590	5590
Parishes	26	26	26	26	26	26
Bandwidth (km)	40	40	40	40	40	40
Mean dep. var.	3.910	6.432	2.238	0.064	2.106	2.316

The table reports a fuzzy geographic RD that instruments the reconstructed administrative PLT land share with location inside the digitized PLT polygon. The dependent variables are $\log(1 + \text{revenue per hectare})$ for the aggregate crop groups and the paper's zero-inclusive transformed yield measures for individual crops. Each column uses the relevant annual- or perennial-crop questionnaire and assigns zeros only to plots observed in that questionnaire. Coefficients are changes in transformed outcomes, not percentage changes in unlogged productivity.

The doughnut sample excludes the two parishes closest to the cutoff. Every specification uses a 40 km bandwidth and controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield. Parentheses report parish-clustered standard errors. The first-stage and fuzzy coefficients use robust bias-corrected inference following [Calonico et al. \(2014\)](#). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A9: OLS-Conley spatial-correlation robustness

	Perennial	Annual	Rice	Cacao	Banana	Main other perennials
OLS + Conley	0.646 (0.467)	-0.255* (0.152)	-0.340* (0.204)	-0.017 (0.015)	-1.470*** (0.414)	1.632*** (0.200)
Observations	5590	4364	4364	5590	5590	5590
Parishes	26	26	26	26	26	26
Bandwidth (km)	40	40	40	40	40	40

Each outcome uses its relevant crop questionnaire, and zeros are assigned only to plots observed in that questionnaire. Perennial-questionnaire membership is defined before the crop-age restriction; missing ages remain, and only reported ages of 37 years or more are excluded. Cacao and banana yields divide plot-level total production by total usable crop area. The coefficient is the PLT-side jump from a triangular-weighted ordinary-least-squares approximation to the preferred local-linear RD within the 40 km window.

The specification excludes the two parishes closest to the cutoff and controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield. Parentheses report Conley standard errors that allow spatial correlation within 25 km. The spatial covariance estimate with a 40 km cutoff is not positive semidefinite for banana yield, so the table uses the shorter common cutoff that remains defined for every outcome. Because this ordinary-least-squares spatial-correlation check differs from the main robust bias-corrected estimator following [Calonico et al. \(2014\)](#), its coefficients need not equal the main RD estimates. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A10: Distance-based boundary-discontinuity comparison

	Main RD	Boundary RD doughnut	Boundary RD no doughnut
Perennial	-0.172 (0.299)	0.297* (0.168)	0.979*** (0.233)
Annual	-0.524*** (0.165)	-0.349*** (0.125)	-0.044 (0.101)
Rice	-0.661*** (0.210)	-0.948*** (0.241)	-0.774*** (0.093)
Cacao	-0.047*** (0.012)	-0.044*** (0.009)	-0.006 (0.013)
Banana	-1.645*** (0.539)	-0.896* (0.472)	-0.589* (0.321)
Main other perennials	1.239*** (0.189)	1.773*** (0.109)	1.586*** (0.088)
Observations	varies by outcome	varies by outcome	varies by outcome
Bandwidth (km)	40	40	40
Doughnut	yes	yes	no

Note: The first column reports the paper's main robust bias-corrected local-linear geographic RD estimates following [Calonico et al. \(2014\)](#). The second and third columns report a distance-based boundary-discontinuity estimator following Cattaneo, Titiunik, and Yu (2026), using a 40 km window with and without the doughnut exclusion. Each outcome uses its relevant crop questionnaire, and zeros are assigned only to plots observed in that questionnaire. Perennial-questionnaire membership is defined before the crop-age restriction; missing ages remain, and only reported ages of 37 years or more are excluded. Cacao and banana yields divide plot-level total production by total usable crop area.

All columns control for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield. Parentheses report standard errors clustered by parish. The boundary estimator treats the design as a boundary-discontinuity problem rather than a canonical one-dimensional local-linear RD, so its coefficients need not equal the main estimates. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A11: Residualized local-randomization inference in the preferred window

	Perennial	Annual	Rice	Cacao	Banana	Main other perennials
Residualized difference in means	0.097	0.029	-0.114	0.002	-0.637	0.223
Randomization p-value	0.156	0.456	0.022	0.616	0.014	0.000
Large-sample p-value	0.206	0.502	0.021	0.543	0.014	0.001
Observations	5590	4364	4364	5590	5590	5590
Left of cutoff	2419	2440	2440	2419	2419	2419
Right of cutoff	3171	1924	1924	3171	3171	3171

Note: The table reports randomization inference for the outcome-specific questionnaire samples in the preferred 40 km doughnut window. I assign zeros only to plots observed in the questionnaire relevant to each outcome. I first remove the variation explained by latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield. I then compare the resulting residuals. I define perennial-questionnaire membership before applying the crop-age restriction. I retain missing ages and exclude only reported ages of 37 years or more. Cacao and banana yields divide plot-level total production by total usable crop area. The coefficient reports the local difference in means within the window. Its magnitude therefore does not directly match the bias-corrected local-linear RD coefficient in the main tables. The preferred continuity-based RD follows [Calonico et al. \(2014\)](#). This local-randomization procedure follows [Cattaneo et al. \(2014\)](#) and provides finite-sample p-values rather than clustered standard errors.

Table A12: Parish-level randomization inference across fixed windows

Outcome	10 km	15 km	20 km	25 km	30 km	35 km	40 km
<i>Doughnut sample</i>							
Perennial	-0.223 (0.996)	0.903 (0.263)	0.683 (0.125)	0.673 (0.113)	0.406 (0.156)	0.416 (0.140)	0.063 (0.841)
Annual	-0.275 (0.248)	-0.210 (0.267)	-0.052 (0.694)	0.001 (0.999)	0.087 (0.689)	-0.021 (0.917)	0.029 (0.871)
Rice	-0.885 (0.248)	-0.619 (0.091)	-0.122 (0.592)	-0.120 (0.554)	-0.254 (0.163)	-0.329 (0.063)	-0.304 (0.035)
Cacao	-0.035 (0.248)	0.001 (0.887)	0.002 (0.821)	0.005 (0.677)	0.000 (0.967)	-0.002 (0.798)	-0.002 (0.831)
Banana	-1.678 (0.488)	-0.210 (0.858)	-0.268 (0.686)	-0.252 (0.696)	-0.645 (0.356)	-0.593 (0.362)	-1.201 (0.110)
Main other perennials	0.766 (0.248)	0.992 (0.087)	0.791 (0.009)	0.685 (0.027)	0.567 (0.011)	0.650 (0.006)	0.334 (0.151)
Number of parishes	8	10	13	14	20	22	26
<i>No-doughnut sample</i>							
Perennial	0.173 (0.574)	0.739 (0.171)	0.599 (0.114)	0.596 (0.088)	0.373 (0.151)	0.383 (0.132)	0.058 (0.833)
Annual	-0.204 (0.186)	-0.179 (0.153)	-0.039 (0.737)	0.005 (0.959)	0.081 (0.702)	-0.018 (0.921)	0.032 (0.849)
Rice	-0.621 (0.045)	-0.528 (0.021)	-0.194 (0.308)	-0.184 (0.320)	-0.274 (0.099)	-0.335 (0.032)	-0.326 (0.019)
Cacao	-0.018 (0.274)	0.001 (0.914)	0.002 (0.797)	0.005 (0.629)	0.001 (0.932)	-0.002 (0.857)	-0.001 (0.909)
Banana	-0.846 (0.538)	-0.220 (0.838)	-0.444 (0.493)	-0.407 (0.506)	-0.715 (0.294)	-0.657 (0.295)	-1.268 (0.080)
Main other perennials	0.802 (0.090)	0.867 (0.026)	0.703 (0.012)	0.615 (0.030)	0.520 (0.011)	0.595 (0.008)	0.304 (0.170)
Number of parishes	10	12	15	16	22	24	28

Note: Each cell reports the difference in adjusted parish means within the displayed window. Parentheses contain p -values from 9,999 fixed-margin permutations. I use the questionnaire relevant to each outcome and assign zeros only to plots observed in that questionnaire. In the 40 km sample, I first remove the variation explained by latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield. I then collapse the resulting residuals to parish means and apply randomization inference within each window. The doughnut panel excludes the two parishes closest to the IERAC boundary. Treating parishes as the effective units allows each parish to contribute once and avoids treating plots within the same parish as independent. The small number of parishes limits precision. The preferred continuity-based RD follows [Calonico et al. \(2014\)](#). The local-randomization procedure follows [Cattaneo et al. \(2014\)](#).

C.3. Credit-market sample check

Annual crops turn over quickly, so short-horizon informal credit fits them more readily than perennials. Among producers observed in the annual-crop questionnaire, reported borrowing from informal moneylenders is 5.3 percentage points lower on the PLT side, and parish-clustered inference rejects zero at the 5 percent level. Formal borrowing does not change. This result supports the proposed financing channel, but it does not establish that informal credit caused rice specialization.

Table A13: Credit-market evidence among annual-crop producers

	Informal credit	Formal credit
PLT-side jump	-0.053** (0.026) [0.041]	-0.009 (0.018) [0.043]
Observations	4364	4364
Parishes	26	26
Mean dep. var.	0.088	0.087
Bandwidth (km)	40	40

The sample includes agricultural units observed in the annual-crop questionnaire. Informal credit identifies loans from informal moneylenders (*chulqueros*); formal credit includes the Banco Nacional de Fomento, private banks, and credit cooperatives. Entries report the PLT-side jump from the preferred robust bias-corrected sharp RD following [Calonico et al. \(2014\)](#), with the two closest parishes excluded and a 40 km bandwidth.

The outcomes are indicators, so multiplying a coefficient by 100 gives the change in percentage points. Parentheses show parish-clustered standard errors; brackets show Huber–White standard errors. Stars use parish-clustered inference. This annual-crop restriction is a mechanism check, not the baseline credit sample.

Every specification controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.4. Diversification robustness

Table A14: No-doughnut robustness

	Perennial	Annual	Rice	Cacao	Banana	Main other perennials
PLT-side jump	0.502 (0.413) [0.221]	-0.292* (0.170) [0.073]	-0.712*** (0.079) [0.153]	-0.014 (0.017) [0.008]	-1.457*** (0.401) [0.698]	1.300*** (0.185) [0.188]
Observations	6020	4799	4799	6020	6020	6020
Parishes	28	28	28	28	28	28
Bandwidth (km)	40	40	40	40	40	40
Mean dep. var.	3.843	6.475	2.248	0.063	2.061	2.236

Every specification controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield.

This robustness check uses the questionnaire relevant to each outcome and a 40 km bandwidth, but retains the two parishes closest to the cutoff. Zeros are assigned only to plots observed in the relevant questionnaire.

Perennial, annual, and main-other-perennial revenue use $\log(1 + \text{fixed-price revenue per hectare})$, so their coefficients are not percentage changes in untransformed revenue. Rice, cocoa, and banana report physical output per hectare.

Robust bias-corrected estimates follow [Calonico et al. \(2014\)](#) and the bandwidth refinements in Calonico et al. (2020). Parentheses show parish-clustered standard errors; brackets show Huber–White standard errors. Stars use parish-clustered inference.


* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A15: Decomposition of the sugar cane, oil palm, and plantain basket

Specification	PLT-side jump	Standard error	Mean dep. var.
Main basket	1.239***	(0.189)	2.316
+ coffee	0.826***	(0.191)	2.637
+ limon	1.194***	(0.189)	2.413
+ maracuya	0.753**	(0.328)	2.794
+ palmito	1.202***	(0.181)	2.354
+ pina	1.405***	(0.255)	2.413
Coffee only	-0.253***	(0.056)	0.771
Limon only	0.209	(0.128)	0.355
Palmito only	-0.043	(0.034)	0.082
Pina only	0.293	(0.266)	0.160


The main basket contains sugar cane (including cane for sugar and other uses), African oil palm, and plantain. Rows add one crop or retain that crop alone. Every row divides selected-crop revenue by total usable harvested area in perennial crops other than cacao and banana; productive area replaces harvested area when needed, followed by planted area. The dependent variable is $\log(1 + \text{revenue per hectare})$, valued with common Ecuador-wide 2013 producer prices, so its coefficient is not a percentage change in untransformed revenue. The preferred specification uses robust bias-corrected inference following [Calonico et al. \(2014\)](#), the perennial questionnaire, the 40 km bandwidth and doughnut, and the baseline controls.

Figure A6: RD graph: land share in sugar cane, oil palm, and plantain




Note: The dots show binned plot-level means within 100 km of the IERAC border. The lines come from separate linear fits on each side. The outer gray lines mark 95 percent confidence intervals based on standard errors clustered by parish. The outcome is the share of perennial land devoted to sugar cane, oil palm, and plantain among plots observed in the perennial-crop questionnaire. Positive distances denote the PLT side. The corresponding sharp geographic RD uses a 40 km bandwidth, excludes Zapotal and Palenque, and controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield.

Figure A7: RD graph: revenue share from sugar cane, oil palm, and plantain



Note: The dots show binned plot-level means within 100 km of the IERAC border. The lines come from separate linear fits on each side. The outer gray lines mark 95 percent confidence intervals based on standard errors clustered by parish. The outcome is the share of perennial revenue from sugar cane, oil palm, and plantain among plots observed in the perennial-crop questionnaire. Revenue uses fixed Ecuador-wide 2013 producer prices. Positive distances denote the PLT side. The corresponding sharp geographic RD uses a 40 km bandwidth, excludes Zapotal and Palenque, and controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield.

Figure A8: RD graph: number of perennial crops



Note: The dots show binned plot-level means within 100 km of the IERAC border. The lines come from separate linear fits on each side. The outer gray lines mark 95 percent confidence intervals based on standard errors clustered by parish. The outcome is the number of distinct perennial crops cultivated among plots observed in the perennial-crop questionnaire. Positive distances denote the PLT side. The corresponding sharp geographic RD uses a 40 km bandwidth, excludes Zapotal and Palenque, and controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield.

C.5. Landholding structure

Table A16 reports descriptive comparisons of landholding inequality across the two sides of the policy border.

Table A16: Landholding inequality from the agricultural census

	40 km window	60 km window
Pooled plot Gini, expropriation side	0.740	0.770
Pooled plot Gini, PLT side	0.763	0.731
Mean parish Gini, expropriation side	0.718	0.726
Mean parish Gini, PLT side	0.702	0.680
PLT jump in parish Gini	0.100* (0.050)	0.047 (0.040)
Parishes in RD window	26	42
PLT parishes	14	17
Expropriation parishes	12	25

The first four rows describe inequality in plot-level agricultural area. Pooled plot Gini uses all plots on each side of the boundary; mean parish Gini averages parish-specific Gini coefficients. The PLT-jump row estimates an unadjusted local-linear specification in the RD framework of [Calonico et al. \(2014\)](#), with separate slopes on each side and Huber-White standard errors.

Both columns exclude the two parishes closest to the boundary. The estimates offer descriptive evidence on landholding structure rather than a primary test of the paper's mechanism. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.6. Broader regional benchmark

The Manabi benchmark provides supportive regional evidence but does not enter the core RD identification.

C.6.1. Manabi benchmark

The broader regional comparison tests whether the crop ranking extends beyond the RD corridor. I use Manabi plots outside the digitized PLT area as the benchmark. The PLT category combines mapped PLT areas in Guayas, Los Rios, Manabi, Esmeraldas, Pichincha, Santo Domingo de los Tsachilas, and Zona No Delimitada. The expropriation category contains the corresponding areas in Guayas and Los Rios. I report unconditional differences with parish-clustered standard errors. Province fixed effects would absorb the Manabi benchmark. Latitude and longitude controls would require extrapolation across this broad comparison. This exercise remains descriptive and outside the core identification strategy.

Table [A17](#) preserves the relative ranking across reform regimes. PLT areas exceed the non-PLT Manabi benchmark in total perennial revenue and revenue from sugar cane, oil palm, and plantain. Expropriation areas show no statistically distinguishable difference in total perennial revenue and fall below the benchmark for this diversification basket. Expropriation also shows the much larger annual and rice differences. Cacao yield does not differ detectably from the benchmark in either regime. I treat these estimates as a descriptive benchmark rather than a causal comparison.

Table A17: Manabi regional benchmark

	Perennial	Annual	Rice	Cacao	Banana	Main other perennials
PLT	1.461*** (0.232)	0.162 (0.275)	0.554*** (0.180)	0.012 (0.013)	0.589** (0.233)	1.543*** (0.248)
Expropriation	-0.293 (0.298)	1.377*** (0.144)	2.462*** (0.181)	-0.000 (0.013)	1.769*** (0.420)	-0.376** (0.187)
Observations	23,804	19,165	19,165	23,804	23,804	23,804
Parishes	194	194	194	194	194	194

Columns report log revenue per hectare for perennial crops, log revenue per hectare for annual crops, rice yield, cacao yield, banana yield, and log revenue per hectare from sugar cane, oil palm, and plantain. Each outcome uses its corresponding crop questionnaire, and zeros are assigned only to plots observed in that questionnaire. Perennial-questionnaire membership is defined before the crop-age restriction, and cacao and banana yields divide plot-level total production by total usable crop area.

The omitted category contains Manabí plots outside the digitized PLT area. The PLT category contains mapped PLT areas in Guayas, Los Ríos, Manabí, Esmeraldas, Pichincha, Santo Domingo de los Tsáchilas, and Zona No Delimitada. The expropriation category contains the corresponding areas in Guayas and Los Ríos. The unconditional ordinary-least-squares regressions add no covariates. Province fixed effects would absorb the Manabí benchmark, while latitude and longitude controls would require extrapolation across this broad regional comparison. Standard errors are clustered by parish and shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.7. Productivity-reallocation calibration

This back-of-the-envelope adaptation of [Kim and Wang \(2025\)](#) asks a narrow question. Can a negative aggregate agricultural-productivity estimate after agrarian reform coexist with positive diversification and labor-reallocation terms? The answer is yes in the accounting sense. Negative results in established crops can dominate the agricultural average, while diversification and movement out of agriculture contribute positively. The calculation makes those offsets visible. It uses the reallocation formula in [McCaig and Pavcnik \(2018\)](#). It remains suggestive, not a structural estimate or a replacement for the RD.

C.7.1. Model and equations

Consider a two-sector economy with agriculture (a) and non-agriculture (n). Each sector uses a Cobb–Douglas technology,

$$Y_s = A_s L_s^{\alpha_s} K_s^{1-\alpha_s}, \quad s \in \{a, n\}.$$

Aggregate labor productivity is

$$\frac{Y}{L} = \frac{Y_a + Y_n}{L_a + L_n}.$$

Define sectoral output per worker as

$$ARPL_a = \frac{Y_a}{L_a}, \quad ARPL_n = \frac{Y_n}{L_n},$$

and let the agricultural employment share be

$$s_a = \frac{L_a}{L_a + L_n}.$$

Aggregate labor productivity then equals

$$\frac{Y}{L} = s_a ARPL_a + (1 - s_a) ARPL_n.$$

Following [Kim and Wang \(2025\)](#), if b_a is a proportional change in agricultural labor productivity, its direct contribution to aggregate labor productivity is

$$\text{Productivity contribution} = s_a \times b_a.$$

The aggregate residual-productivity coefficient provides the closest available counterpart to b_a , but it remains a proxy based on $\log(1 + \text{revenue})$. Crop-specific coefficients measure zero-inclusive transformed revenue per hectare or crop yields, not proportional changes in agricultural labor productivity. I therefore label $s_a b_a$ a scaled agricultural term and treat the crop rows as specialization scenarios.

Following [McCaig and Pavcnik \(2018\)](#), I hold sectoral output per worker fixed and change only agriculture's employment share. Labor reallocation then contributes:

$$\Delta \left(\frac{Y}{L} \right)^{realloc} = -\Delta s_a (ARPL_n - ARPL_a).$$

I divide by baseline aggregate labor productivity and define

$$g = \frac{ARPL_n}{ARPL_a},$$

which gives

$$\text{Reallocation contribution} = -\Delta s_a \cdot \frac{g - 1}{s_a + (1 - s_a)g}.$$

C.7.2. Calibration

Ecuador's 2000 GGDC/UNU-WIDER Economic Transformation Database (ETD), documented by Kruse et al. (2023), provides the aggregate inputs. Agriculture's employment share is

$$s_a = \frac{1450.226}{1450.226 + 3142.816} = 0.316,$$

and the productivity gap is

$$g = \frac{ARPL_n}{ARPL_a} = 2.458.$$

The paper's supporting structural-transformation input is the PLT-versus-expropriation difference in the agricultural employment share,

$$\Delta s_a^{PLT-EXP} = -0.136,$$

which implies a common reallocation contribution of about

$$0.099.$$

The calculations below use the unrounded inputs.

The supporting pooled canton-level exercise within 50 km of the border provides the employment-share estimate; the core plot-level RD does not. The employment share and productivity gap come from national ETD data rather than local coastal aggregates. The calculation therefore provides an interpretive benchmark, not a separate causal estimate.

I apply the scaling term to aggregate residual productivity, perennial and annual revenue per hectare, rice, cacao, and banana yields, and revenue per hectare from sugar cane, oil palm, and plantain. I calculate the residual-productivity proxy following Montero (2022) and Aragon et al. (2020). Because the crop coefficients measure zero-inclusive transformed revenue per hectare or yields rather than proportional changes in agricultural labor productivity, those rows remain illustrative specialization scenarios.

C.7.3. Results

The aggregate residual-productivity term contributes -0.181 , while labor reallocation contributes 0.099 , leaving an index of -0.082 . Movement out of agriculture offsets more than half, but not all, of the weak within-agriculture result. The specialization scenarios reveal the oppos-

ing margins: sugar cane, oil palm, and plantain produce a strongly positive index, while rice and bananas produce negative indices. Established-crop persistence and diversification therefore pull the aggregate in opposite directions.

Table A18: Why aggregate productivity looks weak: suggestive accounting

Margin	Scaled agricultural term	Reallocation term	Combined index
<i>Panel A. Aggregate benchmark</i>			
Aggregate residual-productivity proxy	−0.181	0.099	−0.082
<i>Panel B. Illustrative specialization scenarios</i>			
Perennial	−0.054	0.099	0.045
Annual	−0.165	0.099	−0.066
Rice	−0.209	0.099	−0.110
Cacao	−0.015	0.099	0.084
Banana	−0.519	0.099	−0.420
Sugar cane, oil palm, and plantain	0.391	0.099	0.490

Note: Panel A combines the aggregate residual-productivity proxy with the labor-reallocation term. The proxy is the robust bias-corrected PLT-side discontinuity following [Calonico et al. \(2014\)](#) in a residual from $\log(1 + \text{revenue})$ regressed on transformed land, labor, temperature, and precipitation. It uses the 40 km bandwidth, excludes the two parishes closest to the boundary, controls for latitude, longitude, border-segment fixed effects, province fixed effects, land-tenure fixed effects, and banana potential yield, and clusters standard errors by parish. The scaled agricultural term multiplies each inserted coefficient by agriculture’s national employment share. The combined index adds the common labor-reallocation term. Panel B mechanically inserts the crop-specific RD coefficients to illustrate the offsetting specialization margins. The Economic Transformation Database supplies Ecuador’s 2000 agricultural employment share and sectoral productivity gap; the authorized canton-year IPUMS aggregate supplies the local employment-share difference. Crop outcomes use their corresponding questionnaires and zero-inclusive transformations. The crop coefficients, scaled terms, and combined indices are not percentage effects on aggregate labor productivity.

C.7.4. Break-even perennial exercise

The break-even exercise finds the inserted perennial coefficient that makes the accounting index positive. It does not identify any crop as the cause of the aggregate perennial coefficient. Using the unrounded employment share and reallocation benchmark, define the index for an inserted coefficient b^{cf} as

$$I(b^{cf}) = 0.315744 \cdot b^{cf} + 0.099023.$$

The index becomes positive whenever

$$I(b^{cf}) = 0.315744 \cdot b^{cf} + 0.099023 > 0,$$

which implies

$$b^{cf} > -0.313617 \approx -0.314.$$

Because the observed perennial margin lies near zero, a modest change in the inserted coefficient makes the index positive. Table A19 reports the arithmetic. The threshold applies to the accounting index; it does not estimate a counterfactual change in perennial or aggregate labor productivity.

Table A19: Break-even perennial accounting scenarios

Scenario	Coefficient inserted	Scaled term	Combined index
Observed perennial coefficient	−0.172	−0.054	0.045
No perennial gap	0.000	0.000	0.099
Small positive insertion	0.100	0.032	0.131
Moderate positive insertion	0.250	0.079	0.178
Large positive insertion	0.500	0.158	0.257
Illustrative diversification-basket insertion	1.239	0.391	0.490

Note: Each row inserts a perennial coefficient b^{cf} into $I(b^{cf}) = 0.315744b^{cf} + 0.099023$. The first number is Ecuador’s unrounded agricultural employment share, and the second is the unrounded labor-reallocation contribution. The diversification basket contains sugar cane, oil palm, and plantain. The combined index becomes positive when $b^{cf} > -0.313617$. These rows are mechanical scenarios, not estimates of counterfactual changes in perennial or aggregate labor productivity. The exercise does not identify any crop as the cause of the aggregate perennial coefficient.

The break-even calculation compares the aggregate perennial coefficient with the reallocation benchmark without assigning the weak aggregate result to one crop. The broader evidence points instead to offsetting patterns: incumbent activities performed better outside the PLT footprint, while the PLT-side environment supported a wider perennial portfolio and labor reallocation.