# The Price of Corruption: Evidence from Lava Jato and the Construction Club in Peru

Carlos Cesar Chavez Padilla

June 2026

**Abstract**

How much does corruption cost in public infrastructure, and who bears it? We study two judicial records of Peruvian road corruption, Odebrecht bribery in road concessions and a sanctioned cartel that rigged national road tenders between 2002 and 2016, and reconstruct from them the universe of tenders and firm roles. Because the competition authority adjudicated each tender, we validate collusion-detection screens against ground truth, and a level screen on the winning bid recovers the verdict where the dispersion screen the literature relies on does not, with areas under the curve of 0.87 against 0.50. Bribed concession projects overran initial budgets by more than half against under a fifth for the rest, and colluded winners bid near the reference-value ceiling while competitive winners gave the State a ten-percent discount. We find no measurable local welfare dividend and no crowd-out of other public investment. The price is a national-budget cost, and level-based screens sharpen its detection.

## 1 Introduction

Corruption in infrastructure is hard to study because the corrupt projects are rarely identified. Peru offers two judicial windows that identify them. The first is the Odebrecht (Lava Jato) record, whose collaboration agreements and prosecutions name the road concessions, the bribes, and the officials who took them. The second is Resolución 080-2021/CLC of the Peruvian competition authority, which sanctions thirty-two construction groups for a decade of bid rigging in the public procurement of national road works. The two records describe distinct mechanisms in the same ecosystem. Lava Jato is bribery in concessions awarded at the top. The Club is cartelized allocation of the tenders below. Together they let us study road corruption with the corruption documented in the rulings, and ask a cost-effectiveness question: did the projects that corruption touched cost more, allocate less competitively, and still deliver commensurate local benefits.

We use both records and reach five results. The first is methodological. Because the resolution adjudicated each tender as colluded or competitive, we validate the collusion-detection screens against that ground truth, and a level screen on the winning bid separates colluded from competitive tenders far better than the dispersion screen the literature relies on, with areas under the curve of 0.87 against 0.50. Second, on fiscal cost, the Odebrecht projects with a documented bribe ran cost overruns averaging more than half of initial investment, against under a fifth for the projects with no documented evidence, a gap robust to the small-sample inference twenty-five projects allow. Third, on allocation, we reconstruct the cartel from the antitrust resolution into a universe of seventy-eight road tenders and the firms that won or covered each one. Wins are diffuse and shared by turns, with no single firm dominant, and the cartel's gain shows in the price, the winning bid on a colluded tender sitting near the reference-value ceiling against a ten-point discount on the competitively bid ones. Fourth, we map the cartel to districts and take it to household welfare. The concession corridors show no detectable household dividend. Districts near cartelized procurement do show adverse outcomes, but so do districts near competitively procured road tenders, so the broad association is a feature of large road projects. Using the cartel's own classification of which tenders it rigged and which it bid competitively, an adverse association specific to the colluded tenders survives a control for road size and is significant within departments, but it loses statistical precision under the tightest within-province comparison, which only eight provinces support, so we read it as suggestive, and not separately identified. Fifth, we are explicit about what the data cannot support. We reconstruct contract allocation and capture, not firm profits, and the firm-level outcome designs used elsewhere fail here because the procurement records omit the large works and the consortia that won them.

This places the price of corruption on the procurement side. The fiscal overruns and the cartelized allocation are corruption-specific by construction, documented in the rulings and absent from a competitive counterfactual. The household side is a roads-and-development result. The roads delivered no compensating local welfare regardless of how they were procured. An adverse association does attach specifically to the colluded tenders and survives holding road size fixed, but it is significant only at the department level and underpowered under the within-province comparison that would separate the corruption from the local geography, so we present it as suggestive and place the cleanly identified price of corruption on the procurement side.

## 2 Institutional background

Peru's road corruption is documented by two separate judicial processes that reach different layers of the same system. Figure 1 places them on a common timeline.

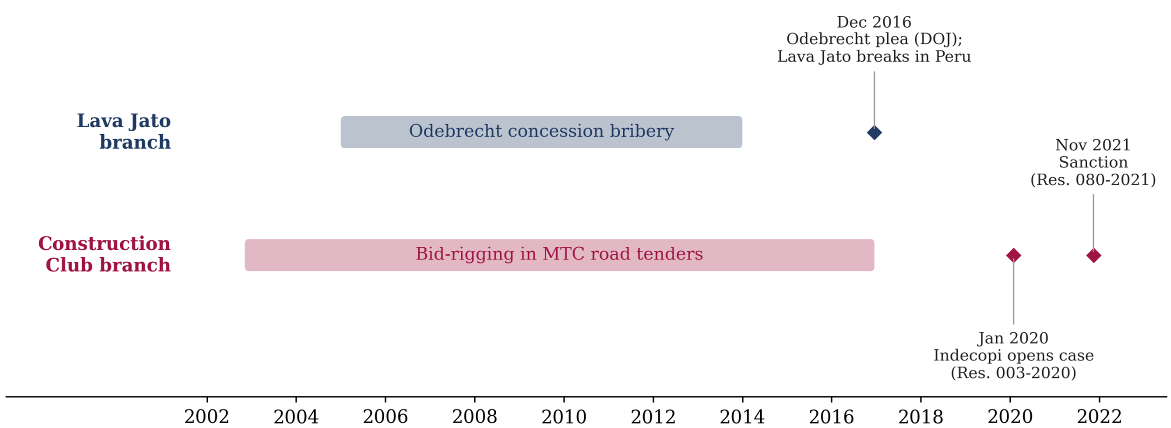


Figure 1: The two judicial branches of Peruvian road corruption.

*Notes.* Conduct periods (bars) and dated events (markers). Construction Club dates are from Resolución 080-2021/CLC (cartel conduct November 2002 to December 2016, opening of the procedure January 2020, sanction November 2021). Lava Jato dates are from the public record (Odebrecht bribery in Peru 2005 to 2014, the December 2016 plea agreement that broke the case).

The first window is the Odebrecht (Lava Jato) case. Odebrecht admitted to paying bribes to obtain public works in Peru over roughly a decade, and the December 2016 plea agreement with the United States Department of Justice, followed by the collaboration agreements with Peruvian prosecutors, named the projects and the officials. The corruption here operated mainly through concessions and megaprojects awarded at the top, the IIRSA corridors and urban works among them, and through the payments that secured them. The contracts were awarded mostly as concessions through the national investment-promotion agency, not through open procurement, and held by special-purpose vehicles financed by securitized debt and state co-financing. Odebrecht admitted roughly US$29 million in bribes in Peru, and four former officials, among them a former president, a former vice-minister of transport, and two regional governors, concentrate the documented payments. The revelation was an economic shock as well as a political one, with public-private infrastructure investment falling from around US$9 billion to under US$1 billion in 2016 and 2017 as projects froze. This is the branch an earlier version of this paper studied, and it identifies a set of road concessions as corrupt by judicial finding, not by inference.

The second window is the *Club de la Construcción*. In Resolución 080-2021/CLC of 15 November 2021 the Peruvian competition authority Indecopi sanctioned thirty-two construction groups for a horizontal bid-rigging agreement in the public procurement of national road works between November 2002 and December 2016. The mechanism was coordination among bidders. The firms agreed in advance on who would win each *licitación pública* run by the Ministry of Transport and Communications and arranged cover bids and abstentions to protect the designated winner. The procedure opened with the Resolución de Inicio 003-2020/ST-CLC of 31 January 2020. This branch reaches the procurement layer the concession cases do not, the ordinary road tenders, and it names the firms and their roles tender by tender.

The two records are complementary, not redundant. Lava Jato identifies the high-level bribery and concession channel but covers a small number of megaprojects. The Construction Club identifies the procurement-cartel channel across many ordinary road tenders and a wide geography. Studied together they document the same road-corruption ecosystem from the political allocation of concessions down to the firm-level rigging of tenders, and they let us measure the procurement-side cost and distortion of corruption against the local outcomes in the places the roads were meant to serve. We keep the two as distinct records throughout and never merge them into a single database.

## 3 Literature review

The paper draws on six literatures, and our design follows from what each needs and what the Peruvian setting can supply. The first is the measurement of corruption, which is hard to observe and usually proxied by perceptions. Olken and Pande (2012) survey the move toward direct measurement, and Olken (2007) is its canonical road-sector instance, comparing official with engineer-estimated costs in Indonesian village projects. We measure in this direct tradition but from a different source, the judicial record. Campos et al. (2021) are the closest antecedent, reading cost overruns of bribed infrastructure off the Odebrecht case across nine countries, and we replicate their comparison on the Peruvian subset and read it as a magnitude. The line between a larger cost and a captured rent is the active-versus-passive waste distinction of Bandiera et al. (2009), and we describe the overruns as excess cost without attributing them to any party.

A second literature studies procurement award and performance. Decarolis (2014) shows the award price is not a sufficient statistic for value once performance and renegotiation are accounted for, which is why we treat the cartel through allocation and documented conduct instead of a single award figure. Coviello et al. (2018) show that discretion and repeated winners shape procurement outcomes, the pattern our reconstruction recovers, and Mironov and Zhuravskaya (2016) detect procurement corruption from firm financial transactions, a design whose data the Peruvian records cannot supply. Bosio et al. (2022) document across 187 countries that restricting procurement discretion helps low-capacity states and hurts high-capacity ones, and Peru is the low-capacity setting in which the discretion the cartel exploited is the object of interest.

A third literature, the one closest to our reconstruction, detects collusion in procurement auctions. Porter and Zona (1993) identify bid rigging in Long Island highway construction from the rank distribution of the bids, the canonical case of a ring that designates a winner while the others enter phony higher bids, which is the structure the cartel's exhibits record here. The modern wave screens for that conduct statistically. Conley and Decarolis (2016) recover bidder groups in collusive auctions, Kawai and Nakabayashi (2022) detect large-scale collusion among Japanese construction firms from rebidding, and Chassang et al. (2022) and Kawai et al. (2023) build robust screens and a regression-discontinuity test for noncompetitive bidding, with Pesendorfer (2000), Bajari and Ye (2003), and Marshall and Marx (2012) the foundations. Our setting inverts their inference problem. Where these designs must detect the designated-winner-and-cover-bid structure from bid data, the competition authority established it tender by tender, so we observe the assignment the screens have to infer. The closest structural match to our two-branch design is Clark et al. (2018), who study bid rigging and entry deterrence in Quebec construction through an investigation into collusion and corruption. We differ in reaching two layers of the same ecosystem through two separate judicial records, concession bribery above and tender rigging below, and in carrying both to household welfare, where procurement studies usually stop.

A fourth literature asks what transport infrastructure delivers locally. Donaldson and Hornbeck (2016) formalize market access as a sufficient statistic for the local gains, the construction our exposure measure uses, while the evidence on incidence is mixed: Faber (2014) finds China's highways concentrated activity in central cities, Ghani et al. (2016) find gains near India's Golden Quadrilateral, Morten and Oliveira (2024) trace roads through trade and migration to welfare, and Asher and Novosad (2020) find new rural roads in India raised neither income nor assets. Our no-dividend result for the concession corridors is consistent with that evidence, and Burgess et al. (2015), who show road building in Kenya followed political favoritism, is the premise our two records make concrete.

A fifth literature studies disclosure. Ferraz and Finan (2008) show that releasing audit results changed electoral outcomes in Brazil, evidence that revelation is itself an economic event, which frames the role of the 2016 Odebrecht disclosure and the 2020 to 2021 Indecopi procedure that we use to date exposure. A sixth literature values politically connected or scandal-exposed firms. Fisman (2001) prices connections from abnormal returns around shocks to a patron, and Colonnelli and Prem (2022) show exposed Brazilian firms reallocate toward productive activity using firm microdata. Both need data Peru does not provide for these firms, whose large works were won by project-specific consortia outside the open procurement records, so we reconstruct contract allocation and capture, not firm profits or returns.

Against these literatures the paper makes one methodological contribution and two substantive ones, all resting on a data contribution. The headline is that we validate collusion-detection screens against an adjudicated ground truth. Because Indecopi classified each tender as colluded or competitive, we can run the screens this literature has to defend by inference and score them against the truth, and we find that a level screen on the winning bid recovers the verdict far better than the dispersion screen the literature most often uses (Section 7). This is possible only because of the data contribution, two independent judicial records of the same road-corruption ecosystem brought into a single analysis and a universe of road tenders and firm roles reconstructed from the antitrust resolution and not previously assembled, the realized assignment the detection literature must infer. On that foundation the paper measures the procurement-side price of the corruption, the larger fiscal overruns on the bribed projects and the concentrated cover-bid allocation and overcharge on the cartelized tenders, without overclaiming a rent it cannot observe. And it carries the corruption to households, showing that the concession corridors delivered no detectable local dividend and that the adverse association with cartelized procurement, while present, cannot be separated from the road and is reported as bounds. The reading is one of cost-effectiveness, documented road corruption raised cost and distorted allocation without a commensurate local development return.

# 4 Data

*The judicial record of bribes and overruns.* Treatment in Module 1 is a bribe documented in a court ruling, from the Odebrecht case dataset of Campos et al. (2021), restricted to the twenty-five Odebrecht projects in Peru. The dataset codes, per project, whether a bribe was paid, the bribe amount, and the cost overrun as a fraction of initial investment, with sector and initial and final investment. Ten projects carry a confirmed documented bribe, five are judicially implicated without a project-specific payment, and ten have no documented evidence. The corridors these projects built are mapped in Figure 2.


Figure 2: Odebrecht road concession corridors by documented-bribe status.
*Notes.* Documented-bribe corridors (warm: IIRSA Norte, IIRSA Sur T2 and T3) and the rest (cool). *Source.* MTC national road network (2017), bribe coding from Campos et al. (2021).

*Road network and the accessibility measure.* Module 2 is built on the MTC national road network for 2017 (*Red Vial Nacional*), which records each segment's surface. We node the network at intersections so that crossing roads connect at junctions, giving a graph of about 34,700 nodes whose giant component holds 99% of them. The network is sparse, as a national trunk network with few alternate routes in the Andes and Amazon is, and we report its connectivity. Travel time on each edge is its length divided by a surface-specific free-flow speed. District origins are the inhabited cells of the WorldPop 2017 population raster, covering the dispersed rural population and not only the district center, each carrying an explicit off-network feeder travel time to the nearest road node so that a remote settlement does not inherit a corridor node's access. Destinations are the 196 provincial economic centers, weighted by province population. The corridor interventions are coded by physical type, with the documented unpaved IIRSA routes paved (the conservative shock), the coastal Panamericana duplications and the Sierra rehabilitation treated separately, and corridor-specific pre-upgrade speeds, all documented with their basis and uncertainty. Table 1 collects the parameters.

Table 1: Market-access measure: parameters and corridor coding.

| Element | Specification |
|---|---|
| *Free-flow speed by surface (km/h)* | |
| Rigid pavement / asphalt | 80 / 70 |
| Basic pavement / cobble | 50 / 45 |
| Gravel (afirmado / sin afirmar) | 35 / 25 |
| Earth track (trocha / emboquillado) | 20–22 |
| Off-network feeder | 15 |
| *Gravity and geography* | |
| Decay parameter $\theta$ | 1 and 2 |
| Origins | WorldPop 2017 inhabited cells, population-weighted |
| Destinations | 196 provincial economic centers, population-weighted |
| *Corridor upgrade coding (pre-upgrade speed)* | |
| Conservative (IIRSA paving) | documented gravel routes at 22–30, paved post |
| Coastal Panamericana duplication | treated separately |
| Sierra rehabilitation | treated separately |

*Notes.* Travel time on each network edge is its length divided by the surface-specific free-flow speed, with an explicit off-network feeder time from each populated origin cell to the nearest road node. The accessibility gain of a corridor differences market access before and after the documented upgrade, under the conservative, midpoint, and expansive shock definitions. The Huánuco–Lima PE-3N defect is disclosed in Section 4. *Source.* MTC *Red Vial Nacional* 2017 and WorldPop 2017.

*Route validation and a known defect.* We validate the network against twenty known origin-destination routes across the coast, Andes, and Amazon. Eighteen are plausible against known travel times, one correctly remains unreachable by road (Iquitos), and one central-highway route, Huánuco to Lima, fails because of a break in the PE-3N corridor that forces an implausible detour. The defect is in the source network and is not yet repaired. Until it is, the travel-time-to-Lima measure is reported as secondary and the affected central-highland districts are flagged, while the market-access gradient, which differences the same network before and after the cartel upgrades, is largely unaffected.

*Household welfare.* Module 3 uses the national household survey (ENAHO), *sumaria* module, 2004–

2023, aggregated to a district–year panel of real per-capita household income and consumption and the poverty rate, deflated and person-weighted. ENAHO is designed to be representative above the district, so the district cells carry survey sampling noise, which the district and region-by-year fixed effects and the clustered inference are meant to accommodate. A district long-difference across the 2007 and 2017 population censuses provides a two-period check on the construction-employment question.

*Concession records.* For the institutional note on concession continuity we use the OSITRAN performance reports, which show whether each road concession was rescinded or interrupted after 2016. We use these to establish that the contracts continued operating, not to infer the persistence of any rent.

*Public investment.* The fiscal-incidence test uses the MEF Invierte.pe project registry, public investment by district and function with initiation and completion dates, to test whether the cartel overcharge crowded out non-transport investment.

*Asset declarations.* The appendix exploratory exercise uses the JNE candidate sworn asset declarations, linked by national identity number across the 2018 and 2022 municipal elections.

### 4.1 Descriptive statistics

Table 2 summarizes the household panel. The estimation sample is 21,938 district-year cells covering 1,692 districts over 2004 to 2023. Mean log per-capita income is 8.40 and mean log consumption 8.25, each with the dispersion across districts that a national panel spanning Lima and the rural Andes and Amazon would show, and the mean district poverty headcount is 0.42. Each cell pools a median of fourteen surveyed households, the sampling noise the district and region-by-year fixed effects and the clustered inference absorb.

Table 2: Summary statistics for the household panel.

| | Mean | SD | Median | N |
|---|---|---|---|---|
| *Panel A. District-year outcomes (2004–2023)* | | | | |
| Log household income | 8.40 | 0.71 | 8.45 | 21,938 |
| Log household consumption | 8.25 | 0.61 | 8.29 | 21,938 |
| Poverty rate | 0.42 | 0.30 | 0.38 | 21,938 |
| Households surveyed per cell | 25.9 | 41.0 | 14.0 | 21,938 |
| *Panel B. Districts by exposure cell* | | | | |
| Neither corridor nor cartel | | | | 1,127 |
| Lava Jato corridor only | | | | 165 |
| Construction Club cartel only | | | | 332 |
| Both branches | | | | 68 |
| All districts | | | | 1,692 |

*Notes.* The estimation panel is the national household survey (ENAHO) collapsed to district-year cells, 21,938 cells across 1,692 districts for 2004 to 2023. Income and consumption are per capita and in logs, the poverty rate is the headcount, and the last row counts surveyed households underlying each cell. Panel B classifies a district by whether it is ever within ten kilometers of a Lava Jato concession corridor, of a Construction Club cartel tender, of both, or of neither. The cartel-only and both cells, 400 districts, are the ones the household analysis identifies from. *Source.* ENAHO *sumaria*, the Odebrecht concession corridors, and Resolución 080-2021/CLC.

Panel B places districts in the four exposure cells that organize the analysis. Of the 1,692 districts, 1,127 are near neither branch of the corruption, 165 lie near a Lava Jato concession corridor only, 332 near a Construction Club cartel tender only, and 68 near both. The household identification comes from the 400 districts ever exposed to the cartel procurement, and the colluded-versus-competitive contrast of Section 6.2 narrows further to the eight provinces that contain both a colluded and a competitive tender. The thinness of that overlap is the binding constraint on the household side, and it is the reason the local results are read as bounds.

### 4.2 The reconstructed cartel universe

The corruption in Peruvian road infrastructure has two judicially documented arms, and this paper uses both. The first is the Odebrecht (Lava Jato) record of bribery behind a set of road concessions and megaprojects, which Section 2 describes. The second, which we introduce here, is the *Club de la Construcción*, a cartel that rigged the public tenders for national road works. The two are distinct mechanisms in the same ecosystem. Odebrecht bribery operated mainly through concessions awarded by ProInversión and the officials who took the payments. The Club operated through bid rigging in the *licitaciones públicas* run by the Ministry of Transport and Communications. We treat them as two records, not one database.

#### 4.2.1 The antitrust resolution

Indecopi, the Peruvian competition authority, sanctioned the cartel in Resolución 80-2021/CLC-INDECOPI of 15 November 2021 (Expediente 001-2020/CLC). The procedure had opened with the Resolución de Inicio 003-2020/ST-CLC of 31 January 2020, and the imputed firms were notified in early February 2020. The resolution declares that thirty-two economic groups engaged in a horizontal bid-rigging agreement, coordinating cover bids and abstentions in the public procurement of road works at the national level between November 2002 and December 2016, and it fines them in units of the *Unidad Impositiva Tributaria* (UIT). The conduct periods are firm specific. Sixteen groups participate from the cartel's start in November 2002, and sixteen enter later, between 2004 and 2014. The resolution is the authoritative source for who was in the cartel, for each firm's role in each tender, and for the award outcomes. We take winners, members, and roles from it throughout.

#### 4.2.2 Reconstructing the process universe

The resolution documents the cartel's allocation through firm-specific evidence exhibits, in which each firm's own records list the tenders it participated in and the consortium that took each award. Appendix A documents the full reconstruction and validation of both judicial datasets, the Construction Club cartel here and the Lava Jato concessions of Section 6.1. We read three structured exhibits directly from the resolution at the page level: the Cosapi exhibit (CSP18, twenty-five processes), the Obrainsa exhibit (OBN10, thirty-eight processes), and the JJC exhibit (JJC01, fifty-two processes). The remaining named evidence (Johesa, GyM, San Martín and others) is narrative and not a process table, and we use it only to corroborate participation, not to add processes.

Merging the three exhibits requires care because the same tender appears under different code notations across firms, for example a *clásico* number in one exhibit and an MTC code in another. We match on a conservative full code signature and confirm each candidate on project name and winner before merging, so that distinct sub-tramos and distinct lots are never collapsed by code similarity alone. The union is *seventy-eight distinct processes* spanning 2002 to 2016. For each process we record the designated winning consortium, its member firms, and the value figures the resolution reports. For each firm in each process we record a role: winning-consortium member, cover bidder, cover bidder of uncertain target (the resolution's category for support whose beneficiary the firm did not recall), or participant without support.

#### 4.2.3 Cartel governance, deviations, and enforcement

The resolution records not only who won and who covered but how the cartel governed itself, and this matters for the mechanism. A mechanical rotation table would leave little room for a sustained overcharge, but the evidence describes a durable institution that managed entry, assignment, and discipline over fourteen years. Table 3 collects the documented episodes. The assignment of each process was preceded by coordination meetings, general and one-on-one, in which firms declared their interest in particular works, and a standing coordinator organized the smaller firms. The arrangement survived friction. When one of the largest members withdrew in 2007, the remaining firms continued to coordinate and then arranged its reincorporation, which the competition authority reads as the persistence of the agreement and not its collapse. Discipline was explicit: an internal email records a signed side-agreement and the understanding that a member who broke ranks would cease to belong. Deviations occur in the record, eight tenders in which a cartel firm bid without supporting the designated winner, and one award that passed on appeal from the disqualified designate to another cartel member, but they are deviations within a functioning system, not its breakdown.

Table 3: Cartel governance, deviations, and enforcement.

| Episode type | Period | What the resolution documents | Evidence |
|---|---|---|---|
| Coordination role | 2002–2016 | Grupo Plaza acted as the standing coordinator of the small and medium firms, named in the declarations of Cosapi and JJC executives. | Aranda, Martínez (p. 362) |
| Pre-assignment meetings | 2002–2016 | The assignment of each process was preceded by meetings in which members declared their interest in a given work, held in named restaurants and offices, both general and one-on-one. | Testimony (p. 158, 283) |
| Temporary exit and re-entry | 2007–2008 | JJC withdrew in 2007 while the remaining members kept coordinating, and other firms then coordinated JJC's reincorporation into the agreement. | GyM18, GyM20 (p. 362) |
| Signed side-agreement and enforcement | 2011 | An internal email records a signed agreement with an outside representative and the threat that a member who broke the agreement would cease to be part of the arrangement. | CSP13 (p. 1387) |
| Documented deviations | 2002–2016 | Eight tenders in which a cartel firm participated without supporting the designated winner, the resolution's no-support category. | JJC01 (p. 312–313) |
| Reassignment on appeal | 2010 | The designated winner of one tender was disqualified and the award passed on appeal to another cartel member, the coordination rerouting rather than failing. | CSP18 (p. 315) |

*Notes.* Episodes coded from the narrative and the evidence exhibits of Resolución 080-2021/CLC, with the resolution's own evidence codes and page references. The episodes describe how the cartel was governed rather than how often it failed, and the competition authority finds that the coordination continued across the deviations and the temporary exit. *Source.* Resolución 080-2021/CLC.

The allocation the governance produced looks like division by turns more than dominance. Documented wins are diffuse, with a winner Herfindahl of 0.03 and no group taking more than nine percent, consistent with a sharing arrangement in which members take turns. Two regularities shape who takes which turn. Allocation tracks capacity, since participation and wins correlate at 0.72, and it tilts by territory, since groups with three or more wins concentrate 61 percent of them in a single geographic zone, the market-division pattern that Marshall and Marx (2012) describe. This governance is the channel that makes a local effect conceivable. A cartel that endures, meets, assigns, and enforces can hold the winning price near the ceiling the State will pay, year after year, on the works that build a district's roads. The price of that durability is the overcharge documented in Section 6.2, and the mechanism by which it could reach households is not the road, which a competitive contract would also have built, but the public money absorbed in the margin between the colluded price and the competitive one, money that did not fund the additional or better works a competitive program would have delivered for the same budget. Whether that channel is large enough to move household welfare is the question Section 6.2 cannot answer cleanly, but the governance evidence is why the question is worth asking.

#### 4.2.4 External validation against procurement records

We validate the universe against Peru's open procurement data (the OCDS export of the SEACE system). An early attempt matched poorly for 2015 and 2016 until we traced the cause to our own processing and not the data. For the post-2015 procurement regime the process code lives in the record's `officialNumber` and the work name in its `description`, and an earlier extract had kept only the title field. Re-extracting the raw OCDS with the full fields recovers the missing processes. Of the seventy-eight processes, *seventy-three (94 percent) are found* in the enriched procurement layer, and *sixty-three (81 percent) are high-confidence matches* on the licitación type, the MTC sub-unit, and the work name. The unmatched remainder is four pre-2004 processes that predate digital procurement coverage and a small set with a resolution-versus-SEACE code numbering mismatch. The procurement data confirms that these are real tenders of the right entity, year, work, and order of magnitude.

Two limits of the procurement data shape how we use it. First, *SEACE is not a usable source of winners* for the large road licitaciones. The award and supplier records are empty for the big works, so the winner-to-firm link is absent for exactly the processes that matter. Winner identity therefore comes from the resolution, and SEACE serves only to validate existence, code, entity, work, date, and the order of magnitude of the value. Second, for the same reason we *do not construct a cartelized-share-of-procurement measure.* A firm-level share would divide cartelized value by the firm's total public-contract value, but the denominator from SEACE awards is too incomplete and uneven for the cartel firms, fourteen of the thirty-two have no award records at all and the large-works winners appear in only a handful, so the share would be biased upward for the most relevant firms. We use a measure built entirely from the judicial record instead.

#### 4.2.5 The cartel-exposure measure

Our firm-level measure of cartel involvement is the set of processes, values, and roles the resolution attributes to each firm. For firm $f$ we record the number of cartelized processes it won, the number in which it acted as a cover bidder, the reference value of the processes it won where the resolution reports it, and the duration of its documented participation. *This measure is not a share of total procurement activity. It is a judicially documented exposure measure derived from the antitrust resolution.* It is predetermined with respect to the 2020-2021 sanction but it is not exogenous to firm size, capacity, or political connection, so we read it as exposure intensity, not randomly assigned treatment.

The reconstruction shows a concentrated cartel. Ingenieros Civiles y Contratistas Generales (ICCGSA) carries the largest fine at 77,265 UIT and the most won processes, followed by Obrainsa at 72,796 UIT, and the documented winning is concentrated in a small number of groups (ICCGSA, Johesa, Construcción y Administración, Energoprojekt, Obrascon Huarte Lain), while many firms appear mainly as cover bidders. Obrainsa, for instance, wins few of the processes in which it is documented and acts as a cover bidder in most. This pattern of designated winners surrounded by coordinated cover bids is the cartel mechanism, and it is the object the firm-side evidence reconstructs. We document the cartel's allocation of public road tenders from the judicial record and do not estimate firm outcomes.

The household analysis uses the district counterpart of this measure, the proximity-weighted exposure of each district to the geocoded tenders, whose construction is summarized in Table 4.

Table 4: Construction of the Construction Club cartel-exposure measure.

| Element | Specification |
|---|---|
| Spatial unit | District (centroid), national household survey panel |
| Cartel tenders geocoded | 76 of the 78 reconstructed road tenders |
| Proximity kernel | $\exp(-d_{ij}/\,10\text{ km})$ summed over nearby tenders |
| Distance bands (robustness) | 5 km, 10 km (baseline), 20 km |
| Alternative intensity | count of tenders within the band |
| Value weighting (robustness) | reference value of nearby tenders |
| Treatment timing | first nearby tender year, post from that year +2 |
| Districts ever exposed (10 km) | 400 (332 cartel-only, 68 both branches) |
| Standardization | per one standard deviation of exposure |

*Notes.* The exposure measure is built entirely from the judicial record, not from a procurement share, for the reasons given in Section 4.2. Two of the seventy-eight tenders could not be geocoded to a road segment and are dropped from the spatial measure. The decay, band, count, and value variants are reported in Appendix E. *Source.* Resolución 080-2021/CLC and the MTC national road network.

### 4.3 Mapping local exposure

To take the two records to households we measure, for each district, its exposure to each branch of the corruption, and the two measures cover different geographies (Figure 3).

For the concession branch we use the accessibility improvement each district received from the Lava Jato corridors, a network market-access calculation that compares travel times on the actual road network with a counterfactual in which the cartel segments carry their pre-upgrade condition (Section 5), together with log corridor kilometers and a travel-time measure. This exposure is concentrated on the eight built concession corridors of panel A.

For the cartel branch we build a district measure from the seventy-eight reconstructed tenders, geocoding each process from its endpoint towns to district centroids and reaching seventy-six of the seventy-eight across twenty-one departments (Appendix D). For each district we compute a proximity-weighted exposure and date its treatment by the earliest nearby process. The counts depend on the distance band: across all districts, 276 fall within five kilometers of a cartel process, 437 within ten, and 809 within twenty. Our analytic exposure uses the ten-kilometer band, under which 400 districts across twenty-three regions are exposed in the household panel. Panel B maps the wider twenty-kilometer footprint for visibility, and Appendix Figure 14 traces it growing from 258 districts by 2008 to 809 by 2016. Either way the exposure is far broader than the eight concession corridors and most of it lies away from the Lava Jato corridors, the independent geographic variation that lets the two branches tell distinct household stories in Section 6.

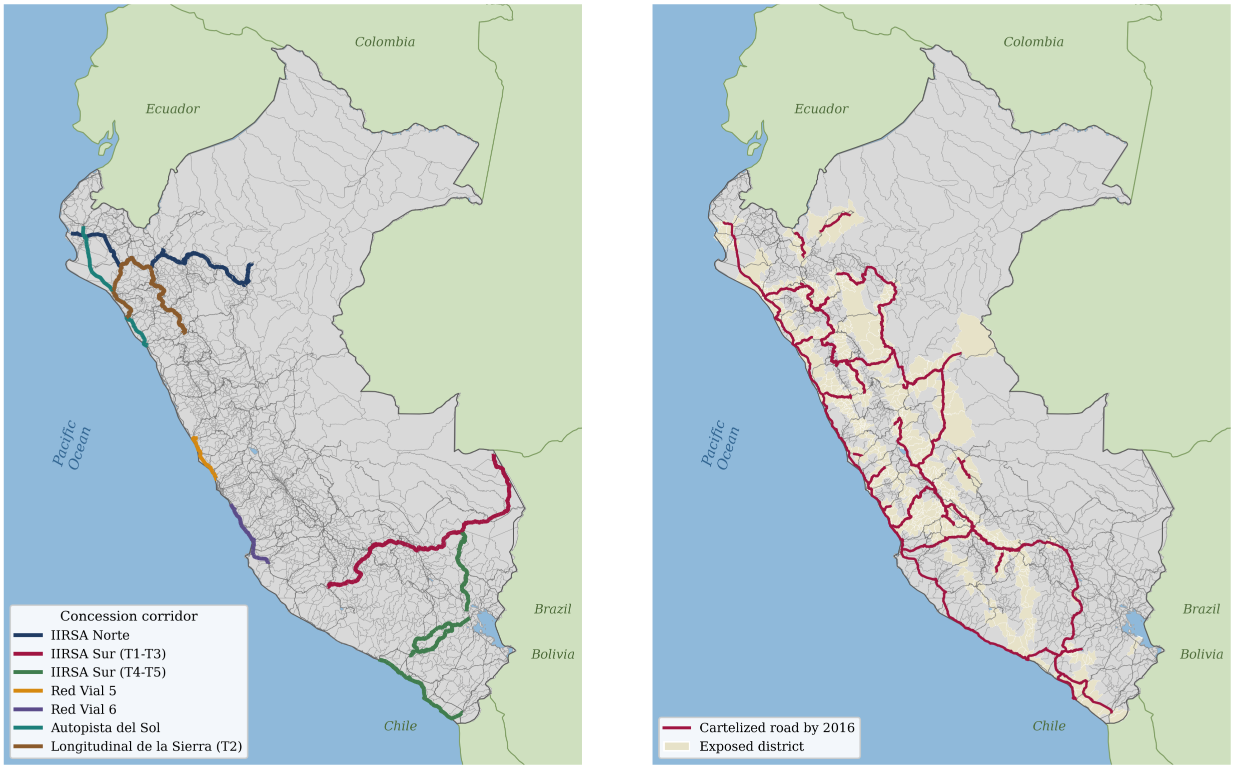


Figure 3: Geography of the two records. Panel A: Lava Jato concession corridors. Panel B: Construction Club cartel exposure by district.

*Notes.* Panel A shows the built concession corridors over the national road network. Panel B shades districts within twenty kilometers of a reconstructed cartel process, with the cartelized routes snapped to the actual road network overlaid and distinct from the national network in grey. Neighboring countries in green, the Pacific in blue. *Source.* Resolución 080-2021/CLC, INEI district boundaries, MTC road network, OSITRAN concession records.

# 5 Empirical strategy

The paper has five pieces. For the twenty-five Odebrecht projects we compare cost overruns by documented-bribe status,

$$\text{Overrun}_j = \alpha + \beta\,\text{Bribe}_j + X_j'\gamma + \varepsilon_j,$$

as means, medians, and a cross-section with controls, read as the magnitude of a fiscal difference with $N = 25$, not a causal elasticity, and assessed with exact and small-sample inference.

We measure the accessibility each district received from corridor completion as the change in market potential,

$$\Delta \ln MA_d = \ln MA_d^{\text{post}} - \ln MA_d^{\text{pre}}, \qquad MA_d = \sum_m \text{Mass}_m\, T_{dm}^{-\theta},$$

where $T_{dm}$ is travel time on the road network and the pre network sets the cartel segments to their pre-upgrade condition (Section 4). A placebo on random comparable corridors validates that the gradient is specific to the actual corridors, so this is a validation of the measure, not a first stage. We relate accessibility to household outcomes in the ENAHO panel,

$$y_{dt} = \beta\,(\Delta \ln MA_d \times \text{Post}_{dt}) + \alpha_d + \lambda_{r(d)t} + \varepsilon_{dt},$$

with district effects $\alpha_d$, region-by-year effects $\lambda_{r(d)t}$, Post keyed to each corridor's opening, and corridor-level standard errors. The fixed effects do not solve the endogenous placement of roads, and we treat the result as associational.

The cartel exposure enters the same panel. For district $d$ we build a proximity-weighted exposure $\text{Exp}_d^{\text{Club}}$ to the seventy-six geocoded cartel processes and a treatment date from the earliest nearby process, and estimate $y_{dt}$ on $\text{Exp}_d^{\text{Club}} \times \text{Post}_{dt}^{\text{Club}}$ with the same fixed effects, clustering by region. The reconciliation of the two branches is the central specification,

$$y_{dt} = \alpha_d + \lambda_{r(d)t} + \beta_1 \,(\text{Exp}_d^{\text{LJ}} \times \text{Post}_{dt}^{\text{LJ}}) + \beta_2 \,(\text{Exp}_d^{\text{Club}} \times \text{Post}_{dt}^{\text{Club}}) + \beta_3 \,(\text{Exp}_d^{\text{LJ}} \times \text{Post}_{dt}^{\text{LJ}})(\text{Exp}_d^{\text{Club}} \times \text{Post}_{dt}^{\text{Club}}) + \varepsilon_{dt},$$

where $\beta_1$ is the Lava Jato corridor effect, $\beta_2$ the Construction Club cartel effect, and $\beta_3$ their interaction, the outcome in districts exposed to both branches. Because the cartel onset is staggered across years, we estimate the dynamics with four estimators, two-way fixed effects, Callaway–Sant'Anna (Callaway and Sant'Anna, 2021), de Chaisemartin–D'Haultfœuille (de Chaisemartin and D'Haultfœuille, 2020), and Borusyak, Jaravel and Spiess (Borusyak et al., 2024), and rely on the two staggered-robust estimators for the identifying reading. The cartel's geography is not randomly assigned, so we read $\beta_2$ and $\beta_3$ as strong descriptive descriptive associations.

Each coefficient is a distinct estimand. The overrun $\beta$ is a conditional mean difference across twenty-five projects, not an average treatment effect. The accessibility and cartel coefficients are dose-response slopes, the change in a district's outcome per unit of accessibility gain and per standard deviation of cartel exposure, identified off the timing and geography of corridor openings and of nearby cartel tenders. Two treatment clocks run in the panel. The corridor branch separates a construction window, when works are underway, from the opening date when traffic begins, and keys $\text{Post}^{\text{LJ}}$ to the opening. The cartel branch keys $\text{Post}^{\text{Club}}$ to the first year a cartel tender appears near the district. A district can sit on both clocks, and that joint exposure is what $\beta_3$ reads.

The corruption-specific household question needs a comparison that holds the road fixed, and the cartel's own classification supplies one. Among road-exposed districts we separate exposure to the tenders Annex 2 calls colluded from exposure to those it calls competitive, bid by the same firms on the same class of works,

$$y_{dt} = \alpha_d + \lambda_{g(d)t} + \delta_C \,(\text{Exp}_d^{\text{Col}} \times \text{Post}_{dt}) + \delta_K \,(\text{Exp}_d^{\text{Comp}} \times \text{Post}_{dt}) + \varepsilon_{dt},$$

where the contrast $\delta_C - \delta_K$ is the corruption-specific component net of the road, and the geographic effect $\lambda_{g(d)t}$ runs from department-by-year to province-by-year as we tighten the control. The identifying variation is whether a district's nearby tenders were rigged or competitively bid, and its scarcity, only eight provinces contain both, is why Section 6.2 reports this object as bounds.

# 6 Results

The corruption is visible first in how much the projects cost and in how the tenders were allocated, and then, more faintly, in the localities the roads were built to serve. We take the two branches in turn, the Lava Jato concessions and the Construction Club cartel, each carried from procurement through to household welfare.

## 6.1 The Lava Jato concession branch

The Lava Jato branch asks what the bribed concessions cost and whether the corridors they built repaid the localities they crossed. Treatment is a bribe documented in a judicial ruling, taken from the Odebrecht case record assembled by Campos et al. (2021). We code three categories instead of a clean-versus-corrupt binary, a confirmed documented bribe, judicial implication without a project-specific payment, and no documented evidence. "No documented evidence" is not the same as clean, since the investigations reveal projects selectively, and we keep the categories separate.

Of the twenty-five Odebrecht projects in Peru that Campos et al. (2021) code, the ten with a confirmed documented bribe ran cost overruns averaging 53 percent of initial investment, against 17 percent for the projects with no documented evidence, with medians as wide, 47 against 8. In a cross-sectional regression of the overrun on the documented-bribe indicator the gap is +36 percentage points, moving to +25 with controls (Table 5). With twenty-five projects this is a description of the difference in fiscal cost, not an estimated causal elasticity.

Table 5: Documented bribery and cost overruns (descriptive, $N = 25$).

| | (1) bivariate | (2) +controls | (3) legal-or-interm. +controls |
|---|---|---|---|
| Documented bribe | $36.0^{***}$ | $22.1^{*}$ | $20.3^{*}$ |
| | (12.7) | (11.2) | (10.3) |
| Transport | | 20.8 | $26.9^{*}$ |
| | | (13.8) | (14.8) |
| log initial investment | | $11.7^{**}$ | $12.0^{**}$ |
| | | (4.8) | (4.7) |
| N projects | 25 | 25 | 25 |
| $R^2$ | 0.281 | 0.482 | 0.489 |

*Notes.* Cross-sectional regression of the cost overrun (percentage points of initial investment) on the documented-bribe indicator, $N = 25$ Odebrecht projects in Peru. Standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source.* Odebrecht–Peru case data, Campos et al. (2021).

The gap survives the small-sample inference twenty-five projects allow, randomization inference over the bribe labels at $p = 0.005$, a Webb-weight wild bootstrap at $p = 0.005$, a rank-sum test at $p = 0.002$, and a leave-one-megaproject-out range of $[25, 40]$, collected in Table 6. One bias deserves naming, because this gap carries much of the paper's quantitative corruption claim. Documentation is plausibly endogenous to the overrun, since a project that ran far over budget draws more scrutiny and is more likely to surface a documented bribe, which would inflate the gap. The direction of this bias is not clean. The December 2016 cross-country Odebrecht plea agreement that opened the Peruvian record was driven by the United States proceedings and not scrutiny of any single Peruvian project's cost, which limits selection at the level of which cases surfaced, but it does not discipline which Peruvian project drew a documented bribe, and treating the no-evidence projects as clean cuts the other way. We therefore read the gap as a descriptive magnitude, not a signed bound, and measure excess investment relative to initial budgets, not a rent.

Table 6: Small-sample robustness of the overrun gap ($N = 25$).

| Test | Result |
|---|---|
| Baseline gap (documented bribe vs no evidence) | +36 pp |
| With sector and size controls | +25 pp |
| Randomization inference over bribe labels | $p = 0.005$ |
| Webb-weight wild-cluster bootstrap | $p = 0.005$, CI $[10, 62]$ |
| Wilcoxon rank-sum test | $p = 0.002$ |
| Leave-one-megaproject-out range | $[25, 40]$ pp |
| Conditional gap, HC3 standard error | $p = 0.13$ |
| Monotone three-bucket overrun means (%) | 10.5/28.6/52.6 |

*Notes.* Robustness of the cost-overrun gap between documented-bribe and no-evidence Odebrecht projects, $N = 25$. The baseline and conditional gaps are from the regression in Table 5. The three-bucket means are the average overruns for the no-evidence, judicially-implicated, and confirmed-bribe projects. With twenty-five projects the gap is a description robust to small-sample inference, not an estimated elasticity, and the conditional gap weakens honestly under HC3. *Source.* Odebrecht–Peru case data, Campos et al. (2021).

If the bribed corridors cost more, the natural question is whether they delivered more locally, and the answer is that they delivered no detectable dividend. In the ENAHO district panel, estimated against log corridor kilometers, two market-access measures, and a travel-time measure, none of the twelve post-opening estimates on household income, consumption, or poverty is distinguishable from zero, and the nulls are stable to dropping the gold-mining departments, dropping Lima, guarding against spillovers, and leaving out each corridor. For the conservative market-access measure a ten-percent gain in access moves household income by about $-1.5$ percent and consumption by $-1.8$ percent, both indistinguishable from zero, with a corridor-level wild-cluster bootstrap returning an income $p$ of 0.55 and a consumption $p$ of 0.46. The corridors measurably improved accessibility, but the accessibility did not raise local welfare.

A null on average could hide redistribution toward the better-connected places a corridor links, the pattern Faber (2014) finds for China's highways, where activity concentrates in the central cities a road joins. We test it two ways and find little of it. Splitting the corridor-exposed districts by whether they are provincial-capital nodes or peripheral leaves no extra gain at the nodes, with a capital interaction of $+0.003$ on income ($p = 0.84$), so the null is not a local loss masked by an upstream gain captured at the capitals. The districts the corridor physically crosses do fare worse than the adjacent catchment on income, by a relative $-0.035$, consistent with an interregional trunk road that passes through the rural periphery without developing it, though with ten corridors this contrast is imprecisely identified and consumption does not echo it. The reading is that the concession corridors left no local dividend and at most a faint pass-through penalty, not a redistribution toward connected centers, so the null is informative, not mute.

One mechanism behind a costlier project that delivers no more is renegotiation after award, and the regulator's records are suggestive though thin. Across the seven audited concessions the three with a documented bribe average eight distinct contract addenda against five for the rest, consistent with bribed concessions being reopened more often, in the spirit of Decarolis (2014) and Coviello et al. (2018) on renegotiation absorbing value. With seven concessions this is descriptive, not a test. On the cartel side the analogous exercise is not available, because the electronic procurement records carry no addenda data for the large road licitaciones, the same gap that blocks a firm-level outcome design.

### 6.2 The Construction Club cartel branch

The Construction Club branch lets us see the allocation mechanism and its price directly. From the seventy-eight reconstructed tenders (Section 4) we record, for each firm, the processes it won, the processes in which it acted as a cover bidder, and the reference value of the works it won. We measure contract allocation and capture, not firm profits. Table 7 reports the firm-level counts.

Table 7: Cartel allocation and contract capture (top twelve firms).

| Firm | Won proc. | Cover proc. | No-supp. / unc. | Ref. value won (S/. M) | Departments | Indep. evid. |
|---|---|---|---|---|---|---|
| JJC | 14 | 28 | 8 | 1,111 | Amazonas, Ancash, Apurimac +17 | 27 |
| Iccgsa | 11 | 0 | 0 | 1,102 | Amazonas, Apurimac, Ayacucho +6 | 7 |
| CASA | 7 | 0 | 0 | 757 | Arequipa, Cajamarca, Cusco +6 | 4 |
| Johesa | 7 | 0 | 0 | 713 | Ayacucho, Cajamarca, Huancavelica +3 | 2 |
| E. Reyna | 5 | 0 | 0 | 553 | Arequipa, Cajamarca, Lambayeque +4 | 1 |
| Energoprojekt | 5 | 0 | 0 | 816 | Cajamarca, Cusco, Huancavelica +3 | 3 |
| Obrainsa | 4 | 38 | 0 | – | Amazonas, Ancash, Apurimac +16 | 18 |
| Constructora Malaga | 4 | 0 | 0 | 933 | Apurimac, Ayacucho, Lima +3 | 3 |
| Queiroz Galvao | 4 | 0 | 0 | 559 | Cajamarca, Cusco, Huanuco +2 | 3 |
| Altesa | 3 | 0 | 0 | 23 | Ancash, Cajamarca, Cusco +1 | 2 |
| Aramsa | 3 | 0 | 0 | 406 | Huanuco, Pasco, Ucayali | 1 |
| Camargo Correa | 3 | 0 | 0 | 374 | Cajamarca, Huanuco, Lambayeque +1 | 2 |
| *Reconstructed processes* | | | | | | 78 |
| *Share of wins, top 5 firms* | | | | | | 0.44 |
| *Firms with documented cover bids* | | | | | | 3 |
| *Herfindahl index of wins* | | | | | | 0.06 |
| *Reference value of cartelized tenders (S/. bn)* | | | | | | 6.3 |

*Notes.* Top twelve firms by documented involvement, with summary statistics over all thirty-two sanctioned groups. Cover bids are documented only for the three firms with structured apoyo exhibits. The full firm table is in Table 12. Won and cover columns count processes in which the firm was the designated winner or a documented cover bidder. Reference value won is a lower bound given partial value coverage. *Source.* Resolución 080-2021/CLC.

The structure is one of role division more than win concentration. Wins themselves are fairly diffuse, with a Herfindahl index of 0.06 and the top five groups taking 0.44 of documented wins, so no single firm dominates the awards. What the cartel organizes is the assignment of roles, a designated winner for each tender surrounded by coordinated cover bids. JJC wins fourteen of the seventy-eight tenders and Ingenieros Civiles y Contratistas Generales eleven, and the thirty-seven tenders with a reported reference value sum to about 6.3 billion soles. The cover-bid counts carry a caveat, since they are documented in full only for the three firms with structured bid exhibits, Cosapi, Obrainsa, and JJC, so the apparent concentration of cover bids in those firms partly reflects which records we could read. We therefore read the cover-bidding as coordinated support around designated winners (Figure 4), not as a full ranking of which firms covered most.

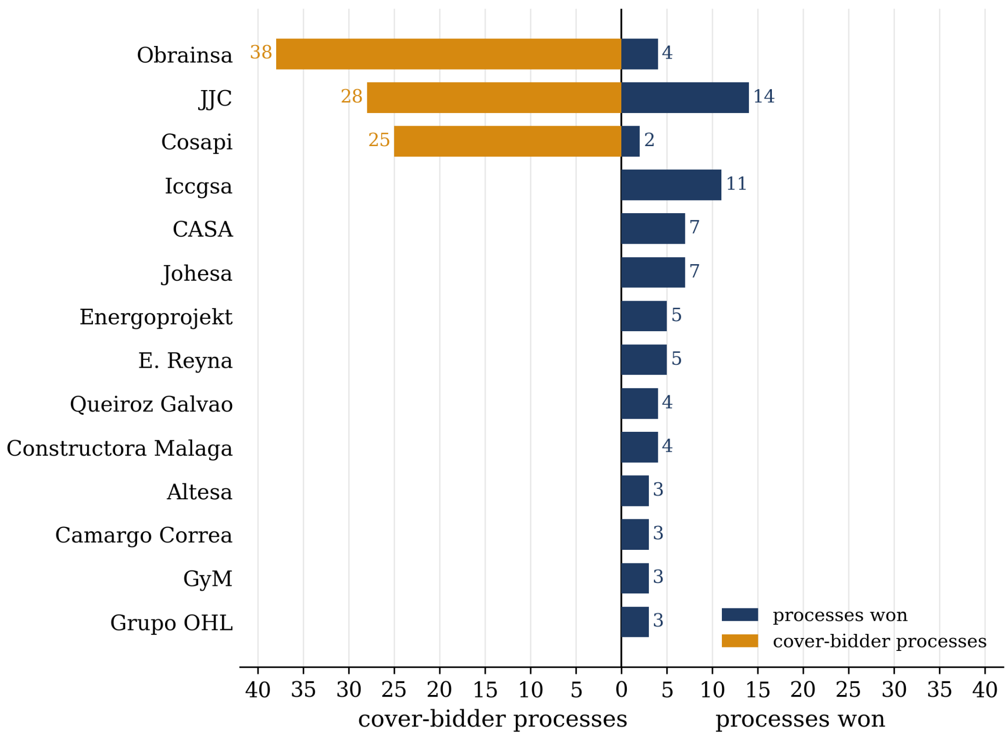


Figure 4: Role concentration: cover-bidder and won processes by firm.
*Notes.* Processes won (right) and processes in which the firm was a documented cover bidder (left), top firms by total documented involvement. A co-participation network of the firms is in Appendix Figure 12. *Source.* Resolución 080-2021/CLC.

The same records show what the coordination cost. Annex 3 of the resolution reports, for each sanctioned process, the competition authority's own estimate of the illicit benefit, and these sum to about 780 million soles against reference values of about 31 billion, a median estimated overcharge near four percent of the reference value. The bid record makes the mechanism visible. Among the offers Annex 2 classifies, the winning bid on a colluded tender is a median 108 percent of the reference value, at or near the ceiling the rules allow, while the winning bid on a competitive tender by the same firms is a median 90 percent, a ten-point discount to the State. The cartel did not only assign the winners, it erased the competitive discount, the designated-winner-and-cover-bid pattern that Porter and Zona (1993) and Conley and Decarolis (2016) recover statistically from bid distributions and that the resolution here establishes tender by tender. Colluded tenders also drew fewer bidders, a median of three against four. Figure 5 shows the two distributions barely overlapping.

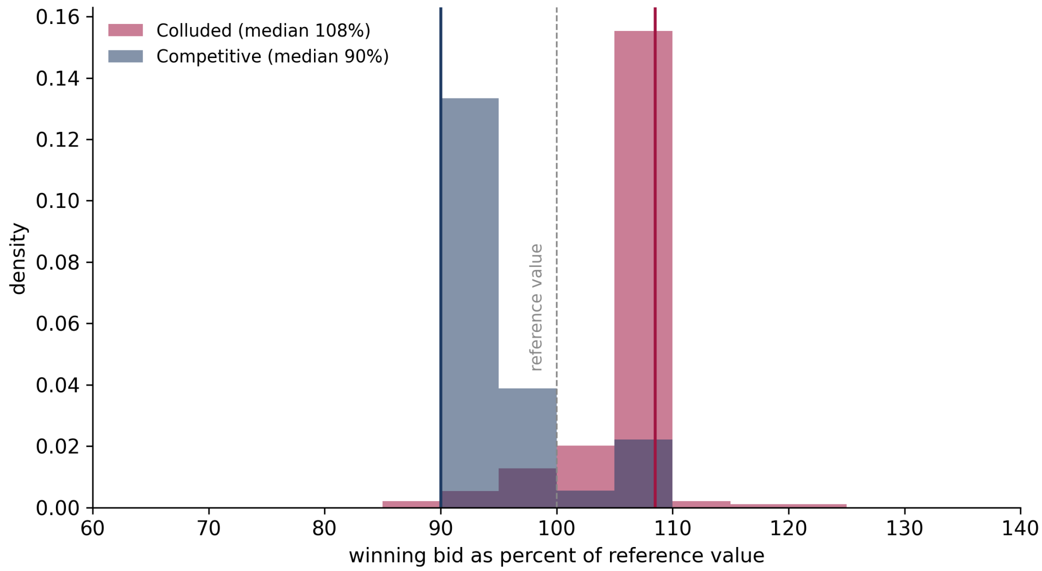


Figure 5: Winning bid as a percent of the reference value, colluded versus competitive tenders. *Notes.* Distribution of the winning bid as a percent of the reference value, for the tenders Annex 2 classifies as colluded (188 winning bids) and competitive (36), bid by the same cartel firms. Vertical lines mark the medians, 108 and 90 percent. The dashed line is the reference value. Bids are trimmed to the 40 to 160 percent range to drop transcription outliers. *Source.* Anexo 2 of Resolución 080-2021/CLC.

That this level gap is the genuine signal of collusion here, and not an artifact, is something we can verify and not assert, because Indecopi adjudicated each tender. Section 7 takes the detection literature's screens to these bids and scores them against that ground truth.

The overcharge is the price of the corruption on the tenders the cartel controlled. Whether it also reached the households around those tenders is the harder question, and we turn to it now. The broad cartel-procurement exposure is associated with adverse outcomes. Table 8 places the corridor and cartel measures in the same district panel with district and region-by-year fixed effects and region-clustered standard errors. The corridor measure is null throughout, while the cartel measure is negative per one standard deviation of exposure, $-0.016$ on income and $-0.017$ on consumption, and survives controlling for the corridor measure. Read alone this looks like a cartel effect, but nothing in the comparison yet separates the corruption from the road.

Table 8: Lava Jato corridor exposure versus Construction Club cartel exposure.

| | (1) Lava Jato corridor only | (2) Construction Club cartel only | (3) Both branches | (4) Both + interaction |
|---|---|---|---|---|
| *Log household income* | | | | |
| Lava Jato corridor exposure | −0.016 | | −0.016 | −0.009 |
| | (0.018) | | (0.018) | (0.018) |
| Construction Club cartel exposure | | −0.016** | −0.016** | −0.011* |
| | | (0.007) | (0.007) | (0.006) |
| Lava Jato × Construction Club | | | | −0.027*** |
| | | | | (0.005) |
| *Log household consumption* | | | | |
| Lava Jato corridor exposure | −0.019 | | −0.018 | −0.012 |
| | (0.016) | | (0.016) | (0.016) |
| Construction Club cartel exposure | | −0.017*** | −0.016*** | −0.012** |
| | | (0.005) | (0.005) | (0.005) |
| Lava Jato × Construction Club | | | | −0.023*** |
| | | | | (0.005) |
| *Poverty rate* | | | | |
| Lava Jato corridor exposure | 0.015 | | 0.015 | 0.010 |
| | (0.014) | | (0.014) | (0.014) |
| Construction Club cartel exposure | | 0.010 | 0.009 | 0.006 |
| | | (0.006) | (0.006) | (0.005) |
| Lava Jato × Construction Club | | | | 0.016*** |
| | | | | (0.004) |
| District FE | Yes | Yes | Yes | Yes |
| Region × year FE | Yes | Yes | Yes | Yes |
| Observations | 21,900 | 21,900 | 21,900 | 21,900 |

*Notes.* Coefficients are post-treatment effects per one standard deviation of exposure. Lava Jato corridor exposure is the concession-corridor market-access measure on the corridor opening clock. Construction Club cartel exposure is the proximity-decayed district exposure to the 76 geocoded Indecopi processes on the first-nearby-process clock. All columns include district and region-by-year fixed effects, with standard errors clustered by region. Stars denote $p < 0.1$, $p < 0.05$, $p < 0.01$.

The cartel measure compares districts near a cartelized tender to districts near no such tender, which includes districts near no major road tender at all. Building the identical proximity exposure for the Ministry of Transport national road licitaciones of the same period that are *not* in the cartel set produces the same association, −0.016 on income and −0.018 on consumption, indistinguishable from the cartel coefficient. The adverse association at this resolution is a feature of being near a large road tender, competitively or corruptly procured, and the broad measure recovers the road effect.

The cartel's own classification allows a sharper comparison. Annex 2 labels each tender the cartel firms entered as a colluded or a competitive offer, the same firms bidding on the same class of works, so comparing district exposure to colluded tenders against exposure to competitive ones holds the firms and the road class fixed and varies only whether the tender was rigged. Table 9 enters both exposures together. With district and department-by-year fixed effects and province-clustered standard errors, colluded exposure is adverse on all three outcomes, $-0.027$ on income, $-0.024$ on consumption, and $+0.013$ on poverty, each significant under a province wild-cluster bootstrap. Competitive exposure is null throughout, and unlike the broad dose the contrast survives a control for the size of the nearby roads, where the colluded coefficients hold or strengthen.

Table 9: Colluded versus competitive road exposure, by geographic fixed effect.

| | Colluded road exposure | Competitive road exposure |
|---|---|---|
| *Log household income* | | |
| District + department×year | −0.027 | 0.009 |
| | (0.028) [0.036] | (0.620) |
| District + province×year | −0.023 | 0.013 |
| | (0.165) | (0.649) |
| within both-arm provinces | −0.035 | 0.018 |
| | (0.474) | (0.331) |
| + road-size control | −0.032 | 0.008 |
| | (0.066) [0.100] | (0.684) |
| *Log household consumption* | | |
| District + department×year | −0.024 | −0.004 |
| | (0.024) [0.037] | (0.805) |
| District + province×year | −0.018 | 0.002 |
| | (0.220) | (0.935) |
| within both-arm provinces | −0.049 | 0.019 |
| | (0.153) | (0.021) |
| + road-size control | −0.045 | −0.009 |
| | (0.002) [0.008] | (0.533) |
| *Poverty rate* | | |
| District + department×year | 0.013 | 0.007 |
| | (0.033) [0.044] | (0.310) |
| District + province×year | 0.010 | 0.001 |
| | (0.229) | (0.927) |
| within both-arm provinces | 0.021 | 0.003 |
| | (0.264) | (0.826) |
| + road-size control | 0.027 | 0.011 |
| | (0.004) [0.022] | (0.146) |
| Observations (department, province rows) | 4,659 | |
| Observations (within both-arm provinces) | 414 | |

*Notes.* Coefficients are post-treatment effects per one standard deviation of proximity-decayed exposure, from a single regression entering colluded and competitive road exposure together. Colluded and competitive tenders are classified by the *tipo de oferta* column of Anexo 2 of Resolución 080-2021/CLC and are bid by the same cartel firms on the same class of national road works. Each panel adds a finer geographic control, from department-by-year to province-by-year to a within-province comparison restricted to the eight provinces that contain both a colluded and a competitive tender (27 districts). The road-size control adds the reference-value-weighted total of nearby road exposure at the department-by-year level. Analytic $p$-values clustered by province are in parentheses, wild-cluster bootstrap $p$-values in brackets where computed. The colluded minus competitive difference is not distinguishable from zero at the department level (wild-cluster $p$ of 0.17, 0.38, and 0.60) and is borderline only for consumption within the both-arm provinces ($p = 0.07$). *Source.* ENAHO *sumaria* district panel matched to Anexo 2.

The contrast does not survive the tightest geographic comparison, and the reason is informative and not fatal. Figure 6 traces the colluded and competitive coefficients across the fixed-effect ladder.

The colluded roads sit disproportionately in interior departments whose incomes rose through the commodity boom, while the competitive tenders cluster in a smaller set of coastal and urban provinces, so a comparison that crosses regions can confound the corruption with secular regional divergence. Moving from department-by-year to province-by-year fixed effects leaves the colluded coefficient almost unchanged in size but grows the standard error until it is no longer significant, and restricting to the eight provinces that contain both a colluded and a competitive tender returns the same adverse signs at similar magnitude with no precision, resting on twenty-seven districts. What changes across the ladder is precision, not the point estimate.

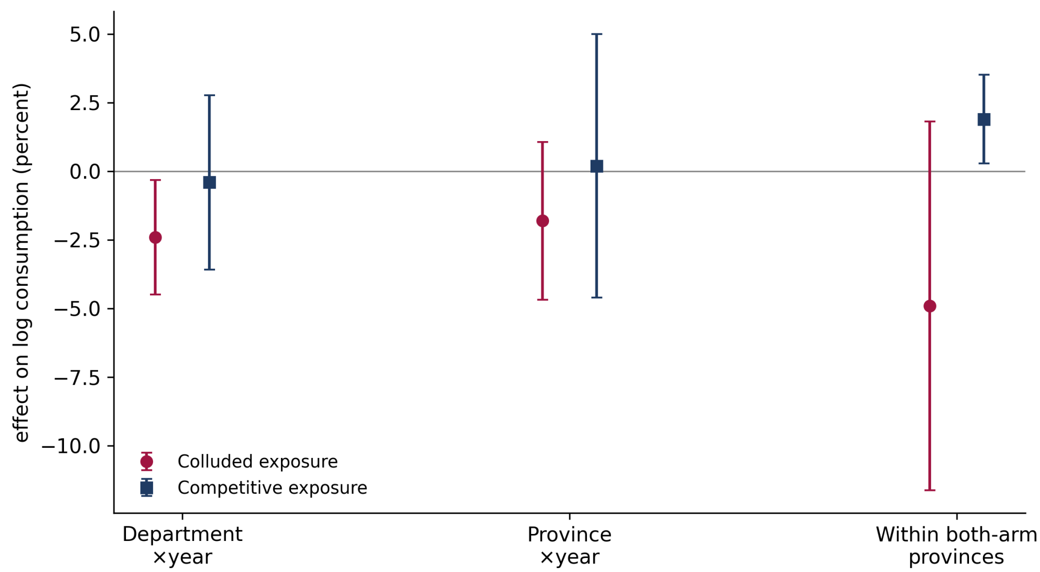


Figure 6: Colluded versus competitive exposure on consumption, across the geographic fixed-effect ladder.

*Notes.* Effect of colluded and competitive road exposure on log household consumption, per one standard deviation, at three geographic fixed-effect levels, with 95 percent province-clustered confidence intervals. The colluded point estimate stays adverse and of similar size while the interval widens, the precision loss from only eight provinces containing both kinds of tender. *Source.* ENAHO *sumaria* matched to Anexo 2.

Because the clean comparison is underpowered, we report the household incidence as bounds, and three exercises mark out what the data support. At the department level, where the colluded coefficient is significant, the minimum detectable effect at eighty percent power is about 3.4 percent for income, 2.9 percent for consumption, and 1.7 points for poverty, so the design rules out adverse effects larger than these but cannot certify smaller ones. For selection on unobservables, the Oster (2019) statistic for the colluded coefficient, benchmarked on the competitive exposure and the road-size control, is 0.51 for income, 0.24 for consumption, and 0.48 for poverty, all below one, so unobserved selection weaker than the observed controls would suffice to move the coefficient to zero, the formal counterpart of the concern that the cartel chose which tenders to rig. For parallel trends, the Rambachan and Roth (2023) relative-magnitudes restriction on the broad-exposure event study (Figure 7) lets the robust confidence interval include zero once a post-treatment deviation reaches about half the largest pre-treatment violation. The three agree. What is not point-identified is the downstream welfare *incidence*, not the price of the corruption, which the overcharge and the overruns identify directly. The household data narrow the range of that incidence, the adverse point estimates are not small and survive the road-size control, but they cannot rule out zero, and they cannot certify the effect as corruption-specific against the regional geography.

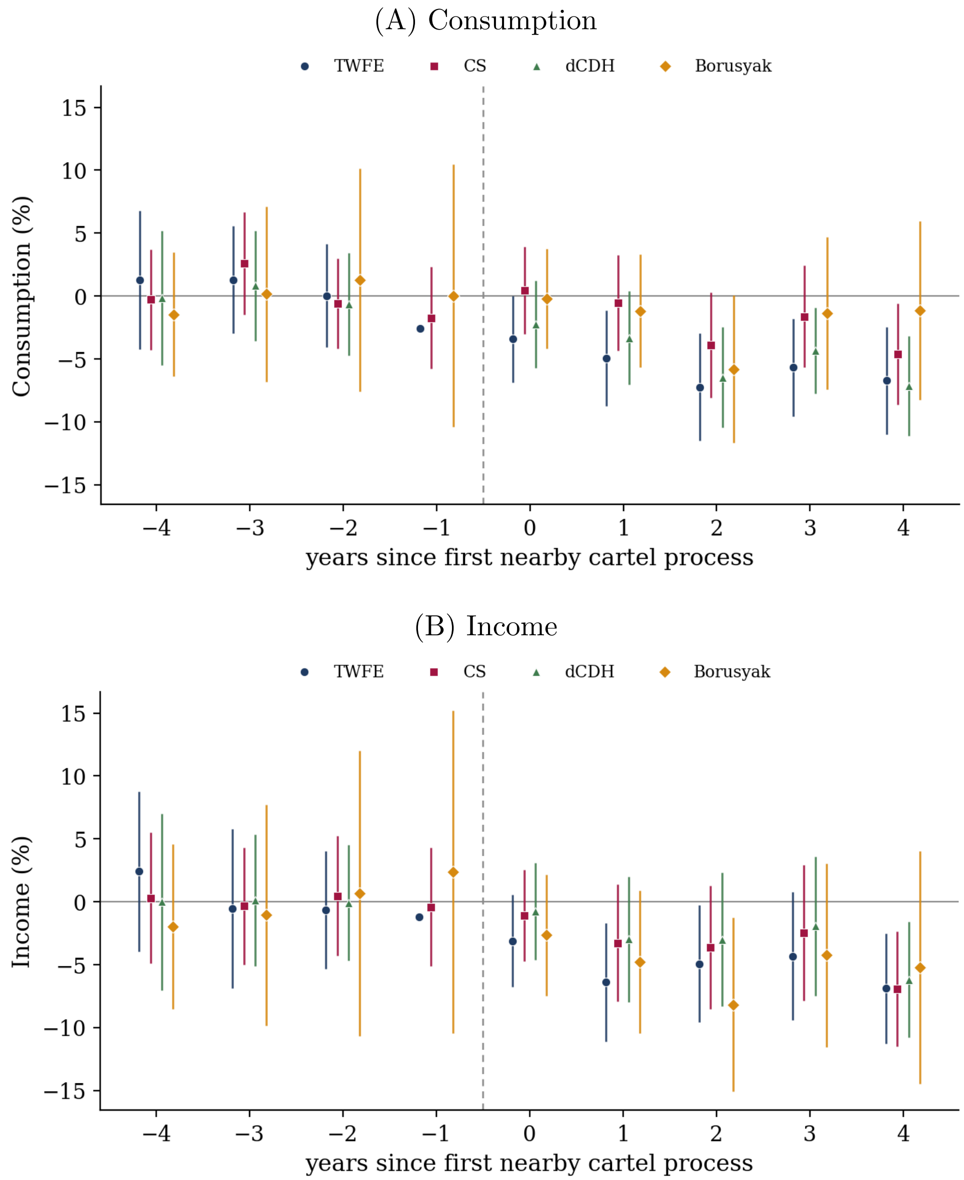


Figure 7: Household consumption (A) and income (B) by four estimators, road-tender exposure. *Notes.* Event study of the broader road-tender exposure association, normalized to each estimator's pre-period mean, with district and region-by-year fixed effects and region-clustered standard errors. Under the two estimators robust to staggered timing the pre-trends are flat and the post path grows adverse. The binary staggered estimators cannot anchor the colluded-versus-competitive contrast, because the competitive control group lacks common geographic support, a diagnostic in Appendix E. *Source.* ENAHO.

Where, then, does the cost of the corruption land? It is identified on the procurement side, in the larger overruns on the bribed concessions and the cartel's overcharge on the tenders it controlled, but its incidence on the localities is harder to find. The household evidence bounds and does not identify a local welfare effect, and a natural alternative, that the overcharge crowded out other local public investment, also fails to appear. Using the Invierte.pe project registry we measure public investment by function and ask whether districts and regions more exposed to the colluded tenders executed less non-transport investment after exposure. Figure 8 shows they did not. The effect on non-road investment is indistinguishable from zero at the district level, and while negative in point estimate at the region level it is imprecise and not echoed by the non-road share, so it does not survive as displacement, while at the national level transport and non-transport investment rose together. The reading that fits all of this is that the price of the corruption is a national-budget cost, visible in the overruns and the fiscal wedge, that did not measurably displace local welfare or local investment. The roads were paid for dearly out of the national treasury, and the locality the road crossed neither gained a dividend nor lost its other public spending.

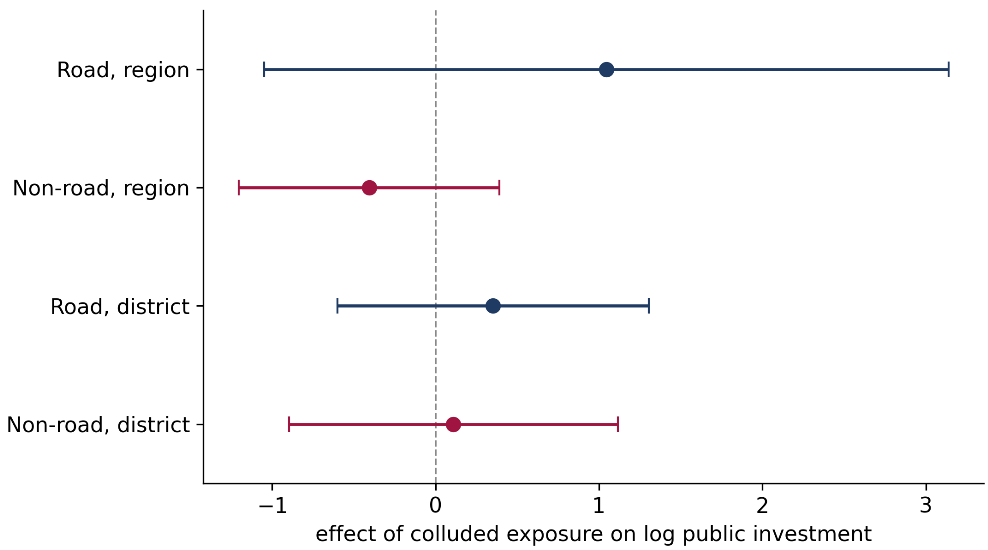


Figure 8: Fiscal crowd-out: effect of colluded-cartel exposure on public investment, by function and margin.

*Notes.* Each point is the coefficient on colluded-cartel exposure interacted with post-exposure, in a panel of public investment (Invierte.pe initiation-year project cost) with two-way fixed effects, district and region-by-year at the district margin and region and year at the region margin, clustered at the coarser unit. Bars are 95 percent confidence intervals. Non-road investment (the crowd-out outcome) is indistinguishable from zero at both margins, while road investment carries a positive point estimate that confirms the exposure tracks road activity. *Source.* MEF Invierte.pe project registry and Resolución 080-2021/CLC.

# 7 Detecting collusion against the adjudicated record

The collusion-detection literature works without the answer. Porter and Zona (1993), Conley and Decarolis (2016), Kawai and Nakabayashi (2022), and Chassang et al. (2022) build statistical screens to flag the designated-winner-and-cover-bid structure from bid data, and they must defend the screens without ever seeing which tenders were truly rigged. Our setting supplies the missing label. Indecopi adjudicated each tender as colluded or competitive, so we can run the screens on the same bids and score them against the truth, which is an adjudication the literature can rarely perform.

We take three screens to the Annex 2 bids. The first is a level screen, the winning bid as a share of the published reference value, which asks whether the competitive discount was suppressed. The second is the dispersion screen the literature most often uses, the coefficient of variation of the bids, which asks whether the bids were coordinated tightly. The third is a thin-competition screen, the number of bidders. Instead of a single cutoff, we treat each as a continuous classifier and trace its full receiver operating characteristic against Indecopi's classification across the 145 tenders with at least one classified bid, 117 colluded and 28 competitive.

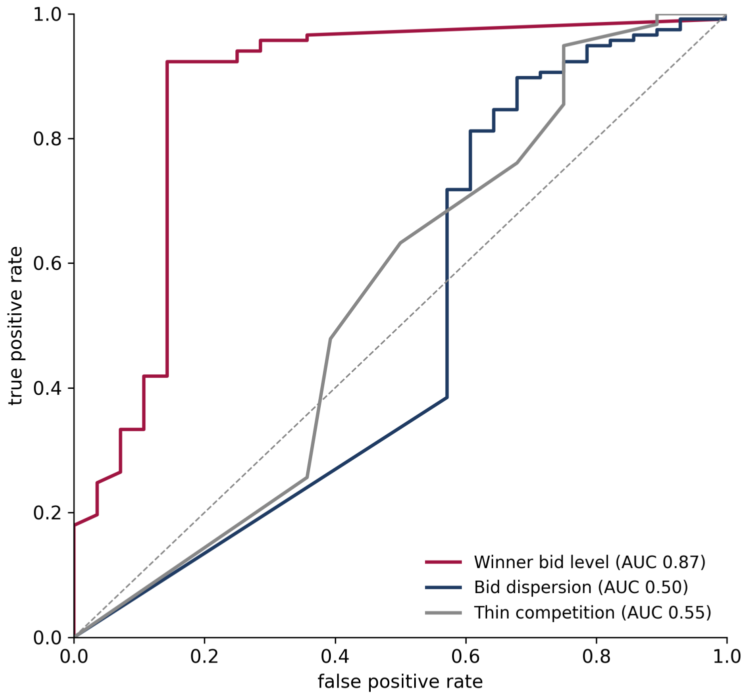


Figure 9: Receiver operating characteristic of three collusion screens against the adjudicated label. *Notes.* Each curve traces the true-positive against the false-positive rate as the screen's threshold varies, scoring tenders as colluded. The diagonal is chance. The area under the curve (AUC) is 0.87 for the winner-bid level, 0.50 for bid dispersion, and 0.55 for thin competition, with single-bid tenders carrying their zero dispersion so that all 145 tenders enter each screen. *Source.* Anexo 2 of Resolución 080-2021/CLC.

The level screen dominates. Its area under the curve is 0.87, against 0.50 for the dispersion screen, no better than chance once the single-bid tenders that carry zero dispersion are included, and 0.55 for thin competition. Restricted to the multi-bid tenders where dispersion is naturally measured the dispersion screen reaches only 0.62, still far below the level screen, so its weakness is not an artifact of that treatment. The advantage is statistically significant and not a feature of these particular tenders. Bootstrapping over tenders, the level screen's AUC exceeds the dispersion screen's by a gap whose 95 percent interval is $[0.23, 0.51]$, and its margin over the thin-competition screen is $[0.21, 0.45]$. An analytic DeLong test for the difference between correlated areas under the curve agrees, rejecting equality of the level and dispersion screens at $z = 5.3$ ($p < 0.001$) and of the level and thin-competition screens at $z = 5.3$ ($p < 0.001$). Table 10 reads off illustrative operating points at simple thresholds, where the level screen classifies 88 percent of tenders correctly while the dispersion screen at its median cut does no better than a coin flip.

Table 10: Collusion screens validated against the adjudicated label.

| Screen | Accuracy | False positive | False negative |
|---|---|---|---|
| Winner bids above the reference value | 0.88 | 0.14 | 0.12 |
| Low bid dispersion (CV below median) | 0.47 | 0.21 | 0.61 |
| Few bidders (three or fewer) | 0.61 | 0.50 | 0.37 |

*Notes.* Each screen predicts a tender as colluded, and the prediction is compared to Indecopi's adjudicated classification across 145 tenders with at least one classified bid (117 colluded, 28 competitive). Accuracy is the share correctly classified, the false-positive rate is among competitive tenders, the false-negative rate among colluded tenders. *Source.* Anexo 2 of Resolución 080-2021/CLC.

The ranking is informative for the literature, not only for this case. Where procurement bids are anchored to a published reference value, the cartel's gain is to erase the discount the State would otherwise capture, so the signal lives in the level of the winning bid and not in the spread of the losing ones that a dispersion screen is built to detect. A screen tuned to variance can miss a cartel that coordinates on price instead of rank. We can say this because we observe the assignment that Chassang et al. (2022) and Kawai et al. (2023) must infer, and the judicial record turns a methodological debate into a measurement.

# 8 Conclusions

Two judicial records let us study Peruvian road corruption with the corrupt projects identified and not inferred. The Odebrecht case names the bribery behind a set of road concessions, and the Indecopi resolution names the firms that rigged the procurement of national road works for over a decade. Combining both with procurement and household data, we measure the procurement-side cost and distortion of the corruption and ask what the exposed localities received in return.

The procurement side is costly and distorted, and this is where the price of corruption is identified. Documented-bribe concession projects ran cost overruns above half of initial investment against under a fifth for the rest, and the cartel allocated road tenders through a small set of repeat winners surrounded by coordinated cover bids. The household side is a roads-and-development result. The roads built in this period delivered no compensating local welfare, and this holds whether they were competitively or corruptly procured. Districts near cartelized tenders and near competitively procured tenders show the same broad outcomes, so that association is a road effect. Using the cartel's own classification of the tenders it rigged against those the same firms bid competitively, an adverse association specific to the colluded tenders survives a road-size control and is significant within departments, but it loses precision under the tightest within-province comparison, which only eight provinces support, so we read the household evidence as suggestive.

We are deliberate about scope. We reconstruct contract allocation and capture from the antitrust record, not firm profits, because the procurement data omit the large works and the consortia that won them, and the firm-level outcome designs used elsewhere fail here for the same reason. The cartel chose the geography of the tenders it rigged, and the competitively bid road works that would anchor a clean within-province comparison sit in only eight provinces, so the household data can show that the colluded tenders carry an adverse association robust to road size but cannot pin it down separately from the local geography. What the two records jointly support is a cost-effectiveness conclusion. The corruption is plainly visible and identified in the cost and the allocation of the projects, the household evidence points the same way without being decisive, and the roads it procured brought no commensurate local return.

## References


Sam Asher and Paul Novosad. Rural roads and local economic development. *American Economic Review*, 110(3):797–823, 2020.

Patrick Bajari and Lixin Ye. Deciding between competition and collusion. *Review of Economics and Statistics*, 85(4):971–989, 2003.

Oriana Bandiera, Andrea Prat, and Tommaso Valletti. Active and passive waste in government spending: Evidence from a policy experiment. *American Economic Review*, 99(4):1278–1308, 2009.

Kirill Borusyak, Xavier Jaravel, and Jann Spiess. Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies*, 91(6):3253–3285, 2024.

Erica Bosio, Simeon Djankov, Edward Glaeser, and Andrei Shleifer. Public procurement in law and practice. *American Economic Review*, 112(4):1091–1117, 2022.

Robin Burgess, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró i Miquel. The value of democracy: Evidence from road building in kenya. *American Economic Review*, 105(6): 1817–1851, 2015.

Brantly Callaway and Pedro H. C. Sant'Anna. Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230, 2021.

Nicolas Campos, Eduardo Engel, Ronald D. Fischer, and Alexander Galetovic. The ways of corruption in infrastructure: Lessons from the odebrecht case. *Journal of Economic Perspectives*, 35(2):171–190, 2021.

Sylvain Chassang, Kei Kawai, Jun Nakabayashi, and Juan Ortner. Robust screens for noncompetitive bidding in procurement auctions. *Econometrica*, 90(1):315–346, 2022.

Robert Clark, Decio Coviello, Jean-François Gauthier, and Art Shneyerov. Bid rigging and entry deterrence in public procurement: Evidence from an investigation into collusion and corruption in quebec. *Journal of Law, Economics, and Organization*, 34(3):301–363, 2018.

Emanuele Colonnelli and Mounu Prem. Corruption and firms. *Review of Economic Studies*, 89(2): 695–732, 2022.

Timothy G. Conley and Francesco Decarolis. Detecting bidders groups in collusive auctions. *American Economic Journal: Microeconomics*, 8(2):1–38, 2016.

Decio Coviello, Andrea Guglielmo, and Giancarlo Spagnolo. The effect of discretion on procurement performance. *Management Science*, 64(2):715–738, 2018.

Clément de Chaisemartin and Xavier D'Haultfœuille. Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996, 2020.

Francesco Decarolis. Awarding price, contract performance, and bids screening: Evidence from procurement auctions. *American Economic Journal: Applied Economics*, 6(1):108–132, 2014.

Dave Donaldson and Richard Hornbeck. Railroads and american economic growth: A "market access" approach. *Quarterly Journal of Economics*, 131(2):799–858, 2016.

Benjamin Faber. Trade integration, market size, and industrialization: Evidence from china's national trunk highway system. *Review of Economic Studies*, 81(3):1046–1070, 2014.

Claudio Ferraz and Frederico Finan. Exposing corrupt politicians: The effects of brazil's publicly released audits on electoral outcomes. *Quarterly Journal of Economics*, 123(2):703–745, 2008.

Raymond Fisman. Estimating the value of political connections. *American Economic Review*, 91 (4):1095–1102, 2001.

Ejaz Ghani, Arti Grover Goswami, and William R. Kerr. Highway to success: The impact of the golden quadrilateral project for the location and performance of indian manufacturing. *Economic Journal*, 126(591):317–357, 2016.

Kei Kawai and Jun Nakabayashi. Detecting large-scale collusion in procurement auctions. *Journal of Political Economy*, 130(5):1585–1629, 2022.

Kei Kawai, Jun Nakabayashi, Juan Ortner, and Sylvain Chassang. Using bid rotation and incumbency to detect collusion: A regression discontinuity approach. *Review of Economic Studies*, 90 (1):376–403, 2023.

Robert C. Marshall and Leslie M. Marx. *The Economics of Collusion: Cartels and Bidding Rings*. MIT Press, Cambridge, MA, 2012.

Maxim Mironov and Ekaterina Zhuravskaya. Corruption in procurement and the political cycle in tunneling: Evidence from financial transactions data. *American Economic Journal: Economic Policy*, 8(2):287–321, 2016.

Melanie Morten and Jaqueline Oliveira. The effects of roads on trade and migration: Evidence from a planned capital city. *American Economic Journal: Applied Economics*, 16(2):333–367, 2024.

Benjamin A. Olken. Monitoring corruption: Evidence from a field experiment in indonesia. *Journal of Political Economy*, 115(2):200–249, 2007.

Benjamin A. Olken and Rohini Pande. Corruption in developing countries. *Annual Review of Economics*, 4:479–509, 2012.

Emily Oster. Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, 37(2):187–204, 2019.

Martin Pesendorfer. A study of collusion in first-price auctions. *Review of Economic Studies*, 67 (3):381–411, 2000.

Robert H. Porter and J. Douglas Zona. Detection of bid rigging in procurement auctions. *Journal of Political Economy*, 101(3):518–538, 1993.

Ashesh Rambachan and Jonathan Roth. A more credible approach to parallel trends. *Review of Economic Studies*, 90(5):2555–2591, 2023.

# A Reconstruction of the two judicial datasets

The paper's primary contribution is a dataset built from two independent judicial records. Because neither was assembled before, this appendix documents how each was reconstructed and validated, so the universe of tenders, firm roles, and project costs can be reproduced from the public rulings.

## A.1 The Construction Club cartel

The source is Resolución 080-2021/CLC of the Peruvian competition authority, a document of about 2,630 pages that sanctions thirty-two economic groups for rigging the procurement of national road works between November 2002 and December 2016. The ruling does not present a single allocation table. The cartel's conduct is recorded instead in firm-specific support exhibits, each listing the tenders in which one firm submitted cover bids in favor of the designated winner. We transcribed the three structured exhibits page by page from the rendered pages, not the noisy text layer: the Cosapi exhibit (twenty-five processes), the Obrainsa exhibit (thirty-eight), and the JJC exhibit (fifty-two). Each row records the process code, the project, the bidding consortium, the winner, and whether support was given.

We normalized the process codes, which appear in two parallel notations in the ruling, a classic licitación form and a Ministry-of-Transport subunit form, and merged the three exhibits into a deduplicated union of seventy-eight distinct road tenders. Merges were made only on a matching code together with a compatible project name and a compatible winner, and cross-notation candidates and within-exhibit collisions were resolved by hand against the ruling's narrative. Two cases illustrate the discipline. The Obrainsa exhibit lists the Túnel Callao consortium as the winner of licitación 7-2011, which the ruling states verbatim is an error referring to a different project, so we record the winner the resolution names. Licitación 11-2010 is the one documented appeal case, where the designated winner was disqualified and the award passed on appeal to another cartel member, which we code as a within-cartel reassignment, not an outside win. Across the seventy-eight tenders only one has an outside winner, which is why a design based on tenders the cartel tried and failed to control is not available.

Two annexes supply the quantities. Annex 2, on pages 2578 to 2625, is an offer-level bid database of 571 rows carrying for each bid the process code, the bidding group, the winner, the reference and offered values, and the ruling's classification of the offer as colluded or competitive, which is the variation the household contrast in Section 6.2 uses. Annex 3 reports, for each sanctioned process and firm, the reference value, the estimated illicit benefit, and the base fine, 193 rows in total. The illicit-benefit column is the competition authority's own forensic estimate of the overcharge and is the basis for the cartel-rent figure in Section 6.2. We validated the reconstructed universe against the national electronic procurement records, matching about four in five of the post-2004 codes on the full process signature for existence and reference value, with the pre-2004 shortfall a documented consequence of the electronic system beginning in 2004. Winners and roles are always taken from the resolution and never from the procurement records, which have structural gaps in the award data for large works. The reconstruction yields the firm-process role table, the proximity-decayed district exposure measure, and the per-process illicit benefit used in the main text.

## A.2 The Lava Jato concessions

The Lava Jato branch is built from the Odebrecht record, whose collaboration agreements and prosecutions name the road concessions, the bribes, and the officials who received them. From this record we identify the set of national road concessions and which carry a documented bribe. Project fiscal cost follows Campos et al. (2021), whose comparison of bribed and non-bribed infrastructure costs we reproduce on the Peruvian subset, reading the overrun as a magnitude on the twenty-five projects the comparison allows. We treat the 2016 disclosure of the Odebrecht agreements as the event that dates exposure for the concession branch.

For the local analysis we date when each corridor's improved works opened to traffic. We mined the regulator's annual reports, about one hundred documents, for opening dates and retained, for each corridor, the first verifiable opening of improved works, not the contract or exploitation date, accepting a date only when it agreed across independent reports and rejecting dates that recurred across many concessions as report artifacts. The corridor accessibility measure follows Donaldson and Hornbeck (2016). We build a noded road network, assign population-weighted origins from a gridded population surface and explicit feeder times to the network, and compute the change in market access to provincial-capital destinations from corridor-specific documented pre-improvement speeds. We disclose one defect in the underlying network, a structural break on the central highway between Huánuco and Lima that the raw network renders as a long detour, which we flag and which largely cancels in the accessibility difference. The household panel merges this exposure onto the national household survey at the district level.

# B Reconstructing the cartel from Resolución 080

The cartel universe is read directly from Resolución 080-2021/CLC-INDECOPI (Expediente 001-2020/CLC), the 2,630-page antitrust ruling, not its noisy optical-character text. The resolution documents the cartel through firm-specific evidence exhibits in which each firm's own records list the tenders it participated in and the consortium that took each award. Three of these are structured process tables, and we transcribe them at the page level: the Cosapi exhibit (CSP18, twenty-five processes), the Obrainsa exhibit (OBN10, thirty-eight processes), and the JJC exhibit (JJC01, fifty-two processes). The remaining named evidence (Johesa, GyM, San Martín and others) is narrative or email, not a process table, and we use it only to corroborate participation.

Merging the three exhibits requires care because the same tender appears under different code notations across firms, a *clásico* number in one exhibit and an MTC code in another. We match on a conservative full code signature and confirm each candidate on project name and designated winner before merging, so that distinct sub-tramos and distinct lots are never collapsed by code similarity alone. Several cases were resolved by returning to the resolution text: process LP 7-2011-MTC/20 appears in one exhibit with an erroneous consortium that the resolution states corresponds to a different project, the LPN 8-2005 process splits into two awarded tramos with different winners, and the third member of the 6-2011 consortium was corrected against the resolution. The union is seventy-eight distinct processes spanning 2002 to 2016. For each firm in each process we record a role, winning-consortium member, cover bidder, cover bidder of uncertain target, or participant without support, and these roles and the designated winners are the source for the firm-level allocation in Section 6.2. Winners and roles always come from the resolution, never from the procurement data.

# C Validation against the procurement records

We validate the seventy-eight processes against Peru's open procurement data, the OCDS export of the SEACE system. An early attempt matched poorly for 2015 and 2016, which we traced to our own processing, not the data: for the post-2015 regime the process code lives in the record's `officialNumber` and the work name in its `description`, and an earlier extract had kept only the title field. Re-extracting the raw OCDS with the full fields recovers the missing processes. Of the seventy-eight, seventy-three (94 percent) are found in the enriched procurement layer and sixty-three (81 percent) match at high confidence on the licitación type, the MTC sub-unit, and the work name. The unmatched remainder is four pre-2004 processes that predate digital procurement coverage and a small set with a resolution-versus-SEACE code numbering mismatch.

Two limits of the procurement data shape how we use it. The award and supplier records are empty for the large road licitaciones, so SEACE is not a usable source of winners for the works that matter, and winner identity therefore comes from the resolution while SEACE serves only to confirm existence, code, entity, work, date, and the order of magnitude of the value. For the same reason we do not construct a cartelized-share-of-procurement measure, because the per-firm award denominator is too incomplete for the cartel firms. The exposure measure is built entirely from the resolution.

# D Geocoding and district exposure

We geocode each process from the endpoint towns named in its work title, matching the place names to the centroids of the INEI district polygons and resolving capital-town aliases whose district carries a different name (Aguaytía to Padre Abad, Tingo María to Rupa-Rupa, Yauri to Espinar, and similar). Seventy-six of the seventy-eight processes are geocoded across twenty-one departments, the two exceptions being a pre-2004 process with no usable endpoint and an urban grade-separated interchange with no corridor. For each of the 1,891 INEI districts we compute the count, reference value, and cover-bidder count of cartel processes within five, ten, and twenty kilometers, a distance-decayed intensity at each scale, and a treatment year equal to the earliest nearby cartel process. For the maps we snap each process to the actual MTC road network and route between its endpoints, so the cartelized routes follow real roads instead of straight lines. The endpoint-based geocoding is an approximation, and the proximity bands and decay scales in Appendix E show the household result is not an artifact of any single bandwidth.

# E Household robustness

Table 11 reports the post-treatment effect of the cartel-exposure measure on household income and consumption across the robustness battery: proximity measured by a distance-decayed intensity and by a count within five, ten, and twenty kilometers, district weighting by household count, alternative clustering, the three treatment cohorts, and the leave-one-region-out range. Consumption is stable throughout. Income is the same in sign and magnitude but loses significance under province and two-way clustering and when the two densest regions are dropped together, which is why the main text reads income as the more fragile outcome.

Table 11: Household robustness, cartel exposure (post-treatment, per one standard deviation).

| | Income | | Consumption | |
|---|---|---|---|---|
| Specification | $\beta$ | $p$ | $\beta$ | $p$ |
| Decay, 5 km | −0.015 | 0.074 | −0.017 | 0.004 |
| Decay, 10 km (baseline) | −0.016 | 0.027 | −0.017 | 0.003 |
| Decay, 20 km | −0.019 | 0.013 | −0.016 | 0.005 |
| Count within 5 km | −0.019 | 0.045 | −0.023 | 0.003 |
| Count within 20 km | −0.026 | 0.011 | −0.021 | 0.012 |
| Weighted by households | −0.021 | 0.042 | −0.021 | 0.002 |
| Cluster by province | −0.016 | 0.054 | −0.017 | 0.007 |
| Cluster two-way | −0.016 | 0.074 | – | – |
| Cohort 2004–2008 | −0.015 | 0.158 | −0.016 | 0.022 |
| Cohort 2009–2011 | −0.022 | 0.175 | −0.019 | 0.082 |
| Cohort 2012–2016 | −0.028 | 0.250 | −0.027 | 0.137 |
| Leave-one-region-out | [−0.020, −0.014] | | [−0.020, −0.014] | |

*Notes.* Post-treatment effect per one standard deviation of the cartel-exposure measure, district and region-by-year fixed effects, standard errors clustered by region unless noted. Cohorts restrict the treated set to districts first exposed in the period shown, keeping the never-exposed as controls. The two-way consumption cell is omitted because the variance estimator is singular in that specification. *Source.* ENAHO *sumaria* district panel and Resolución 080-2021/CLC.

The staggered-design estimators also pass their native pre-trend checks. The joint placebo test of the de Chaisemartin–D'Haultfœuille estimator does not reject flat pre-trends for income ($p = 1.00$), consumption ($p = 0.58$), or poverty ($p = 0.35$), and the Borusyak pre-trend test is likewise not rejected ($p \approx 0.78$ to 0.85). The one pre-trend that fails is two-way fixed effects on the poverty rate ($p = 0.04$), the specification the staggered-design critique targets, which is a further reason the main text reads the result off the staggered-robust estimators.

### E.1 Why the colluded-versus-competitive contrast is read off the fixed-effect ladder

The colluded-versus-competitive contrast in Table 9 cannot be estimated with the binary staggered-robust estimators, because the competitive control group lacks common geographic support. The competitive tenders fall in a small set of coastal and urban departments, while the colluded cohorts span about twenty mostly interior departments, most of which contain no competitive tender at all. The Callaway–Sant'Anna group-time estimates are correspondingly incoherent. Cohort effects on log income range from +0.41 to −0.18 and reach economically implausible magnitudes at longer horizons, with raw group-time effects as large as +0.67, and the cohorts that read most positive are precisely those whose districts were already growing fastest before any treatment, with pre-treatment income slopes of eleven to thirteen percent a year against about five percent in the controls. The positive aggregate is a weighted average of cross-regional comparisons that violate parallel trends, not a treatment effect. The de Chaisemartin–D'Haultfœuille estimator, which rests on more local comparisons, returns a null with clean placebo tests. We therefore read the contrast off the geographic fixed-effect ladder of the main text, which conditions on region-by-year variation and returns a stable point estimate, and we treat the binary staggered estimates as uninformative here, not evidence either way.

# F Raw cartel event studies and poverty

Figure 10 reports the cartel-exposure event studies for income and consumption without the pre-period normalization used in the main text, each estimator on its own native baseline. Callaway–Sant'Anna and de Chaisemartin–D'Haultfœuille have flat raw pre-trends and an adverse post-treatment path, two-way fixed effects corroborates the sign, and the Borusyak estimator has noisier raw pre-period point estimates and is sensitive to the normalization, which is why the main text relies on the two staggered-robust estimators.

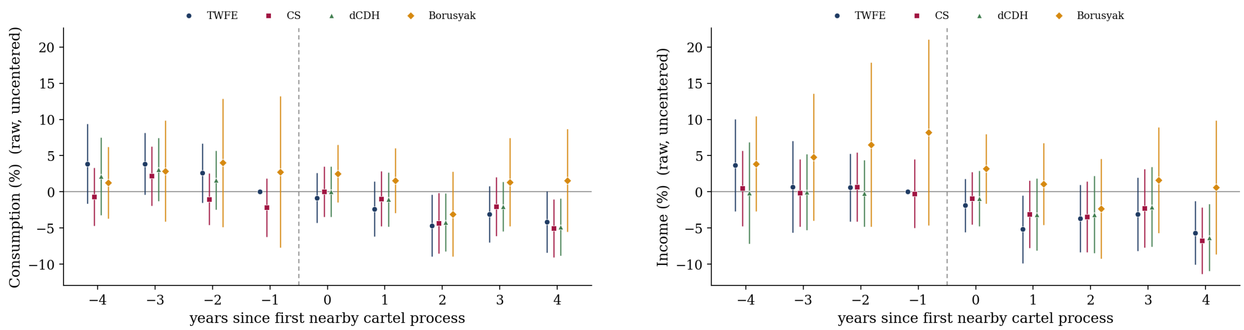


Figure 10: Raw, uncentered cartel event studies: consumption (left) and income (right).
*Notes.* Coefficients shown without pre-period normalization. *Source.* ENAHO *sumaria* district panel and Resolución 080-2021/CLC.

Figure 11 reports the poverty event study, which is directionally adverse but weaker than consumption and income.

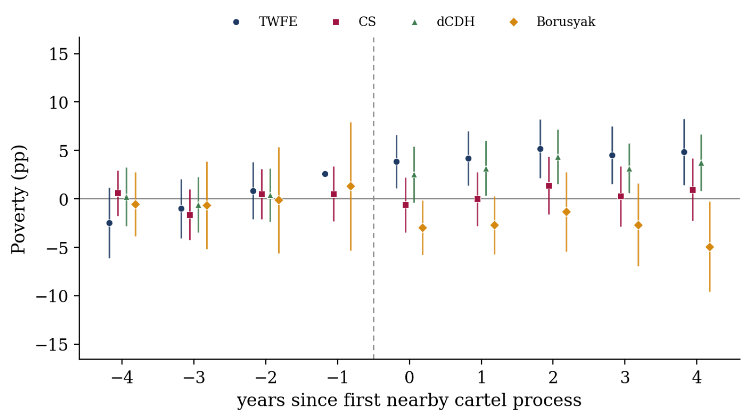


Figure 11: Poverty-rate event study by four estimators, cartel exposure.
*Notes.* Normalized to each estimator's pre-period mean. *Source.* ENAHO *sumaria* district panel and Resolución 080-2021/CLC.

# G Additional figures and the full firm table

Table 12: Cartel allocation and contract capture, all firms.

| Firm | Won proc. | Cover proc. | No-supp. / unc. | Ref. value won (S/. M) | Departments | Indep. evid. |
|---|---|---|---|---|---|---|
| JJC | 14 | 28 | 8 | 1,111 | Amazonas, Ancash, Apurimac +17 | 27 |
| Iccgsa | 11 | 0 | 0 | 1,102 | Amazonas, Apurimac, Ayacucho +6 | 7 |
| CASA | 7 | 0 | 0 | 757 | Arequipa, Cajamarca, Cusco +6 | 4 |
| Johesa | 7 | 0 | 0 | 713 | Ayacucho, Cajamarca, Huancavelica +3 | 2 |
| E. Reyna | 5 | 0 | 0 | 553 | Arequipa, Cajamarca, Lambayeque +4 | 1 |
| Energoprojekt | 5 | 0 | 0 | 816 | Cajamarca, Cusco, Huancavelica +3 | 3 |
| Obrainsa | 4 | 38 | 0 | – | Amazonas, Ancash, Apurimac +16 | 18 |
| Constructora Malaga | 4 | 0 | 0 | 933 | Apurimac, Ayacucho, Lima +3 | 3 |
| Queiroz Galvao | 4 | 0 | 0 | 559 | Cajamarca, Cusco, Huanuco +2 | 3 |
| Altesa | 3 | 0 | 0 | 23 | Ancash, Cajamarca, Cusco +1 | 2 |
| Aramsa | 3 | 0 | 0 | 406 | Huanuco, Pasco, Ucayali | 1 |
| Camargo Correa | 3 | 0 | 0 | 374 | Cajamarca, Huanuco, Lambayeque +1 | 2 |
| Conalvias | 3 | 0 | 0 | 496 | Cajamarca, Cusco, Lambayeque +1 | 2 |
| Grupo OHL | 3 | 0 | 0 | 600 | Apurimac, Ayacucho, Cusco | 1 |
| GyM | 3 | 0 | 0 | 571 | Apurimac, Ayacucho, Huancavelica | 2 |
| Upaca | 3 | 0 | 0 | 384 | Ancash, Huancavelica, Lima | 1 |
| Cosapi | 2 | 25 | 0 | 409 | Apurimac, Arequipa, Ayacucho +11 | 19 |
| Aterpa | 2 | 0 | 0 | 71 | Amazonas, Lima | 0 |
| Conciviles | 2 | 0 | 0 | 329 | Cajamarca, Cusco, Lambayeque | 2 |
| Constructora TP | 2 | 0 | 0 | 490 | Apurimac, Ayacucho, Cusco | 1 |
| JCCG | 2 | 0 | 0 | 182 | Apurimac, Ayacucho | 0 |
| Superconcreto | 2 | 0 | 0 | 404 | Cajamarca, La Libertad, Lambayeque +1 | 1 |
| Andrade Gutierrez | 1 | 0 | 0 | 469 | Lima | 1 |
| CyM | 1 | 0 | 0 | 110 | Ayacucho, Huancavelica | 0 |
| Eivisac | 1 | 0 | 0 | 442 | Ayacucho | 1 |
| Grupo Plaza | 1 | 0 | 0 | 163 | La Libertad | 1 |
| OAS | 1 | 0 | 0 | – | Lima | 1 |
| San Martin | 1 | 0 | 0 | 357 | Pasco, Ucayali | 0 |

*Notes.* The table reports contract allocation and capture within the Construction Club, reconstructed from Resolución 080-2021/CLC. It does not measure firm profits. Won and cover-bidder columns count processes in which the firm was the designated winner or a documented cover bidder. Reference value won sums available reference values of won processes and is a lower bound given partial value coverage. Independent evidence counts processes corroborated by more than one firm exhibit.

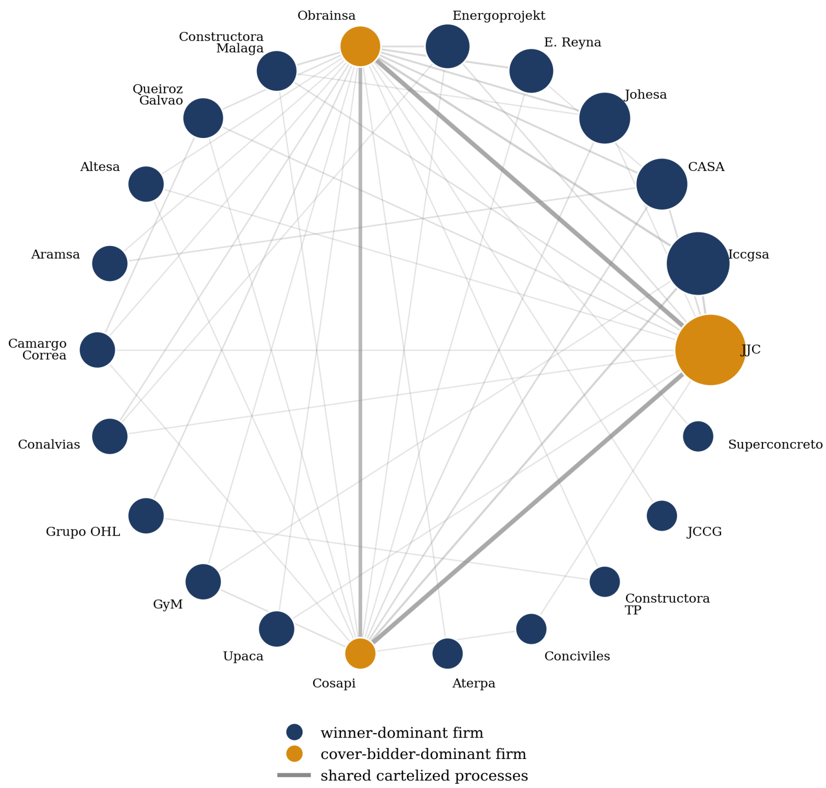


Figure 12: Cartel co-participation network.

*Notes.* Firms that participated together in the same cartelized tender are linked, with link weight the number of shared processes (shown for two or more). Node size is processes won. Navy nodes are winner-dominant firms, ochre nodes cover-bidder-dominant. *Source.* Resolución 080-2021/CLC.

The firm-role matrix in Figure 13 gives the same allocation as counts.

| | Won | Cover | No support | Depts | Processes |
|---|---|---|---|---|---|
| JJC | 14 | 28 | 8 | 20 | 50 |
| Obrainsa | 4 | 38 | 0 | 19 | 42 |
| Cosapi | 2 | 25 | 0 | 14 | 27 |
| Iccgsa | 11 | 0 | 0 | 9 | 11 |
| CASA | 7 | 0 | 0 | 9 | 7 |
| Johesa | 7 | 0 | 0 | 6 | 7 |
| E. Reyna | 5 | 0 | 0 | 7 | 5 |
| Energoprojekt | 5 | 0 | 0 | 6 | 5 |
| Constructora Malaga | 4 | 0 | 0 | 6 | 4 |
| Queiroz Galvao | 4 | 0 | 0 | 5 | 4 |
| Altesa | 3 | 0 | 0 | 4 | 3 |
| Aramsa | 3 | 0 | 0 | 3 | 3 |
| Camargo Correa | 3 | 0 | 0 | 4 | 3 |
| Conalvias | 3 | 0 | 0 | 4 | 3 |
| Grupo OHL | 3 | 0 | 0 | 3 | 3 |
| GyM | 3 | 0 | 0 | 3 | 3 |

Figure 13: Firm-by-role matrix: won, cover, no-support, departments, processes.
*Notes.* Each column is colour-normalized to its own maximum, with the raw count shown. Firms ordered by total documented involvement. *Source.* Resolución 080-2021/CLC.

The cartel footprint accumulates over time as more tenders are let, shown in Figure 14.

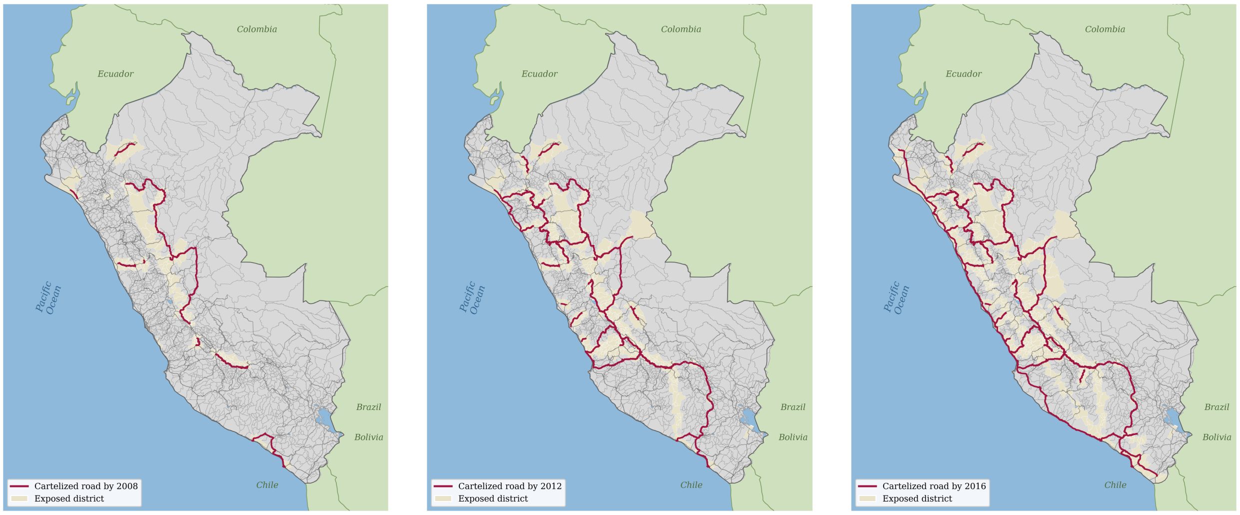


Figure 14: Evolution of cartel exposure: districts and cartelized routes reached by 2008, 2012, and 2016.

*Notes.* Cumulative districts within twenty kilometers of a cartel process let by the year shown, with the cartelized routes let by that year. *Source.* Resolución 080-2021/CLC, MTC road network.

# H Exploratory evidence from candidate asset declarations

As an exploratory exercise we ask whether the elected authorities of the exposed districts show any change in declared wealth, using the sworn asset declarations candidates file with the electoral authority (JNE). We link declarations by national identity number across the 2018 and 2022 municipal elections and regress the change in declared property count on the district's corridor exposure, with department fixed effects and mayors and councillors run separately.

The design is weak in ways that should be stated before the result. Only individuals who filed in both elections are observed, which is a selected-survivor sample, and reappearing in 2022 is endogenous to electoral success, investigation, and wealth itself. The 2018–2022 window substantially postdates the construction of most corridors, so it cannot speak to wealth accumulated during contracting or building. Declared property counts do not capture concealed ownership, family-held assets, or cash. And treatment varies geographically, not at the individual level, so the inference must reflect the assignment level, not the count of individuals.

With that assignment level recognized, exposure does not move the declared wealth of local authorities, but the estimate is imprecise. The property-count gradient is +0.027 for mayors and +0.302 for councillors, neither distinguishable from zero, and the minimum detectable effect is about 0.17 of a standard deviation for mayors once errors are taken at the corridor assignment level instead of treating filers as independent. That is far larger than the 0.03 a naive calculation would suggest. The honest conclusion is narrow: we detect no association between corridor exposure and changes in declared assets in this selected sample, but the design is neither sufficiently powered nor temporally aligned to test whether project rents reached local officials. We do not read this as evidence that the rent did or did not diffuse.

Table 13: Declared property change and corridor exposure, 2018–2022 (exploratory).

| | District mayors | | District councillors | |
|---|---|---|---|---|
| | Δprop | Δinc | Δprop | Δinc |
| exposure (log corridor km) | 0.0272 | -3298.2010 | 0.3024 | 16592.3619 |
| | (0.0413) | (12970.7293) | (0.2056) | (14937.6907) |
| MDE/SD (80%) | 0.029 | | 0.038 | |
| N authorities | 3228 | 3228 | 4883 | 4883 |

*Notes.* Change in declared property count, selected-survivor sample of authorities filing in both 2018 and 2022. MDE/SD at the corridor assignment level is about 0.17 for mayors. Standard errors in parentheses.
*Source.* JNE candidate sworn asset declarations, linked by national identity number.

# I Descriptive accounting paths of three publicly reporting entities

Of the cartel only three entities file financial statements with the securities regulator: the implicated parent Aenza (formerly Graña y Montero) and two concession vehicles, Norvial and H2Olmos. We report their post-2016 accounting paths as case evidence, with no counterfactual interpretation. After the 2016 revelation Aenza's return on equity fell well below its pre-revelation level, while the two listed concession vehicles show no comparable decline over the same years. We do not construct a comparison group or attribute these movements to the scandal. Return on equity responds to write-downs, changes in book equity, debt restructuring, provisions, and asset sales as much as to operating performance, so these paths are illustrative of the firms' reported accounts and not a measure of any causal effect of the corruption.