Quality Screening through Market Access in Public Infrastructure: Evidence from Rwanda

Nicolas Jimenez *

Job Market Paper

September 24, 2026

UPDATED FREQUENTLY. CLICK HERE TO ACCESS THE MOST RECENT VERSION.

Abstract

Governments often cannot contract on the quality of what they procure. When quality is non-contractible, screening may improve performance by restricting which firms can compete. Many governments screen contractors indirectly, conditioning eligibility for larger contracts on the execution of smaller ones. I study this trade-off in Rwanda's public procurement of infrastructure, focusing on one key measure of quality: project completion time. I link the near-universe of electronic procurement records for 2016–2024 to records of government payments that trace contract execution. I estimate a dynamic auction model in which low- and high-quality firms compete in first-price auctions and completed work determines access to larger contracts. High-quality firms complete work and accrue experience necessary for promotion faster than low-type firms. Therefore, high-quality firms value winning tenders more than low-quality firms, which induces them to bid more aggressively, partially offsetting their higher costs. I conduct two sets of counterfactuals. First, I modify the eligibility criteria and highlight the government's trade-off between price and quality. Second, I change the government's promotion policy to estimate the value of screening on past performance.

JEL classification: D44, C57, L74, D82, H57.

Keywords: public procurement, dynamic auctions, adverse selection, public infrastructure.

*Yale University. nicolas.jimenez@yale.edu. I am grateful to Lauren Falcao Bergquist, Nicholas Ryan, and Katja Seim for their advice, encouragement, and support. This paper has been greatly improved by comments from Anna Croley, Amit Khandelwal, Kyungho Lee, Shoki Kusaka, Rohini Pande, and participants at the Yale Development and Industrial Organization lunches. The IGC Rwanda Office, especially Christian Lippitsch and Mphatso Kumwenda, provided assistance in obtaining data access. Andres Henriquez and James Ahabyona provided excellent research and data collection assistance. OpenAI Codex and Claude Code were used to program code and edit the text for clarity. This project is the product of a collaboration with the Rwanda Public Procurement Authority (RPPA), and I thank Jimmy Christian Byukusenge, Francine Gatarayiha, and Martin Mukunzi for facilitating data access. I gratefully acknowledge funding for this project from the International Growth Centre (IGC), the Yale Economic Growth Center (EGC), and the National Science Foundation Graduate Research Fellowship Program (GRFP). All errors are my own.

1 Introduction

Public infrastructure (e.g., roads, sewage systems, electricity grids) is critical for economic development, but governments frequently struggle to ensure its timely completion: 60% of World Bank infrastructure projects are delayed by one year or more (Espinoza and Presbitero, 2022). Infrastructure projects are awarded by governments through public procurement markets and carried out by private construction firms with heterogeneous *costs* and *quality*. I focus on a specific margin of contractor quality: timely execution that meets the required technical specifications. Governments often cannot observe this quality when awarding contracts, nor enforce performance once contracts are signed.[1] In practice, therefore, quality is non-contractible. Low state capacity exacerbates this friction. Since official discretion is prone to abuse where capacity is low, screening is often implemented using coarse, rules-based proxies (Bosio et al., 2022), even though such rules may forgo information that officials could otherwise use (Decarolis et al., 2025).

One proxy that is available even to a low-capacity government is past performance. Public procurement is a repeated market in which a set of suppliers competes for contracts over time, so the record of past contracts can be used to screen contractors indirectly. When quality cannot be fully observed or contracted upon at the award stage, governments can instead condition eligibility for future – and often larger – contracts on the execution of earlier ones. These restrictions may allocate more valuable contracts to higher-quality suppliers, but they also reduce the number of eligible bidders and weaken competition. Eligibility rules therefore create a tradeoff between improving contract execution and containing procurement costs.

In this paper, I ask three questions. First, can eligibility rules that condition access to future contracts on past execution screen suppliers on persistent, unobserved quality? Second, how do the dynamic incentives these rules create—the continuation value of winning—shape equilibrium bidding and the allocation of contracts? Third, what is the cost of this screening in forgone competition, measured in procurement prices and completion delays?

I study these questions in the context of Rwanda's public procurement of infrastructure (henceforth *Works*) contracts. Public procurement is a large share of Rwanda's economy, accounting for approximately 11% of GDP in 2017/18 (World Bank, 2020). Although Works tenders make up only 13% of awarded tenders, they account for more than half of procurement value (Rwanda Public Procurement Authority, 2023). In this setting, both tenders and construction firms are assigned to a category on an ordered ladder which determine which tenders each firm is eligible to bid on. Promotion to a higher category is granted on the basis of past executed contracts. Eligibility restrictions of this kind are common, but their consequences for procurement outcomes remain poorly understood.[2] Studying these policies empirically requires observing three objects that are

[1] For example, penalties for delays are not collected when courts are slow (Coviello et al., 2018), and winning bidders may renegotiate the terms on which they won (Ryan, 2020).

[2] The World Bank's Standard Procurement Document for Works makes experience a pass/fail qualification criterion, requiring a minimum number of similar contracts "satisfactorily and substantially completed" (Section III, criterion 4.2(a)), and EU Directive 2014/24/EU permits requiring "a sufficient level of experience demonstrated by suitable references from contracts performed in the past" as a requirement for participation (Art. 58(4)). Registry systems

rarely available together: the rules linking past execution to future eligibility, firms' positions and progression along the eligibility ladder, and their performance after receiving a contract. Rwanda is particularly well suited to overcoming these measurement challenges. First, published regulations codify the criteria linking past execution to future eligibility, providing an observed institutional benchmark for the promotion policy. Second, firms hold sector-specific categories that are recorded over time, allowing me to trace each supplier's progression along the ladder. Third, government payments to firms appear in tax records, allowing me to construct a measure of post-award payment clearing—an administrative proxy for execution—beyond the point at which most procurement records end.

I rely on two novel administrative datasets: procurement and Value-Added Tax (VAT) records. The procurement records are held by the Rwanda Public Procurement Authority (RPPA) in its electronic procurement platform, and cover the near-universe of public Works tenders between June 2016 and October 2024 — approximately 2,000 tenders, with information about the advertised project, all submitted bids, bidder identities, evaluation and award outcomes, signed contracts, and amendments. I link these to firms' category request histories, which allow me to reconstruct a monthly running category for each firm and sector, and hence the set of tenders for which a firm was eligible at the date of each bid. Like most administrative procurement records, the Umucyo data ends at contract award and does not track contract completion. I therefore measure execution from tax records, which record all government payments to firms. Because payment follows certification of completed work, the resulting series traces realized contract execution. This yields a quarterly panel of government payments for 98 percent of the bidders in the procurement data, from 2013 through 2024. I link both datasets using each firm's taxpayer identification number.

Using the linked data sources, I first document three empirical facts about this setting that highlight the relevance of the entry eligibility policy and the tradeoff it imposes. First, contract execution delays are pervasive. I document that only 20 percent of firms' project portfolios are fully paid by their scheduled end, a third remain incompletely paid two years later. Second, portfolio completion speed is a persistent firm characteristic. A firm's completion speed in its first three quarters with outstanding work predicts its subsequent payment clearing: a one-standard-deviation higher predetermined completion score predicts a 10.4 percentage-point higher clearing rate in later quarters. I also show that most of the variation in project outcomes is a product of the firm, rather than of the project or the buyer: once portfolio size is held fixed, buyer and match characteristics account for little of that variation in completion speed, while firm fixed effects raise the explained share from 19 to 43 percent. Third, although faster completion is not used to award contracts directly, the category ladder converts completed experience into access to larger projects. Contracts are awarded via a first-price sealed-bid auction among eligible bidders. Therefore, conditional on the bid rank, no firm characteristics – a firm's past experience and completion speed – matter for their contract award. But firms' completed contracts and completion speed matter for their likelihood

closer to Rwanda's—a persistent class that caps the value of work a firm may bid for, upgraded on past executed contracts— exist in Kenya, Tanzania, Zambia, South Africa, Singapore, Spain, and the Philippines.

of category promotion. I show the positive and increasing relationship between firms' completed contract value, measured by payments data, and their likelihood of upgrading. Promotion in turn expands access: the median tender in the top category is more than one hundred times larger than in the bottom category, and, within the same firm, sector, and fiscal year, tenders pursued after promotion are 73 percent larger.

To interpret these facts and quantify the tradeoff, I specify and estimate a dynamic auction model that maps closely to this setting (Jofre-Bonet and Pesendorfer, 2003; De Silva et al., 2023). In the model, construction firms differ in a permanent, privately known execution type, carry a backlog of unfinished work, and bid for tenders whose eligibility depends on their category. A firm's type governs the rate at which it completes awarded work, and therefore the speed at which the value it wins becomes the paid experience. Each firm carries an individual state: the category that determines which tenders it may enter, the unfinished work (backlog), and the paid experience it has accumulated (which matters for promotion). In each quarter a firm observes the tenders for which it is eligible, may participate in one or more. For each tender it participates in, the firm submits a bid (for each) knowing its own cost draw and its rivals' public states, but not their types. Winning therefore imposes a tradeoff: it delivers current revenue and, in the future, may move the firm up the category ladder, but it also adds unfinished work that may raise its costs and lowers its participation in future auctions until that work clears. The firm's optimal bid balances the two, and takes the familiar form of a construction cost plus a static markup, less the discounted gain in continuation value.

Estimation proceeds sequentially, following the tradition of two-step methods of dynamic structural models (Hotz and Miller, 1993; Bajari et al., 2007; Aguirregabiria and Mira, 2007). First, I estimate the policy environment, in particular the RPPA's category promotion policy, based on the empirical relationship between accrued paid experience and category promotion. I treat the RPPA's promotion rule as an exogenous policy rule that firms account for when forecasting how winning an auction changes their eligibility for future tenders (Narayanan, 2025). Not all firms participate in all auctions they are eligible for; I estimate a descriptive firm participation policy that matches firms to auctions based on firm- and project-level observable characteristics. I assume there are two firm types, *High* and *Low*, and that High-type firms clear their backlog faster. I then recover firms' execution type as a finite mixture using firm's quarterly backlog clearing rate, participation in auctions, and bids. Identification of the mixture relies on a conditional independence assumption on the joint distribution of its backlog clearance, participation, and bidding given its type, as well as sufficient cross-sectional and panel variation in these three margins (Hall and Zhou, 2003; Kasahara and Shimotsu, 2009). I estimate the mixture and the type-specific policies jointly by Expectation–Maximization, following Arcidiacono and Miller (2011), who first applied the algorithm to dynamic models with persistent unobserved heterogeneity.

The second stage recovers the primitives that govern firms' bidding behavior. At the auction stage, firms observe rivals' identities, but not rivals' types. They form beliefs about rivals conditional on public states by updating a prior of the (category-specific) rival types using the information from observing the realized set of entrants. [3] Using these beliefs, I construct hazards for each bidder in each auction, which summarize the level of competition in that auction and determine the static markup. Given the beliefs and the estimated policies from the first stage, I compute the continuation value of winning from firms' conditional choice probabilities (Hotz and Miller, 1993; De Silva et al., 2023). I invert the bid first-order condition to decompose each bid into cost, static markup, and dynamic value (Guerre et al., 2000). Identification of cost parameters comes from the difference between the observed bid and the separately-identified static and dynamic markup components (Jofre-Bonet and Pesendorfer, 2003). This estimation stage yields structural cost primitives, static markups, and dynamic markups/markdowns by firm type.

I estimate that around 20% of firms are high-quality type. At the median backlog, high-type firms complete 42% of their outstanding project work per quarter, compared with 22% for low-type firms. Because high-type firms convert awards into completed work, and thus into access to larger contracts, faster than low-type firms, they value winning more. This larger continuation value leads them to bid more aggressively, which offsets most of their cost disadvantage.

I first use the model to quantify how the dynamic continuation value shapes firms' bids and the allocation of contracts. Promoting firms at random within their category, independently of completed work, removes most of this incentive and raises procurement prices by 8.2 percent. Thus, the current regime is partly self-financing: the government pays lower prices today precisely because the prospect of future market access makes firms mark down their bids.

I then use the model to evaluate a relaxation of the eligibility restriction. Allowing firms one category below a tender's requirement to bid on the smaller half of projects reduces procurement spending by 2 percent without significantly increasing delay. Extending the same relaxation to all projects reduces spending by 6 percent but adds 0.2 quarters, or about 2.3 weeks, of expected delay.

Section 2 describes Rwanda's public works procurement market, the categorization policy, and the data, including the construction of the payment-based measure of contract execution from VAT withholdings. Section 3 documents the three empirical facts. Section 4 develops the dynamic auction model, Section 5 describes estimation, and Section 6 presents identification. Section 7 reports the structural estimates, and Section 8 the counterfactuals: the value of the dynamic incentive and relaxations of the eligibility rule. Section 9 concludes.

1.1 Literature Review

This paper primarily contributes to three strands of literature.

The first is the literature on the potential tradeoff between prices and quality outcomes in procurement. A large literature documents the tradeoff between procurement prices and execution quality that this screening is meant to resolve (Decarolis, 2014; Bajari et al., 2014; Lewis and Bajari, 2011; Carril et al., 2026). The closest paper to mine is Andreyanov et al. (2024), who show that scoring vendors on audited past performance can raise realized quality sharply without raising

[3] This restriction on beliefs keeps the model tractable but rules out learning of others' types over time, and is similar to that in Backus and Lewis (2025).

prices using a static auction model and comparing a scoring regime to a first-price auction regime. Relative to this paper and the rest of the literature, I use a dynamic auction model that embeds the government's quality screen as a participation restriction to study how dynamic incentives for future market access shape firm participation and bidding decisions.

The second strand studies procurement when quality is imperfectly contractible. Theory shows why quality is under-supplied if it cannot be contracted. Incomplete contracts make adaptation costly (Bajari and Tadelis, 2001), providers shade non-contractible quality (Hart et al., 1997), and a buyer who values quality faces a lemons problem at the award stage, so that optimal procurement restricts price competition (Manelli and Vincent, 1995; Lopomo et al., 2023). Several empirical papers have provided clear examples: weak enforcement leaves delay penalties unenforced (Coviello et al., 2018), allows firms to win by bidding below cost and renegotiating (Ryan, 2020), exposes sellers to buyer hold-up (Ryan, 2026), and makes discretion in awards a trade-off with corruption (Decarolis et al., 2025). Relative to this work, I study a rules-based substitute for enforcement, in which the government makes access to future, larger contracts depend on firm performance. I quantify the value of that future access to firms of each type and the price the government pays for it through reduced competition.

The third is the literature on the structural estimation of dynamic auctions in public works procurement. The seminal paper by Jofre-Bonet and Pesendorfer (2003) estimates the first dynamic auction game in which capacity and backlog link current bids to a dynamic state. A small but growing literature has since studied other aspects of dynamic procurement auctions, including participation (Groeger, 2014), unobserved auction heterogeneity (Balat, 2015), endogenous asymmetry between bidders (Saini, 2012), and subcontracting (Jeziorski and Krasnokutskaya, 2016). The closest paper is De Silva et al. (2023), which develops a dynamic oligopoly model with backlog, entry, exit, and experience accumulation in large markets. I use the discrete analogue of their framework, and extend it in two main ways. First, I introduce Rwanda's category ladder as the state that determines tender eligibility and thus governs firms' continuation value. Second, I use an Expectation–Maximization algorithm to estimate persistent latent firm types, following Arcidiacono and Miller (2011).[4] Unlike the papers in this literature, my government payments data allows me to measure backlog and project completion directly over time, rather than impute it using contract schedules.

2 Setting and Data

2.1 Public Works Procurement in Rwanda

Public procurement is the primary instrument of infrastructure provision in Rwanda. Public Works tenders[5] are administered by procuring entities (PEs)—ministries, districts, hospitals, schools, and

[4]Outside procurement, recent papers that use similar approaches to estimate unobserved permanent heterogeneity include Cui (2025) and Narayanan (2025).

[5]A *tender* is an auctioned project (e.g., the paving of a highway segment). A tender may be subdivided into *lots*, which are advertised together but bid on and awarded separately; the lot is therefore the unit of competition, and I treat each lot as an independent auction throughout. Because 85 percent of tenders contain a single lot, the text

other public bodies—each of which advertises and administers its own tenders under the regulation of the Rwanda Public Procurement Authority (RPPA). The suppliers are construction firms, the overwhelming majority of which (more than 95% in the bidding data described below) are registered as formal firms in Rwanda.[6] Throughout the paper I focus on the three main Works sectors – Buildings, Roads and Bridges, and Drinking Water Supply –, which constitute 90% of all tenders in my sample.

A Works tender proceeds through a sequence of steps recorded in my data. The PE advertises the tender—divided, where appropriate, into lots—on *Umucyo*, Rwanda's electronic procurement platform, which the RPPA launched nationally in 2017. Umucyo covers roughly 85–90 percent of public tenders; the remainder are tendered offline and are not observed in my data. Firms submit sealed bids, alone or as part of a joint venture (JV) consisting of several firms. Bid evaluation proceeds in two stages: first, bids that do not meet administrative and technical responsiveness criteria are disqualified. Second, among the qualified bids, the contract is then awarded to the lowest evaluated price. After award, the winning firm signs a contract specifying a final contract value and a start and end date. Subsequent changes to value or duration are recorded as contract amendments. During execution, the firm files an invoice upon completing each portion of the works, the PE certifies the invoice, and payment is released. At the moment of payment, public institutions withhold Value-Added Tax (VAT). This institutional feature underlies my measurement of contract execution, and I return to it below. The RPPA additionally maintains a public blacklist of firms which commit serious breaches of contract and are fined or barred from future tenders.

Three features make this setting well suited to studying quality screening. First, electronic procurement records track the near-universe of Works tenders from advertisement through amendment. Second, tax records can be used to measure when firms actually paid for Works contracts. Third, experience-based eligibility rules that govern market access are public, and can be estimated using the first and second components. I now turn to these rules, which create the supplier categorization system.

Supplier Categorization System Since 2014, public procurement of infrastructure (Public Works) in Rwanda uses a formal category system to determine firm eligibility for each tender. In conversations, high-level RPPA officials described two objectives for the categorization policies. First, to mitigate delays and project abandonment of large and complex projects by firms without the capacity to complete them. Second, to reserve smaller projects for small and medium firms to grow and gain experience. The advent of electronic procurement in 2016–17 in Rwanda facilitated the widespread adoption and accurate application of the categorization system by every PE, as well as stringent enforcement of firm eligibility criteria. The system is generally seen as a success within the RPPA, although questions remain about the degree to which it encourages firm growth due to

refers to auctioned projects as tenders rather than lots when the distinction is immaterial.

[6]This share includes subsidiaries of multinational firms owned (or partially owned) by foreign firms. Unfortunately, my data does not include ownership information, so I am unable to accurately distinguish between foreign and Rwandan firms.

low rates of firm promotion.

The categorization system classifies tenders and firms. When advertised by the commissioning PE, each tender is assigned to an ordered category ($A > B > C > D > E > F$) based on the project's engineering cost estimate (ECE). Given an ECE, the classification is automatic through the electronic system, unless in rare occasions where the PE official requests an alternative categorization.[7] In these cases, the PE must submit justification their alternative categorization and the RPPA must approve it. On the other side of the market, firms are also classified in categories (also A to F). When registering on the Umucyo system to compete for public procurement tenders in a given sector (e.g., Buildings), a firm can request to be categorized in A, B, C, D, or E, or will be placed in category F by default. The RPPA's decision to approve firms' categorization requests is based on four criteria for which the firm must submit documentation: 5-year previous revenue, number and value of previously executed public contracts, owned equipment, and permanent professional staff hired at the firm. Moving up the category ladder requires meeting higher thresholds on every margin. The requirements for a given category are broadly similar across sectors. Appendix Table 6 reports the turnover and past executed-tender requirements for the three sectors studied in the paper. During the sample period, categories D and E required one executed tender, while categories A–C required two, each worth at least 60 percent of the prospective tender value (Rwanda Public Procurement Authority, 2014). [8]

The category system determines bidding eligibility. During the sample period studied in this paper (2016–2023), a firm could bid for tenders in its own category and up to two categories below. For example, a category B firm could bid for projects in categories B, C, and D. The eligible set for a lot is determined by the lot category and each firm's registered category at the time of bidding.[9]

Eligibility is sector-specific, and firms are categorized separately for each sector they register in. A firm may hold one category in Buildings and a different category in Roads and Bridges.

2.2 Data

I use administrative records from two institutions. From the RPPA, I use the records of its electronic procurement platform, *Umucyo*, for the period September 2016 – October 2024: the procurement records, covering advertisement, bid submission, evaluation, award, contract signing, and contract amendments, and the categorization request records. From the Rwanda Revenue Authority (RRA), I use VAT declaration data covering 2013–2024, from which I recover a firm-quarter measure of government payments, as described below.

The two sources are linked by firms' Taxpayer Identification Numbers (TIN). The procurement records define the auctions, eligible firms, signed commitments, and contract schedules; the VAT

[7]For instance, if the project is more complex than the ECE suggests. Miscategorization in either direction occurs for 3.9% of projects in our sample.

[8]In Roads and Bridges, categories A and B are subdivided into $A1/A2$ and $B1/B2$, respectively; I collapse these into A and B for simplicity.

[9]In October 2023 the cross-category rule was tightened. Firms could continue to bid for tenders in their own category, but to bid in lower-category tenders, they must joint venture with a firm in the tender's category.

records define the realized payment timeline. I supplement these with Rwanda's Open Contracting Data Standard (OCDS) records for 2016–2024, which identify public contracts, supplier IDs, contract dates, and contract values after the tendering process. I use the OCDS records to identify the non-Works (Goods, Services, and Consulting) contracts held by the Works firms in my sample and to net the associated payment flow out of the VAT measure of Works execution.

2.3 RPPA Procurement Records

Between June 2016 and October 2024, Umucyo records 2,211 Works tenders (2,493 tender-lots). The RPPA Umucyo data represents the near-universe of procured Public Works tenders during the sample period. For each tender, Umucyo records the tender notice, tender lots, submitted bids, bidder identities, technical evaluation records, selected winners, signed contracts, and amendments. The three sectors named above collectively make up 90% of projects: Buildings (60%), Roads and Bridges (16%), and Drinking Water Supply (12%). I restrict the sample to projects awarded through low-price sealed-bid auctions (National or International Competitive Bidding), consisting of 85% of lots. After these sample restrictions, I retain valid bids submitted by lead bidders and joint-venture partners. A contract is signed after about 90 percent of auctions.[10]

From the raw procurement records I construct a tender-lot-bid panel. For each tender-lot, the panel links project characteristics (sector, engineering estimate cost), the procuring entity, bidder (anonymized) TINs, bid prices, joint venture members, evaluation outcomes (disqualifications), selected winner, signed contracts, contract value, contract start and end dates, and contract amendments. To this panel, I link the bidder's registered category at the time of bid submission, as described below.

Firm Category Histories I construct firm-sector category histories from firms' category requests submitted to the RPPA. The records identify the firm TIN, request and response dates, sector, requested category, and approval decision. I convert approved requests into a monthly running category for each firm-sector, which I use to assign bidder eligibility at the bid date and to measure category upgrades. Appendix Table 7 shows that category eligibility requirements are followed during my sample period: 99.5 percent of bids come from firms in the tender category or no more than two categories above it.

2.4 RRA VAT Declarations and Government Payments

While the RPPA data contains detailed information from tender advertisement to tender award, it does not contain detailed information about contract execution. Because payment follows the certified-invoice process described in Section 2.1, government payments trace realized execution. To measure the speed of contract execution, I therefore use data on government payments to firms from Value-Added Tax (VAT) declarations from 2013–2024. Public institutions retain 18 percent

[10]Some tenders are closed without a winner being selected. In a small number of cases, a winner is selected but no contract is signed.

VAT on taxable supplies under public tenders, and the retention is automatically recorded when firms file VAT declarations. I measure government payments to firms from the VAT retained by public institutions (Rwanda Revenue Authority, 2025), and construct VAT-implied government payments for firm i in quarter t as

$$Y_{it}^{\text{tot}} = \frac{\text{VAT retained by public institutions}_{it}}{0.18}.$$

Firms file VAT monthly or quarterly, depending on firm size; I aggregate monthly filings to the quarterly payments. I am able to match 98% of unique bidders in the procurement dataset to the VAT data. This yields a firm-quarter panel of government payments from all sources.

2.5 From Payments to Execution and Experience

Allocating payments to the Works portfolio. The VAT payment data record the total payments across all contracts in each time period; they do not identify the contract associated with each payment. It is common for firms to have multiple ongoing contracts: in more than half of firm-quarter pairs in which a firm has an ongoing Works contract, the firm also has another Works or non-Works contract. I therefore combine the VAT data with the timing and value of observed Works contracts from Umucyo to construct a firm-level measure of payments attributable to the Works portfolio.

I describe this process in detail in Appendix A.2. In summary, I construct the total amount of government payments by firm-quarter from all sources and subtract from it the payment amount that cannot be attributed to Works contracts. To measure this share unattributable to Works, I use the value of non-Works contracts from OCDS records, and impute the value of offline tenders that I do not observe in my data.[11] A firm's open *portfolio* in a quarter is the value of Works contracts it has not yet been paid for, i.e., its *backlog* carried over from the previous quarter plus the value of new Works contracts awarded to it that quarter. Payment for Works contracts is defined as the full remaining amount, up to the value of the firm's open portfolio. In each period, the value of Works payments reduces the portfolio; whatever remains unpaid carries over as the next quarter's backlog.

Among the 4,674 firm-quarters with open portfolios, Works payments are positive in 44 percent. Conditional on a positive payment, the median is RWF 32 million and the mean is RWF 225 million.

Portfolio execution and experience. For firm-quarters with an open portfolio, I measure quarterly portfolio *clearance* as Works payments received in the quarter divided by the open portfolio. It ranges from zero, when none of the portfolio is paid that quarter, to one, when all of it is. As firms clear their portfolios they accumulate *paid experience*, the cumulative value of Works payments they have received. Firms' promotion up the category ladder depends on the total value of paid

[11]The value of unaccounted payments changes measured Works payments in only 15.6% of quarters when firms have active Works contracts. Moreover, it affects relatively few firms: setting unaccounted payments to zero would increase attributed Works payments by RWF 106.5 billion, but five firms account for 81% of this difference.

experience. Across firm-quarters with an open portfolio, paid experience has a median of RWF 37 million and a mean of RWF 753 million; over the panel, firms accumulate RWF 462 billion of it.

2.6 Validation

I validate the measure of Y^W as payment completion using independently collected project reports submitted by third party firms hired to supervise the contract execution and approved by the procuring entity. A field team collected and coded available project reports and handover documents for 196 tenders. Among these, 124 contracts contained a dated final report in the relevant sectors: 67 are Buildings, 22 are Roads and Bridges, and 35 are Drinking Water Supply. While this sample is selected based on report availability – projects are larger and recent – they provide an independent measure of project completion and delays that is useful to validate the main measure using VAT payments data.

Completion and payment timing. VAT payments are positively associated with project completion. Among 108 contracts for which both reported completion and contemporaneous payment clearing are observed, a 10 percentage point increase in reported physical completion is associated with a statistically significant 5.1 percentage point increase in cumulative Y^W, conditional on contract value, scheduled duration, sector fixed effects, and report-year fixed effects (Appendix Table 5).

For a small sample of 11 contracts, the timing of the firm's other contracts allows us to attribute payments directly to the focal project. Median cumulative payment rises from 24 percent of contract value in the final-report quarter to 100 percent one quarter later, when 6 of the 11 projects are fully paid. By then, nine have received at least 70% of the value. Although the sample is small and potentially selected, this pattern suggests that payments follow documented project completion relatively quickly (Appendix Figure 12).

Delays. I measure delays by comparing the report date with the contract's scheduled end date. Among 121 contracts for which we collect report dates, delayed contracts are 13.3 percentage points more likely not to have reached 90 percent payment one quarter after the scheduled end date (standard error 9.5 percentage points). This coefficient is positive, as expected, but not statistically significant (Appendix Table 5). Taken together, the results support interpreting Y^W as a measure of payment clearing that is associated with physical progress but follows documented completion with a lag.

3 Empirical Evidence

This section presents three facts about project execution and market access in Rwanda's infrastructure public procurement market. First, payments on public Works contracts frequently lag behind their scheduled timeline, and some contracts remain unpaid well after their scheduled end. Second, portfolio clearing speed differs persistently across firms: performance in a firm's first three observed quarters predicts its subsequent payment clearing. Third, the category ladder converts experience, i.e., completed projects, into access to larger tenders. Paid experience predicts promotion, higher categories contain larger projects, and firms win larger contracts after promotion. Taken together, these facts motivate a model with persistent firm types, backlog, accumulated paid experience, and category-dependent market access.

3.1 Fact 1: contract execution delays are pervasive

To measure delays, I use data on the realized contract payment amount and timing (from the VAT data) and the scheduled payment end date (from the procurement contract data). In Panel A of Figure 1, I plot the scheduled payment timeline assuming payments were equally spaced over the contract duration, and compare it to the realized payment timeline. I do this for both the portfolio and the project-attributable samples, which yield nearly identical results. The graph illustrates that, by the scheduled contract end date, firms have received approximately 54 percent of contract value, 74 percent two quarters later, and 82 percent after eight quarters.

However, the average payment timeline hides significant delays and project non-completion. Panel B reports the share of contracts or portfolios that have reached full payment. At the scheduled end, only 22 percent of project-attributable contracts, and 20 percent of portfolios, are fully paid. The shares rise to 50 and 47 percent after two quarters, and to 65 and 64 percent after eight quarters. Almost one-third of portfolios are not fully paid, highlighting that delays are not driven only by small unpaid balances on otherwise completed portfolios. These facts suggest that public procurement contracts in Rwanda suffer from substantial delays and non-completion.

Figure 1: Payment clearing around the scheduled end

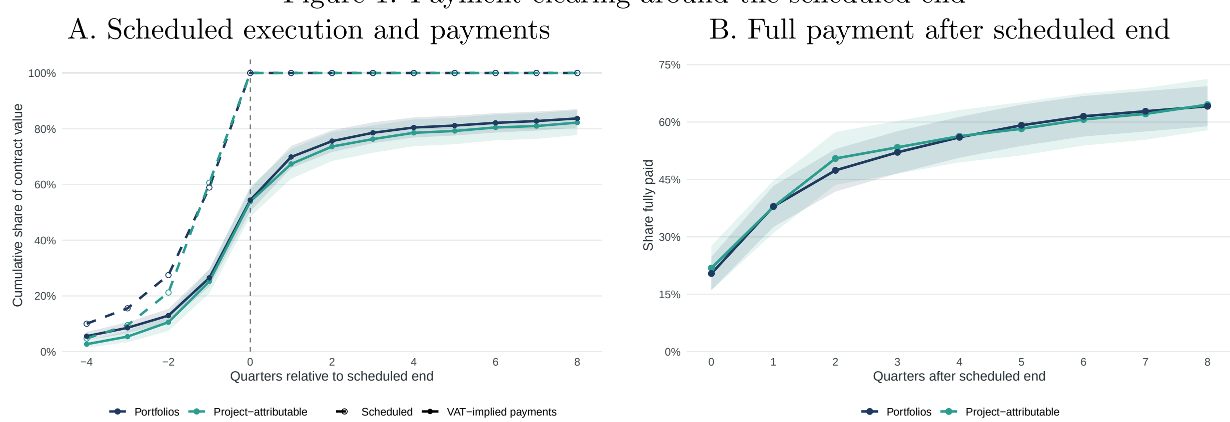

Notes: Panel A reports cumulative scheduled execution and VAT-implied payments relative to the scheduled end. Panel B reports the share of contracts or portfolio episodes that reach full payment. The sample contains 206 project-attributable contracts and 382 portfolio episodes. Shaded bands are 95 percent confidence intervals clustered by firm.

3.2 Fact 2: portfolio completion speed is a persistent firm characteristic

I next ask whether portfolio completion speed contains a persistent firm component by testing whether firms' completion rate in its first few periods predict its later outcomes. For each firm, I construct a completion score from its first three quarters with outstanding work (i.e., positive backlog). I first residualize the quarterly backlog clearing rate on the value of work available for payment, new awards, the number of active contracts, and calendar quarter fixed effects. I then average these residuals within firm, shrink the average toward the sample mean, and standardize it. Panel A of Figure 2 shows that the score has wide dispersion across the firms that win at least one contract.

Panel B shows that the (predetermined) completion score predicts several future clearing outcomes relatively far into the future. A one-standard-deviation higher predetermined score predicts a 10.4 percentage-point higher clearing rate in later quarters. The estimate remains 8.7 percentage points at least four quarters after the score is fixed, and the probability of positive clearing is 7.8 points higher. In the project-attributable sample, the same score predicts a 7.4-point higher project clearing share and an 11.7-point higher probability of full payment by eight quarters after the scheduled end. The persistence across time and across subsequent contracts is suggestive of heterogeneity in firm-level execution quality.

I provide further evidence that execution speed is largely a firm-specific attribute by testing where variation in quarterly portfolio clearing comes from. In Appendix Table 9, I report results of regressions of firm-quarter-level backlog clearing rates, where I progressively add controls to assess the change in adjusted R^2. Calendar quarter fixed effects alone produce an adjusted R^2 of 0.034. Adding the size and composition of the firm's active portfolio raises it to 0.145. Procuring entity composition and the procuring entities' prior completion performance only raise the adjusted R^2 only to 0.149. Observable measures of firm size and experience raise it to 0.193, while firm fixed effects raise it to 0.431. Thus, workload and observable firm characteristics explain some clearing variation, but buyer characteristics explain little; the remaining variation contains a substantial firm-specific component.

3.3 Fact 3: the category ladder converts completed experience into access

Because contracts are awarded through first-price auctions, execution delays and project non-completion are not used directly in bid evaluation and award. Appendix Table 10 shows that the portfolio clearing component does not predict bid disqualification. It also does not predict award conditional on qualification and bid rank.

However, project completion is crucial for determining a firm's future eligibility through the category ladder. As described above, a firm's category determines the size of tenders for which a firm is eligible.[12] Promotion therefore changes the set of projects available to the firm in the future. To examine how completed experience translates into future eligibility, Figure 3 plots next-quarter

[12] The restriction binds in the data: 99.5 percent of classified bids in Buildings, Roads and Bridges, and Drinking Water Supply satisfy the rule that firms bid in their own category or up to two categories below it (Appendix Table 7).

Figure 2: Predetermined clearing and subsequent execution

A. Predetermined clearing component B. Subsequent execution

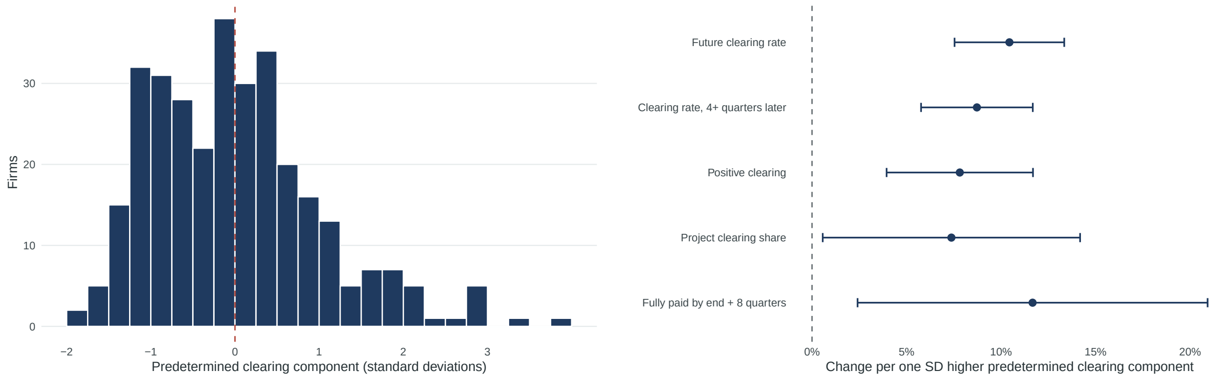

Notes: Panel A plots the standardized, empirical-Bayes-shrunk mean residual clearing rate in each firm's first three positive-backlog quarters. The sample contains 319 firms; the dashed line marks the mean. Panel B reports the change in subsequent outcomes associated with a one-standard-deviation higher score. Bars are 95 percent confidence intervals clustered by firm.

promotion against cumulative paid Works experience, P_{it}, measured relative to the category- and sector-specific value requirement. Promotion rises smoothly with paid experience. This pattern is consistent with left-censored pre-Umucyo and offline experience, as well as other statutory promotion requirements.

I provide further evidence of the value of experience for upgrading by comparing observably similar firms that apply for a category promotion in the same quarter. Within request-quarter-by-sector-by-origin-by-target cells, and controlling for firm age, prior Works contracts, and prior category requests, applicants whose measured paid Works experience meets the stated value requirement are 10.3 percentage points more likely to be approved (Appendix Table 8).

Promotion expands access to larger tenders. Panel A of Figure 4 plots engineering estimates across the ladder. The median tender rises from RWF 30 million in category F to RWF 3.26 billion in category A. The increase is monotone across categories and is two orders of magnitude larger at the top than at the bottom of the ladder.

Firms use the additional access from category promotion. Within the same firm, Works sector, and fiscal year, tenders pursued after an observed category upgrade are 73 percent larger, and winning tenders are 80 percent larger. Panel B plots the corresponding coefficients. Project size is stable over the quarters immediately before promotion and rises when the firm is first observed in the higher category.

Accordingly, category progression is mostly upward: 20.9 percent of the 465 firm-sector pairs in the sample of regular bidders move to a higher category, 78.3 percent remain in the same category, and only 0.9 percent move down (Appendix Figure 14).

Figure 3: Paid experience and category promotion

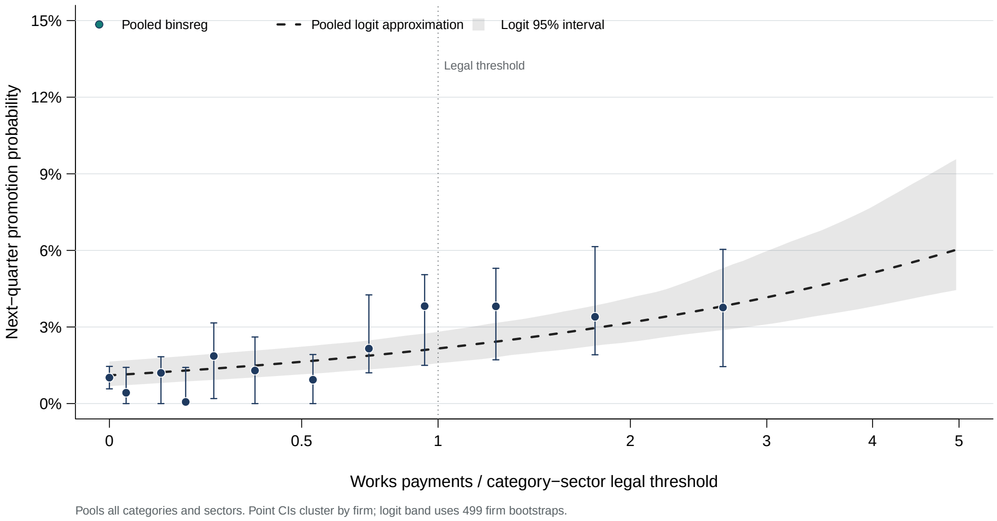

Notes: Points are quantile-spaced bins controlling for Works sector; bars are 95 percent confidence intervals clustered by firm. The dashed line is a pooled logit fit with a firm-bootstrap confidence interval. The sample contains 5,081 firm-quarter transitions for 357 firms in categories F through B, including 96 promotions. The vertical line marks the stated payment-value requirement.

Figure 4: Category promotion and access to larger projects

A. Tender value across categories

B. Project size around promotion

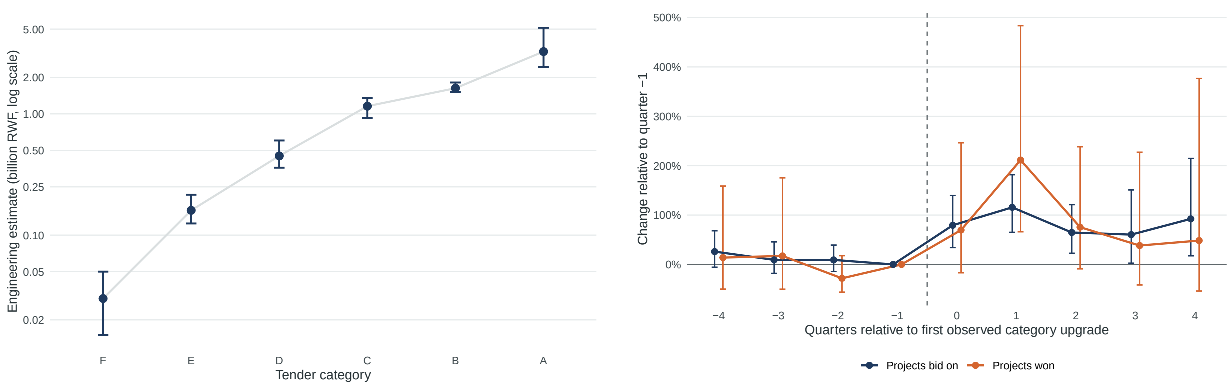

Notes: Panel A reports median engineering estimates and interquartile ranges for 1,710 competitive Works lots before October 2023; values are in billion RWF on a log scale. Panel B reports coefficients relative to quarter -1 from regressions of log engineering estimates on event time, absorbing firm-by-sector-by-fiscal-year fixed effects. Bars are 95 percent confidence intervals clustered by firm.

3.4 Mapping the evidence to the structural model

The three facts describe the RPPA's indirect screening mechanism. Firms differ persistently in their execution speed. While execution speed is not used directly in awarding contracts, it does determine how paid experience accumulates. Paid experience then translates into category promotion, which opens access to larger tenders and potentially increases firm revenue:

firm type $\longrightarrow$ backlog clearing / project completion $\longrightarrow$ paid experience $\longrightarrow$ category promotion $\longrightarrow$ access to lar

Therefore, the category system can select firms with faster project execution into larger markets without requiring the government to observe their type directly.

The empirical evidence motivates three features of the model. Firms carry a stock of backlog, which rises when they win contracts and falls as payments clear. Firms differ in persistent execution speed, represented by type θ_i. As firms clear unfinished work, they accumulate paid experience P_{it}, which affects their probability of upgrading categories and hence the tenders they may participate in.

These facts describe how firms progress through the category ladder, but do not show how the ladder affects bidding. In particular, the descriptive evidence cannot separate construction costs from static markups and the continuation value of future market access. I use the structural model to make this decomposition and to quantify the tradeoff between screening firms through eligibility and increasing procurement prices by restricting competition.

4 Model

4.1 Overview

The model is based on the cannonical dynamic auction model of Jofre-Bonet and Pesendorfer (2003), and its moment-based formulation in De Silva et al. (2023). Firms participate repeatedly in first-price procurement auctions, winning adds backlog, and bids reflect the continuation value of an award. I extend these models to capture key features of my setting: I add unobserved and persistent differences in firm execution following Arcidiacono and Miller (2011), and a category ladder that conditions future tender eligibility on past paid experience. These features allow me to quantify how the category ladder screens execution types against the loss of competition from restricting eligibility.

In each period, a firm's category determines the tenders for which it is eligible. Firms differ in a permanent private type $\theta_i \in \{H, L\}$, which governs their speed of project completion. As backlog clears, firms accumulate paid experience and may move up the category ladder. Execution is stochastic conditional on type and state; firms do not choose execution effort. The RPPA category promotion policy and firms' participation policies are exogenous and assumed to be known by all firms.

Conditional on participation, firms draw private costs and optimally choose their bids. When choosing their bid, firms trade off current-period profits, which depend on their costs and competition in that auction, with the net continuation value of winning, which enters the bid as a dynamic markdown. Winning a contract has opposing dynamic effects: it adds backlog, but completing a project expands the firm's future market access. High-type firms complete projects and accrue experience faster than Low-type firms, so, in general, they have larger dynamic value of winning.

Firms know their own type, but the RPPA and rival firms do not observe it. In the each auction, firms observe the identities of participating rivals and each rival's public state. They form beliefs about rival types by integrating over the type distribution in each category, rivals' public states, and the participation policy. When forecasting future competition, firms do not track the full distribution of rival states. Instead, each firm tracks its own public state and a set of aggregate moments that capture competition in each of the categories, as in De Silva et al. (2023).

4.2 Environment

4.2.1 Firms, Tenders, and Eligibility

Time is discrete and measured in quarters, $t = 0, 1, \ldots$. The set of active firms in quarter t is $\mathcal{I}_t$. Public works are divided into sectors $m \in \mathcal{M}$, where $\mathcal{M} = \{\text{Buildings}, \text{Roads}, \text{Water}\}$. Registration categories are

$$\mathcal{K} = \{F, E, D, C, B, A\},$$

ordered from F (lowest) to A (highest); define $R(A) = 1, \ldots, R(F) = 6$, so higher categories have lower ranks. Registration is sector-specific. Firm i's category portfolio is

$$\boldsymbol{\kappa}_{it} = (\kappa_{im,t} : m \in \mathcal{M}),$$

where $\varnothing$ denotes an inactive sector. Observed entry into a new sector is included in the estimated category transition; Appendix C.3 describes its treatment.

Tender-lots, henceforth projects, are indexed by ℓ. A project arriving in quarter t has public characteristics

$$\chi_\ell = (m_\ell, k_\ell, z_\ell, q_\ell),$$

where m_ℓ is the sector, k_ℓ the tender's category, z_ℓ the engineering cost estimate (ECE), and q_ℓ other tender characteristics.

Projects arrive exogenously. In each quarter, the number of new projects in tender category k is a Poisson draw with intensity λ_k, and each arriving project's characteristics χ_ℓ are drawn independently from a category-specific distribution G_k^χ, i.i.d. across projects and quarters and independent of firm states and of the aggregate state.

For a project ℓ in sector $m = m_\ell$, firm i is eligible if it is active in m and the project lies in its category or no more than two categories below it:

$$Q_{i\ell t} = \mathbf{1}\{\kappa_{im,t} \neq \varnothing,\ 0 \leq R(k_\ell) - R(\kappa_{im,t}) \leq 2\}. \tag{1}$$

The realized set of firms eligible for project ℓ is denoted $\mathcal{E}_{\ell t} = \{i \in \mathcal{I}_t : Q_{i\ell t} = 1\}$.

4.2.2 Individual and Aggregate States

Firm i has permanent execution type

$$\theta_i \in \{H, L\},$$

where type H has a stochastically higher quarterly clearance distribution. The realization of θ_i is private to the firm; neither its rivals nor the government observes it.

A firm's individual state is

$$s_{it} = (\boldsymbol{\kappa}_{it}, B_{it}, W_{it}), \tag{2}$$

where B_{it} is the total value of Works projects that remains to be completed, and W_{it} is cumulative awarded Works value. Paid experience, the total value paid to the firm for its completed Works projects, is

$$P_{it} = W_{it} - B_{it}. \tag{3}$$

Rivals' current states are observed and used in the current auction. To forecast future competition, firms summarize the distribution of rival states by an aggregate state $\bar{s}_t$, which captures the intensity of competition in the different categories to which they can upgrade. This moment-based state replaces the full distribution of rival states when firms evaluate future competition (De Silva et al., 2023).[13]

Each firm therefore tracks $(s_{it}, \bar{s}_t)$, conditional on its fixed type θ_i.

4.2.3 Industry Entry and Exit

The active set $\mathcal{I}_t$ evolves through exogenous entry and exit. Entrants draw initial states and permanent characteristics from an exogenous distribution that may depend on $\bar{s}_t$. An incumbent survives with probability $1 - \zeta(s_{it}, \bar{s}_t)$, which depends on its public state and the common state. Exit is absorbing and yields no current payoff, continuation value, or scrap value. The value function is measured before the survival draw.

[13]Section 5 and Appendix C.4.3 give the construction and law of motion of the aggregate state

4.3 Within-Quarter Timing

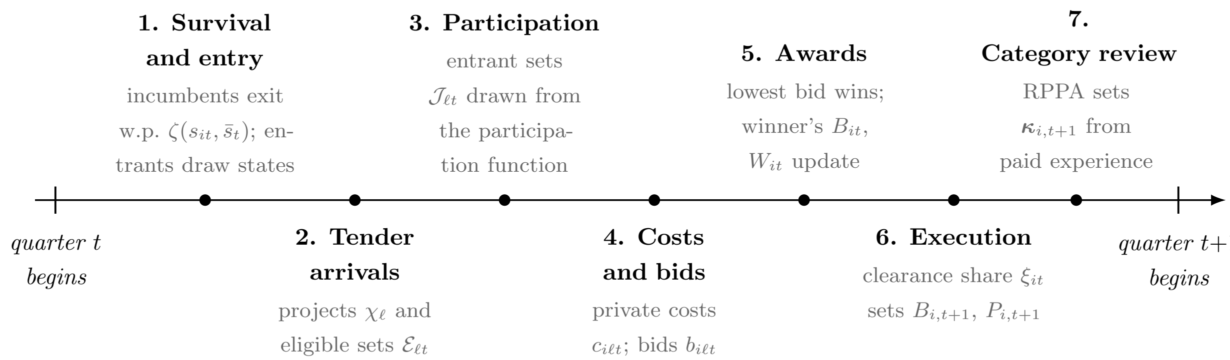

4.4 Execution, Backlog, and Payment Experience

4.4.1 Quarterly Clearance Distribution

Let

$$\Delta W_{it} = \sum_{\ell \in \mathcal{W}_{it}} z_\ell^c \tag{4}$$

be the VAT-exclusive signed-contract value of projects awarded to firm i in quarter t, where z_ℓ^c is contract ℓ's signed value net of VAT. These equations describe the realized states. When evaluating a prospective award at the bidding stage, its signed value z_ℓ^c is not yet observed. The model's win-lose transition therefore uses the project's VAT-exclusive engineering estimate z_ℓ as a proxy for the award increment.[14] The stock of work that can be cleared is

$$X_{it} = B_{it} + \Delta W_{it}. \tag{5}$$

For $X_{it} > 0$, the clearance share is

$$\xi_{it} = \frac{Y_{it}^W}{X_{it}} \in [0, 1], \tag{6}$$

where Y_{it}^W is VAT-exclusive Works payment capped at X_{it}. Quarters with $X_{it} = 0$ are outside the execution risk set.

Conditional on permanent type and post-award stock,

$$\xi_{it} \sim F_\xi(\cdot \mid X_{it}, \theta_i). \tag{7}$$

Since the clearance distribution can depend on X_{it}, the clearance share may vary with the backlog amount; so, all else equal, a firm with more backlog may clear a lower share in each period. For a common stock, type H has a stochastically higher clearance share:

$$F_\xi(x \mid X, H) \leq F_\xi(x \mid X, L) \quad \text{for all } x \in [0, 1]. \tag{8}$$

[14] Like prior papers, I rule out dependence of the backlog on the selected bid.

Execution type changes how quickly a given award becomes paid experience; high-type firms clear work faster. In the empirical implementation, I parametrize the distribution to allow for no payment, partial clearance, and full clearance in each period. Section 5 gives its parametrization.

4.4.2 State Transitions

The continuous states evolve as

$$B_{i,t+1} = (1 - \xi_{it})(B_{it} + \Delta W_{it}), \tag{9}$$

$$W_{i,t+1} = W_{it} + \Delta W_{it}. \tag{10}$$

Paid experience accrues as the firm completes new work:

$$P_{i,t+1} = W_{i,t+1} - B_{i,t+1} = P_{it} + \xi_{it}(B_{it} + \Delta W_{it}) = P_{it} + Y_{it}^{W}. \tag{11}$$

Faster execution reduces next-period backlog and increases paid experience by the same amount. Because promotion depends on paid experience, the transition links execution type to future eligibility.

The category transition occurs after quarter-t execution:

$$\boldsymbol{\kappa}_{i,t+1} \sim \Pi(\cdot \mid \boldsymbol{\kappa}_{it}, P_{i,t+1}). \tag{12}$$

Thus $\boldsymbol{\kappa}_{it}$ governs eligibility in quarter t. Work completed during the quarter first affects eligibility in quarter $t+1$. As described in Section 5.2, the kernel is estimated from observed category histories and held fixed in the firm's problem. It does not condition directly on θ_i; type affects category progression through paid experience.

4.5 Auction Participation

In each quarter, each firm may be eligible to participate in multiple tenders. At the auction stage, for each eligible firm–project pair, let $a_{i\ell t} = 1$ if firm i submits a bid for project ℓ in quarter t. Participation follows the exogenous policy

$$a_{i\ell t} \sim \text{Bernoulli}\left(a_{\theta_i}(\chi_\ell, s_{it}, \bar{s}_t)\right), \qquad i \in \mathcal{E}_{\ell t} \tag{13}$$

which conditions on project-level observables, the firm's individual state, and the aggregate state. For each project, the realized entrant set is $\mathcal{J}_{\ell t} = \{i \in \mathcal{E}_{\ell t} : a_{i\ell t} = 1\}$.

4.6 Construction Costs and Bidding

Private Construction Costs Conditional on participation, firm i observes a private project-specific construction cost distributed as

$$c_{i\ell t} \sim F_c(\cdot \mid \chi_\ell, s_{it}, \theta_i)\,. \tag{14}$$

Here $\chi_\ell = (m_\ell, k_\ell, z_\ell, q_\ell)$ and $s_{it} = (\boldsymbol{\kappa}_{it}, B_{it}, W_{it})$, with relevant category $\kappa_{im_\ell,t}$ and paid experience $P_{it} = W_{it} - B_{it}$. Conditional on $(\chi_\ell, s_{it}, \theta_i)$, cost draws are independent across firms, projects, and quarters.

Information in the Current Auction For the remainder of this subsection, set $m = m_\ell$ and let $\omega_{H,\kappa_{jm,t}}$ denote the population share of type-H firms in rival j's category κ in sector m at time t.

Assumption 1 (Auction information)**.** *In each auction ℓ, bidders compete in a first-price sealed-bid auction, where the lowest bid wins. Each firm observes its own $(\theta_i, c_{i\ell t})$. Entrants observe χ_ℓ, the entrant set $\mathcal{J}_{\ell t}$, and each rival's public state s_{jt}. Rival types and drawn project costs are unobserved. Conditional on rivals' qualification categories, firm types are independent, and category is sufficient for the pre-participation belief:*

$$\mathbb{P}(\theta_j = H \mid \kappa_{jm,t}, s_{jt}, \chi_\ell, \bar{s}_t) = \omega_{H,\kappa_{jm,t}}. \tag{15}$$

Type-specific participation and bid policies may depend on public states.

The assumption rules out learning about type from public states beyond category before participation, while allowing those states to affect behavior (Backus and Lewis, 2025).

Because the two types participate at different rates, observed participation and rival states provide information about the likelihood. Firms use Bayes' rule to update their prior belief about rival j's type, given j's participation and individual state, yielding the posterior belief $\tilde{\omega}_{j\ell t}$:

$$\tilde{\omega}_{j\ell t} = \frac{\omega_{H,\kappa_{jm,t}} a_H(\chi_\ell, s_{jt}, \bar{s}_t)}{\omega_{H,\kappa_{jm,t}} a_H(\chi_\ell, s_{jt}, \bar{s}_t) + (1 - \omega_{H,\kappa_{jm,t}}) a_L(\chi_\ell, s_{jt}, \bar{s}_t)}. \tag{16}$$

Rival distributions. Using the observable information about rival bidders' states and their belief about rival types as in Equation (16), each bidder i forms beliefs about bids that each rival submits. Let $F_{\theta,j}$, $f_{\theta,j}$, and $\bar{F}_{\theta,j} = 1 - F_{\theta,j}$ denote the equilibrium bid distribution, density, and survival function of a type-θ rival j, conditional on χ_ℓ, its observable state s_{jt}, and the aggregate state; I suppress these arguments below.[15]

From i's perspective, rival j's bid has the mixture distribution

$$F_j^{\text{mix}}(b) = \tilde{\omega}_{j\ell t} F_{H,j}(b) + (1 - \tilde{\omega}_{j\ell t}) F_{L,j}(b), \tag{17}$$

[15]The survival function is the probability that j bids above b; i.e., the probability that i beats j with a bid b.

and since this is linear in the distribution function, the density and survival function carry the same weights:

$$\begin{aligned} f_j^{\text{mix}}(b) &= \tilde{\omega}_{j\ell t} f_{H,j}(b) + (1-\tilde{\omega}_{j\ell t}) f_{L,j}(b), \\ \bar{F}_j^{\text{mix}}(b) &= \tilde{\omega}_{j\ell t} \bar{F}_{H,j}(b) + (1-\tilde{\omega}_{j\ell t}) \bar{F}_{L,j}(b). \end{aligned} \tag{18}$$

This mixed distribution governs i's chance of winning, as described below.

Single-bidder auctions. In my setting, 39% of auctions have a single bidder. Similar to a government reserve, I introduce a simulated rival firm that competes with the single bidder on these auctions. This rival is never awarded the contract, but its presence enters the win probability and the hazard as though it could. The simulated rival submits a bid based on a cost that is sampled from the distribution of firm-states in that category. While strong, this assumption allows win probabilities and hazards to remain well-defined for all auctions in the sample.

Winning probability and the i's hazard. When facing rival bidder set $\mathcal{J}_{\ell t}$, bidder i's probability of winning with bid b is

$$G_i(b \mid \mathcal{J}_{\ell t}) = \prod_{j \in \mathcal{J}_{\ell t} \setminus \{i\}} \bar{F}_j^{\text{mix}}(b). \tag{19}$$

Let firm i's hazard, the proportional rate at which i's probability of winning falls as i raises its bid, be

$$H_i(b \mid \mathcal{J}_{\ell t}) = -\frac{\partial \log G_i(b \mid \mathcal{J}_{\ell t})}{\partial b} = \sum_{j \in \mathcal{J}_{\ell t} \setminus \{i\}} \frac{f_j^{\text{mix}}(b)}{\bar{F}_j^{\text{mix}}(b)}. \tag{20}$$

This hazard determines firm i's competitive environment in project ℓ: it rises with each additional rival.

4.7 Dynamic Problem, Bidding, and Equilibrium

Value Function In each period, I assume all tenders are released simultaneously. For each bidder i, I consider each auction separately, conditional on the firm's state at the auction stage, as in the separable benchmark of Gentry et al. (2023). This assumption makes the quarterly flow payoff additive across a firm's auctions. Since the state is measured at the beginning of the quarter, all the firm's bids within the quarter are evaluated at the same state. [16]

Let $V_{\theta_i}(s_{it}, \bar{s}_t)$ denote firm i's beginning-of-quarter value. At the equilibrium bid policy, the random set $\mathcal{W}_{it}$ contains the projects awarded to the firm in quarter t. Because value is measured

[16] In estimation, the assumption bites in quarters where a firm participates in more than one auction – 44% of bidding firm-quarter pairs in the data.

before survival is realized, it satisfies

$$V_{\theta_i}(s_{it}, \bar{s}_t) = [1 - \zeta(s_{it}, \bar{s}_t)]\mathbb{E}\left[\sum_{\ell \in \mathcal{W}_{it}} \big(b_{i\ell t} - c_{i\ell t}\big) + \beta V_{\theta_i}(s_{i,t+1}, \bar{s}_{t+1}) \;\Big|\; s_{it}, \bar{s}_t; \theta_i\right]. \tag{21}$$

The firm exits with probability $\zeta(s_{it}, \bar{s}_t)$ and obtains zero scrap value.[17] Conditional on survival, the expectation integrates over tender arrivals, participation, private costs, rival bids, awards, clearance, category transitions, and the law of motion for $\bar{s}_t$. Entry and rival exit affect the firm through the aggregate state. The transition to $s_{i,t+1}$ follows (9)–(12); in particular, an award adds to backlog before execution and expands eligibility only after it has been completed and paid.

Continuation Value of Winning and Optimal Bids. Consider a project ℓ for which firm i is bidding. Let $s^{\text{win}}_{i,t+1}$ and $s^{\text{lose}}_{i,t+1}$ denote the next-period states when the firm wins and loses project ℓ, respectively. The two states differ only in the focal award: winning adds z_ℓ to backlog and cumulative awards before clearance. The continuation gain from winning is

$$\begin{aligned} D_{i\ell t} &\equiv D(s_{it}, \chi_\ell, \bar{s}_t; \theta_i) \\ &= \beta\mathbb{E}\left[V_{\theta_i}(s^{\text{win}}_{i,t+1}, \bar{s}_{t+1}) - V_{\theta_i}(s^{\text{lose}}_{i,t+1}, \bar{s}_{t+1}) \mid s_{it}, \chi_\ell, \bar{s}_t; \theta_i\right]. \end{aligned} \tag{22}$$

The expectation in (22) integrates over clearance, category transitions, and the law of motion for $\bar{s}_t$, which are common to both states. Winning may accelerate the accumulation of paid experience and category promotion, but it also raises backlog. Because the prospective transition approximates the award increment by z_ℓ, rather than by the bid or the signed contract value, and a single auction does not change the perceived law of motion for $\bar{s}_t$, $D_{i\ell t}$ is fixed with respect to b.

Conditional on the entrant set and the public states of its rivals, the terms in (21) that depend on firm i's bid in project ℓ reduce to

$$G_i(b \mid \mathcal{J}_{\ell t})\,[b - c_{i\ell t} + D_{i\ell t}]\,, \tag{23}$$

where $G_i(b \mid \mathcal{J}_{\ell t})$ is as in Equation (19). The full derivation is presented in Appendix D.2. Using the bid hazard in (20), the necessary First-Order Condition that the bid must satisfy is

$$b_{i\ell t} = c_{i\ell t} \quad + \quad \underbrace{\frac{1}{H_i(b_{i\ell t} \mid \mathcal{J}_{\ell t})}}_{\text{static markup}} \quad - \quad \underbrace{D_{i\ell t}}_{\text{dynamic markdown}} \tag{24}$$

The first term above cost reflects current-auction competition. The second reflects the effect of an

[17] In principle, this normalization is not innocuous, and may bias estimation of other primitives (Aguirregabiria and Suzuki, 2014). Furthermore, some counterfactuals (those that change firms' choice sets or the environment's transition structure) are also not identified without knowledge of the scrap value level (Kalouptsidi et al., 2017, 2021). In practice, the normalization is likely to be second order in this setting. Firms that exit almost always continue filing taxes with positive sales in the private market, meaning exit from the procurement market means redirecting capacity towards private market activity rather than firm death.

award on backlog, paid experience, and future eligibility. Equivalently, construction cost is recovered from

$$c_{i\ell t} = b_{i\ell t} + D_{i\ell t} - \frac{1}{H_i(b_{i\ell t} \mid \mathcal{J}_{\ell t})}. \tag{25}$$

Equation (24) is the link between the model's dynamic mechanism and observed bids (Jofre-Bonet and Pesendorfer, 2003; De Silva et al., 2023). As is standard in first-price settings, I assume each firm's bid is strictly increasing in its private cost draw, conditional on its type and state (Guerre et al., 2000; Athey and Haile, 2007), so that (25) maps each observed bid to a unique cost, given the estimated hazard and continuation value.

Equilibrium A Markov perfect equilibrium would require each firm to track the state of every active rival in each period; the large number of active firms (almost 400) makes this solution concept computationally infeasible. Following De Silva et al. (2023), I instead use the moment-based Markov equilibrium of Ifrach and Weintraub (2017), which builds on related large-market solution concepts in Weintraub et al. (2008) and Weintraub et al. (2010), which approximates the Markov perfect equilibrium. Firms condition on the exact entrant set and rivals' public states in a current auction, but use $\bar{s}_t$ to forecast future competition.

Definition 1 (Moment-based equilibrium)**.** *Fix the exogenous policies: the RPPA category promotion kernel, the auction participation policy, and the entry and exit processes. A moment-based equilibrium in bids consists of type-symmetric bid functions and value functions, firm-specific type beliefs (within auctions), and a perceived law of motion for $\bar{s}_t$ held by all firms, such that:*

1. *in each auction, each firm that submits a bid maximizes (23), given the realized entrant set, rivals' public states, and beliefs over rival types;*
2. *value functions satisfy (21) under the fixed policies and the perceived law of motion;*
3. *beliefs over rival types follow Bayes' rule from the category-specific type shares and the participation policy, as in (16), and beliefs over rival bids are given by type mixtures (18) implied by equilibrium bid policies; and,*
4. *the perceived law of motion for the aggregate state $\bar{s}_t$ is consistent with the law induced by the fixed policies and firms' equilibrium behavior.*

5 Implementation and Estimation

In this section, I describe how I construct the estimation sample and estimate the policy environment. I then describe how I estimate the model in two stages. I first estimate the policy environment, permanent execution types, and type-specific execution, participation, and bid policies. In the second stage, these objects determine auction hazards, static markups, and the transition policies used to compute continuation values. Because continuation values are computed from these first-stage objects alone and do not depend on the construction cost parameters, I solve the value function once and then estimate those parameters from the bidding first-order condition without re-solving firms' policies for each candidate parameter vector. This procedure follows the two-step approach to dynamic models in Hotz and Miller (1993); Bajari et al. (2007); Aguirregabiria and Mira (2007).

5.1 Sample Construction

I estimate the model on auctions conducted between September 2016 and October 2023. I define a firm as *regular* if it submits at least two positive, qualified bids on distinct lots before November 1, 2023. A firm enters the sample if it is regular.[18] This restriction yields a sample of 394 regular firms – the main sample used for analysis. For participation, a regular firm enters the risk set for a tender when it is eligible by category and has submitted a bid during the previous four quarters. The resulting panel contains 116,894 firm–tender pairs. Regular firms represent 56% of all bidders but 95% of all bids and 97% of all contracts won.

I restrict the auction sample to those in which at least one regular bidder competes, which yields 1,319 tender-lots and 3,035 bids. Among these, 460 auctions have a single qualified bidder. For rivals in the sample with an observed state, I evaluate the estimated bid policy at that state; for other rivals, I use the corresponding category and type-cell bid distribution (12.1% of bids). The project execution panel contains 13,440 firm–quarter observations for the 320 firms that win a contract (the remaining 74 firms win no contracts).

The first stage estimates the execution, participation, and bid policies on this firm and auction universe. The second stage then uses the same policies, states, and lot support to recover continuation values and construction costs. All monetary variables used in estimation are expressed in nominal VAT-exclusive Rwandan francs (RWF); where indicated, levels are rescaled to billions of RWF for estimation.

5.2 States and Policy Environment

Firm and Aggregate States In each quarter, I observe the category vector, cumulative awarded value W_{it}, and the current flow ΔW_{it} of Works contracts entering the portfolio. I construct backlog B_{it} recursively from these commitments and the VAT-implied Works payment flow Y_{it}^{W}, and define paid experience as $P_{it} = W_{it} - B_{it}$.[19] Quarterly commitments, backlog, and cumulative awarded value are constructed from signed contract values and Works payments (from VAT data). By construction, P_{it} is also cumulative allocated Works payment.[20]

The participation, bid, and promotion policies use versions normalized by the scale of the firm's

[18]A lead bid is submitted by the firm itself rather than through a Joint Venture (JV). I exclude bids disqualified for administrative, technical, or financial reasons.

[19]I initialize backlog using the estimated scheduled remaining value of contracts active at the beginning of the panel and omit the first four quarters as a burn-in period. The ZOIB specification is estimated from the realized clearance histories and propagates backlog in the value-function calculations and counterfactual simulations.

[20]Appendix C.1 gives the exact construction and its reconciliation with the payment data.

category. For firm i in category $\kappa_{im,t}$ in sector m at time t, I denote

$$\widetilde{B}_{it}(\kappa) = \frac{B_{it}}{S(\kappa)}, \qquad \widetilde{P}_{it}(m,\kappa) = \frac{P_{it}}{T_{m,\kappa}}, \tag{26}$$

where $S(\kappa)$ is the category capacity scale, measured as the median engineering estimate of category-κ tenders in all sectors, and $T_{m,\kappa}$ is the sector-category payment threshold. The numerators are firm totals, while the denominators vary with the firm's category.

The aggregate state summarizes the intensity of competition in each category using the number of realized bidders. I define it by

$$\bar{s}_t = \left(\bar{n}_t^A, \bar{n}_t^B, \bar{n}_t^C, \bar{n}_t^D, \bar{n}_t^E\right), \tag{27}$$

where $\bar{n}_t^k$ is the mean number of qualified bidders per category-k tender lot, pooled across the three Works sectors, during quarters $t-4,\ldots,t-1$. Appendix C.4.3 provides the exact construction and transition.

Industry Entry and Exit Firm entry is marked when the firm submits its first bid. At each period, the set of potential entrants consists of firms that registered but have not submitted a bid. Among the potential entrants, the quarterly entry hazard is

$$h_t^E = \Lambda\left(\gamma_{E0} + \bar{s}_t'\gamma_E\right), \qquad \text{where} \qquad \Lambda(u) = \frac{\exp(u)}{1+\exp(u)}. \tag{28}$$

Entrants' initial states and fixed characteristics are drawn from their empirical joint distribution conditional on the aggregate state. Conditional on $\bar{s}_t$, all potential entrants share the same hazard, so the number entering in quarter t is Binomial in the number of potential entrants.

A share of active firms exit the market each period. Among active firms, the quarterly exit hazard is

$$\zeta_{it} = \Lambda\left(\gamma_{X0} + g_X(s_{it})'\gamma_X + \bar{s}_t'\gamma_{Xs}\right), \tag{29}$$

where g_X contains category indicators, $\log(1+B_{it})$, and $\log(1+P_{it})$. Exit is absorbing. In the implementation, a firm is active from its first to its last observed bid and exits in the quarter of that last bid, with the last four sample quarters treated as right-censored so that every assigned exit is followed by at least four quarters without a bid.

Entry and exit are modelled as exogenous transitions rather than choices made by firms, so these hazards are held fixed in counterfactuals. Appendix C.4.2 reports the estimated hazards.

RPPA Category Transitions I parameterize the transition of each sector component $\kappa_{im,t}$ of the joint category kernel in (12) using a promotion hazard and a destination distribution. For an active sector m and current category $\kappa = \kappa_{im,t}$, the policy evaluates the firm-level paid experience stock against the requirements of that sector-category. After project execution realized in quarter

$t-1$, the probability that a firm is promoted category in that sector is

$$h^R_{m\kappa}(P_{it}) = \Lambda\left[\alpha_{m\kappa} + \delta_0\,\mathbf{1}\{P_{it} > 0\} + \beta_P \log\left(1 + \frac{P_{i,t}}{T_{m,\kappa}}\right)\right]. \tag{30}$$

This probability increases smoothly with paid experience. Conditional on promotion, the destination rule assigns probability

$$p^R_m(\kappa' \mid \kappa, P_{i,t+1}), \qquad R(\kappa') < R(\kappa), \tag{31}$$

to each feasible higher category. Firms moving from E-B are restricted to upwards changes of at most two ranks; moves from F are unrestricted. Moving to category A is an absorbing state. Transitions into a previously inactive sector are included in the same category transition policy. Note that while paid experience is not reset after promotion, its normalization does change with the destination category, and the backlog state is also rescaled by $S(\kappa)/S(\kappa')$ at a category transition. When computing the value of winning, a firm applies these sector-specific transitions to the full category vector $\boldsymbol{\kappa}_{it}$.

5.3 First Stage: Execution Types and First-Step Policies

I assume each firm is one of two permanent latent types, H or L. I assume that firms' execution, auction participation, and bids depend on their type.

Execution For firms with outstanding work $B_{it} + \Delta W_{it} = X_{it} > 0$, I parametrize the distribution of work completed in (7) as a zero–one-inflated Beta distribution:

$$\begin{aligned}
\mathbb{P}(\xi_{it} = 0 \mid X_{it}, \theta_i) &= p_{0\theta_i}(X_{it}),\\
\mathbb{P}(\xi_{it} = 1 \mid \xi_{it} > 0, X_{it}, \theta_i) &= p_{1\theta_i}(X_{it}),\\
\xi_{it} \mid 0 < \xi_{it} < 1, X_{it}, \theta_i &\sim \mathrm{Beta}\Big(\mu^\xi_{\theta_i}(X_{it})\phi, [1 - \mu^\xi_{\theta_i}(X_{it})]\phi\Big),
\end{aligned} \tag{32}$$

The Beta precision $\phi > 0$ is common across types. I implement the three cases as,

$$\begin{aligned}
\operatorname{logit} p_{0\theta}(X_{it}) &= \alpha^\xi_0 + \beta^\xi_0 x^\xi_{it} - \lambda^\xi \mathbf{1}\{\theta = H\},\\
\operatorname{logit} p_{1\theta}(X_{it}) &= \alpha^\xi_1 + \beta^\xi_1 x^\xi_{it} + \lambda^\xi \mathbf{1}\{\theta = H\},\\
\operatorname{logit} \mu^\xi_{\theta}(X_{it}) &= \alpha^\xi_\mu + \beta^\xi_\mu x^\xi_{it} + \lambda^\xi \mathbf{1}\{\theta = H\}, \qquad \lambda^\xi > 0,
\end{aligned} \tag{33}$$

where $x^\xi_{it} = [\log X_{it} - \overline{\log X}]/\mathrm{sd}(\log X)$ is the normalized outstanding work. The coefficients β^ξ_0, β^ξ_1, and β^ξ_μ allow the clearance outcomes to vary with the post-award stock X_{it}. Execution type enters through the parameter $\lambda^\xi > 0$, which lowers the zero-clearance probability and raises both the full-clearance probability and the interior mean. This parameter restriction implements the stochastic ordering in (8). [21]

[21] I also estimated a two-type Bernoulli specification based on indicators about whether quarterly payments met scheduled progress. The main qualitative results were largely robust to this alternative execution process.

Participation For project ℓ, set $m = m_\ell$ and let $\Delta R_{i\ell t} = R(k_\ell) - R(\kappa_{im,t})$ denote the category-rank gap between the firm and the tender. With type L as the omitted category, the participation policy is

$$\mathbb{P}(a_{i\ell t} = 1 \mid \theta_i) = \Lambda\Big(\delta_A + \underbrace{\delta_M \mathbf{1}\{\kappa_{im,t} = k_\ell\} + \delta_{\bar{M}} \overline{\mathbf{1}\{\kappa_{im,t} = k_\ell\}}_i}_{\text{same category}} + \underbrace{\delta_R \Delta R_{i\ell t} + \delta_{\bar{R}} \overline{\Delta R}_i}_{\text{category-rank gap}}$$
$$+ \underbrace{\delta_B \widetilde{B}_{it}(\kappa_{im,t}) + \delta_{\bar{B}} \overline{\widetilde{B}}_i}_{\text{backlog}} + \underbrace{\delta_P \log[1 + \widetilde{P}_{it}(m, \kappa_{im,t})] + \delta_{\bar{P}} \overline{\log(1 + \widetilde{P})}_i}_{\text{paid experience}}$$
$$+ \delta_s' \bar{s}_t + \delta_H \mathbf{1}\{\theta_i = H\}\Big). \tag{34}$$

I control for firm-specific mean terms (denoted with bars) to separate changes in a firm's state from persistent differences across firms (Mundlak, 1978).

Bids Let $\widetilde{b}_{i\ell t} = b_{i\ell t}/z_\ell$ be firm i's bid in project l at time t divided by the engineering cost estimate of project l. Conditional on firm i's private type θ_i, this normalized bid is assumed to be drawn from a type-specific Gamma distribution,

$$\widetilde{b}_{i\ell t} \mid \theta_i \sim \text{Gamma}\left(\nu_{\theta_i}, \frac{\nu_{\theta_i}}{\mu^b_{\theta_i}(\chi_\ell, s_{it}, \bar{s}_t)}\right), \tag{35}$$

where $\nu_\theta > 0$ and the second argument is the rate, so μ^b_θ is the conditional mean which depends on project-level covariates, the firm's individual state, and the aggregate state. For a vector $x^b_{i\ell t}$ containing the standardized normalized states (26), project scale, bidder count, calendar quarter, sector indicators, and the aggregate state, I set

$$\log \mu^b_\theta(\chi_\ell, s_{it}, \bar{s}_t) = \alpha^b_{\theta, \kappa_{im,t}, k_\ell} + (x^b_{i\ell t})' \beta^b_\theta, \qquad m = m_\ell. \tag{36}$$

The intercepts vary by firm-category–tender-category cell. These bid policies recover the type-specific densities and survival functions used in (18).

Firm Type Likelihood For firm i, let $\mathcal{D}^\xi_i$, $\mathcal{D}^b_i$, and $\mathcal{D}^a_i$ denote its execution, normalized-bid, and participation histories. Conditional on type and the covariates in each component, the firm likelihood is

$$\mathcal{L}_i = \sum_{\theta \in \{L, H\}} \omega_\theta \, \mathcal{L}^\xi_i(\theta) \mathcal{L}^a_i(\theta) \mathcal{L}^b_i(\theta). \tag{37}$$

Appendix C.5.1 derives this factorization from the chain rule and the maintained Markov and conditional independence restrictions.

Here $\omega_\theta = \mathbb{P}(\theta_i = \theta)$, with $\omega_H = \Lambda(\alpha^\omega)$ and $\omega_L = 1 - \omega_H$. The firm's type likelihood depends on the product of three likelihoods: for project execution, auction participation, and bidding. Each block is a product over that firm's own observations. Collecting the three cases of (32) into a single product, The execution block depends on the product of the three cases of (32):

$$\mathcal{L}_i^{\xi}(\theta) = \prod_{t\in\mathcal{D}_i^{\xi}} \Big[p_{0\theta}(X_{it})\Big]^{\mathbf{1}\{\xi_{it}=0\}} \Big[\big(1-p_{0\theta}(X_{it})\big)p_{1\theta}(X_{it})\Big]^{\mathbf{1}\{\xi_{it}=1\}}$$
$$\times \Big[\big(1-p_{0\theta}(X_{it})\big)\big(1-p_{1\theta}(X_{it})\big)\, f_{\mathrm{B}}\big(\xi_{it};\, \mu_\theta^{\xi}(X_{it})\phi,\, [1-\mu_\theta^{\xi}(X_{it})]\phi\big)\Big]^{\mathbf{1}\{0<\xi_{it}<1\}}, \tag{38}$$

where f_{B} is the Beta density, and the three components are as in (33). Letting $\eta_{i\ell t}(\theta)$ denote the index inside $\Lambda(\cdot)$ in (34), which differs across types only through the level shift δ_H, the participation block is

$$\mathcal{L}_i^{a}(\theta) = \prod_{(\ell,t)\in\mathcal{D}_i^{a}} \Lambda\big(\eta_{i\ell t}(\theta)\big)^{a_{i\ell t}}\Big[1-\Lambda\big(\eta_{i\ell t}(\theta)\big)\Big]^{1-a_{i\ell t}}. \tag{39}$$

With $\widetilde{b}_{i\ell t} = b_{i\ell t}/z_\ell$ and μ_θ^b from (36), the Gamma density in (35) gives the bid block

$$\mathcal{L}_i^{b}(\theta) = \prod_{(\ell,t)\in\mathcal{D}_i^{b}} \frac{1}{\Gamma(\nu_\theta)}\left(\frac{\nu_\theta}{\mu_\theta^b(\chi_\ell, s_{it}, \bar{s}_t)}\right)^{\nu_\theta} \widetilde{b}_{i\ell t}^{\,\nu_\theta-1}\exp\left(-\frac{\nu_\theta\,\widetilde{b}_{i\ell t}}{\mu_\theta^b(\chi_\ell, s_{it}, \bar{s}_t)}\right). \tag{40}$$

Bids enter conditional on participation, so (39) and (40) together factor the joint distribution of the participation decision and the bid it produces.

EM algorithm Instead of maximizing (37) directly, I use the Expectation–Maximization algorithm in Arcidiacono and Miller (2011). It divides that procedure into two simpler steps, iterated to convergence. The E-step fills in the missing type, computing for each firm the posterior probability that it is type H given its complete history. The M-step then proceeds as if types were observed, re-estimating each policy on a sample in which every observation appears once as each type, weighted by that posterior. Assume iteration r has occurred; then, iteration $r+1$ proceeds as follows.

1. **Expectation step.** Let

$$\Psi = (\alpha^{\omega}, \psi^{\xi}, \psi^{b}, \psi^{a})$$

collect the parameters for the type share and the execution process, normalized-bid, and participation policies. Given $\Psi^{(r)}$, compute the posterior $\widehat{w}_i^{(r)}$ for each firm. Because $\omega_H = \Lambda(\alpha^{\omega})$, this takes a log-odds form,

$$\operatorname{logit}\widehat{w}_i^{(r)} = \alpha^{\omega(r)} + \sum_{c\in\{\xi,b,a\}} \Big[\ln \mathcal{L}_i^{c}\big(H;\psi^{c(r)}\big) - \ln \mathcal{L}_i^{c}\big(L;\psi^{c(r)}\big)\Big], \tag{41}$$

so each block enters the classification of firm i as an additive log-likelihood ratio, and a firm is classified H when its execution, bidding, and participation records jointly favor H over L.

2. **Maximization step.** Given $\{\widehat{w}_i^{(r)}\}$, choose $\Psi^{(r+1)}$ to maximize

$$Q\big(\Psi \mid \Psi^{(r)}\big) = \sum_{i=1}^{N} \sum_{\theta \in \{L,H\}} w_{i\theta}^{(r)} \Big[\ln \omega_\theta + \ln \mathcal{L}_i^{\xi}(\theta) + \ln \mathcal{L}_i^{b}(\theta) + \ln \mathcal{L}_i^{a}(\theta) \Big], \tag{42}$$

with $w_{iH}^{(r)} = \widehat{w}_i^{(r)}$ and $w_{iL}^{(r)} = 1 - \widehat{w}_i^{(r)}$. Because no parameter enters more than one likelihood block, (42) separates into four independent maximization problems, which are solved in sequence (Arcidiacono and Jones, 2003):

(a) the type share in closed form, $\omega_H^{(r+1)} = N^{-1} \sum_i \widehat{w}_i^{(r)}$;

(b) the execution parameters ψ^{ξ} in (33), by weighted maximum likelihood;

(c) the bid policy ψ^{b} in (36), by weighted maximum likelihood within each firm-category–tender-category cell;

(d) the participation policy ψ^{a} in (34), by weighted logit on the eligible risk set.

In (b)–(d) each observation enters twice, once as type L with weight $1 - \widehat{w}_i^{(r)}$ and once as type H with weight $\widehat{w}_i^{(r)}$.

Posteriors are updated with a damping factor of 0.5, and iteration stops when the sup-norm change is below 10^{-3} in the posteriors and 2×10^{-3} in the stacked policy vector.[22]

5.4 Second Stage: Auction Hazards, Continuation Values, and Costs

The first stage provides the type posteriors and the type-specific execution, participation, and bid policies. In the second stage, I use these objects to construct the auction markup in each observed auction, evaluate how winning changes a firm's continuation value, and estimate the parameters of the construction cost function. The policies remain fixed throughout this stage; I evaluate the value they induce rather than re-solving firms' participation and bidding policies for each candidate cost parameter.

Auction hazards and markups. For each firm category, I construct the prior probability that a firm is type H from the estimated posterior masses of firms observed in that category. Once a rival is observed in the entrant set, I update this prior using the type-specific participation probabilities in (16). I then evaluate the two conditional Gamma bid distributions at the rival's public state and combine their densities and survival functions using the type share conditional on participation in (18).

The resulting estimated auction markup is denoted

$$\widehat{\mu}_{i\ell t} = \frac{1}{\widehat{H}_i(b_{i\ell t} \mid \mathcal{J}_{\ell t})}. \tag{43}$$

[22] Any $\Psi^{(r+1)}$ that raises $Q(\Psi \mid \Psi^{(r)})$ above $Q(\Psi^{(r)} \mid \Psi^{(r)})$ also raises the observed log likelihood (37), so each pass weakly improves the fit.

Rival identities are used to recover public states and qualification categories, but not realized types or the econometrician's posterior.

Construction cost specification. I parameterize construction costs relative to the engineering estimate as

$$\begin{aligned}\frac{c_{i\ell t}}{z_\ell} = \; & \underbrace{\alpha^{\text{tend}}_{k_\ell} + \alpha^{\text{firm}}_{\kappa_{im_\ell},t} + \alpha^{\text{sec}}_{m_\ell}}_{\text{category and sector effects}} + \underbrace{\gamma\frac{\log z_\ell}{25} + q_\ell'\beta_c}_{\text{project characteristics}} \\ & + \underbrace{\rho_B \log(1+B_{it}) + \rho_P \log(1+P_{it}) + \rho_{BP}\log(1+B_{it})\log(1+P_{it})}_{\text{backlog and paid experience}} \\ & + \underbrace{\eta_H \mathbf{1}\{\theta_i = H\}}_{\text{execution type}} + \varepsilon_{i\ell t}, \end{aligned} \tag{44}$$

where backlog and paid experience are measured in billions of Rwandan francs. The vector q_ℓ contains BOQ item count, advertised duration, procuring entity type, the lagged material price shock, and missing-value indicators.

Let ϑ collect the coefficients that enter the conditional mean in (44), and write its covariates as $x_{i\ell t}(\theta)$. Conditional on public state and type, $\varepsilon_{i\ell t} \sim N(0, \sigma^2_{i\ell t})$ independently across firm–project matches. Residual dispersion in costs may depend on project size:

$$\log \sigma_{i\ell t} = \alpha^\sigma_0 + \beta^\sigma_z \left(\log z_\ell - \overline{\log z}\right). \tag{45}$$

In a static auction, costs could be recovered by subtracting the estimated markup (43) from firms' bids. In our setting, the bid also reflects the continuation value of winning, which I recover as follows.

Expected current profits. I compute continuation values under the participation and bid policies estimated in the first stage. For each type and point on the (discretized) state grid, I integrate over the tender categories for which the firm is eligible. The arrival intensity λ_k is set to the average number of category-k lots per quarter in the estimation window, and the characteristics distribution G^χ_k is approximated by a 12-point discretized lot distribution, which simulates representative projects and their characteristics from observed projects.[23]

The flow payoff at a state is the expected profit per quarter from bidding. By the bid first-order condition, a firm's profit from an auction at its optimal bid, inclusive of the continuation value of winning, equals its winning probability divided by the hazard (De Silva et al., 2023). Let

$$\pi^A_{i\ell t} = \frac{G_i(b_{i\ell t} \mid \mathcal{J}_{\ell t})}{H_i(b_{i\ell t} \mid \mathcal{J}_{\ell t})}$$

[23]Within each tender category, I sort observed projects by their engineering estimate, divide them into twelve approximately equal-sized strata, and use the median project in each stratum, weighted by the stratum's empirical frequency.

denote this profit for firm i in auction ℓ at time t. The flow payoff at state s is the expected sum of these profits over the auctions the firm bids on in a quarter,

$$\Pi_\theta^A(s) = E\left[\sum_{\ell \in \mathcal{O}_{it}} a_{i\ell t}\, \pi_{i\ell t}^A \,\middle|\, s_{it} = s,\ \theta_i = \theta\right], \tag{46}$$

where $\mathcal{O}_{it}$ is the set of lots arriving in quarter t for which i is eligible, and $a_{i\ell t}$ is the participation indicator. The expectation integrates over project arrivals, participation, the firm's own bid, the realized rival set under the estimated policies, weighted by survival. Let $\mathbf{\Pi}_\theta^A$ denote the vector that stacks $\Pi_\theta^A(s)$ over the state grid.

For every bid, $\widehat{G}_i(b_{i\ell t} \mid \mathcal{J}_{\ell t})/\widehat{H}_i(b_{i\ell t} \mid \mathcal{J}_{\ell t})$ is the estimated winning probability divided by the estimated hazard at that bid. I regress this ratio on project characteristics, the firm's state, the aggregate state, and type, weighting each bid's two type rows by the posterior, and evaluate the fitted mean at the representative projects and grid states. Multiplying by the participation probability and the arrival intensities and summing over categories gives $\mathbf{\Pi}_\theta^A$. States with no observed bids are filled by the fitted regression. The estimated participation and bid policies are held fixed in this stage. Consequently, the transition matrix does not depend on ϑ.

State transitions. Let M_θ^0 denote the "no-award" transition matrix for type θ under the estimated policies. This matrix retains clearance, payments, experience, promotion, exit, and the aggregate-state law of motion, but assumes the firm wins no future contracts. This restriction is necessary because the auction surplus in (46) already contains both the current profit and the continuation value of winning.

For each award realization, I apply the state transitions (as in (9)–(12)): new contract awards enter the backlog, type-specific clearance determines remaining backlog and paid experience, the resulting payment enters the category transition policy, and backlog is renormalized if the firm changes category. Winning a contract therefore has two distinct dynamic effects: unfinished work may raise congestion, and completed work increases the probability of category promotion.

The quarterly discount factor is set to $\beta = 0.95$. Following Jofre-Bonet and Pesendorfer (2003), I evaluate the infinite-horizon Bellman system under the participation and bidding behavior estimated in the first stage. Let $\mathbf{V}_\theta$ denote the vector obtained by evaluating the value function $V_\theta(s, \bar{s})$ at every point on the state grid. Because the estimated policies and the induced transition matrix are held fixed, this vector is

$$\mathbf{V}_\theta = (I - \beta M_\theta^0)^{-1} \mathbf{\Pi}_\theta^A. \tag{47}$$

This is the auction-surplus representation of De Silva et al. (2023). Because neither $\mathbf{\Pi}_\theta^A$ nor M_θ^0 depends on the cost parameters, the Bellman system is solved once per type, and the construction cost parameters are recovered only afterwards. Evaluating (47) solves an infinite-horizon Bellman system.[24]

[24] For numerical validation, I also compute continuation values by simulating the same policy and transition law

I compute the dynamic values for each $\theta \in \{L, H\}$ separately. Thus, the flow payoffs, transition matrix, and value vector—$\mathbf{\Pi}_\theta^A$, M_θ^0, and $\mathbf{V}_\theta$—are type-specific because execution, participation, and bidding policies differ by type.

Continuation value of the focal award. For each observed bid, I create separate rows, $Q_{\theta,i\ell t}^{\text{win}}$ and $Q_{\theta,i\ell t}^{\text{lose}}$, that represent the path if the firm wins and loses the focal award, respectively. Because the signed value is not yet observed at the bidding stage, the winning path adds the project's VAT-exclusive engineering cost estimate z_ℓ to the firm's backlog as a proxy for the prospective signed award increment before drawing clearance. The losing path adds no award. Both integrate over the same clearance, promotion, survival, and aggregate-state transitions. The dynamic gain entering the bid first-order condition is the discounted change in next-quarter continuation value,

$$\Delta V_{\theta,i\ell t} = \left(Q_{\theta,i\ell t}^{\text{win}} - Q_{\theta,i\ell t}^{\text{lose}}\right) \mathbf{V}_\theta, \qquad D_{\theta,i\ell t} = \beta \, \Delta V_{\theta,i\ell t}. \tag{48}$$

The current winning uses the award value of focal project ℓ; the continuation paths integrate over future tender arrivals using representative project values.

Structural estimating equation. Substituting (44), (43), and (48) into the bid first-order condition gives

$$\frac{b_{i\ell t} - \widehat{\mu}_{i\ell t} + D_{\theta,i\ell t}}{z_\ell} = x_{i\ell t}(\theta)'\vartheta + \varepsilon_{i\ell t}, \quad \text{where,} \quad \widehat{\mu}_{i\ell t} = \frac{z_\ell}{\widehat{H}_i(b_{i\ell t}/z_\ell)}. \tag{49}$$

Because $D_{\theta,i\ell t}$ is computed from the value vector before construction cost parameters are estimated, the left-hand side is data once the first stage and the value solve are complete, and ϑ enters only through the cost specification (44) on the right-hand side. Thus the observed bid is separated into construction cost, a static markup, and a dynamic markdown. Each bid contributes one type-L and one type-H row to (49), weighted respectively by $1 - \widehat{w}_i$ and $\widehat{w}_i$. I estimate ϑ by the posterior-weighted heteroskedastic likelihood implied by (45). Firms know their own types, while the econometrician integrates over type using the posterior estimated in the first stage.

5.5 Inference

Inference uses 300 firm-level bootstrap samples. I hold the first stage fixed and re-estimate every second-stage object in each sample, so the reported intervals for continuation values and costs reflect uncertainty in the set of firms, their costs, and the continuation values, conditional on the estimated types and policies. I construct confidence intervals for estimated parameters and the dynamic value of winning ($\bar{D}_L$, $\bar{D}_H$, $\bar{D}_H - \bar{D}_L$) from the draw-level estimates.

forward for five hundred quarters. The results match.

6 Identification Discussion

Identification proceeds in three steps. First, the administrative histories identify the policy functions and state transitions used in the model. Second, repeated execution and auction histories identify the distribution of permanent execution types and the type-specific bid, participation, and project execution policies. Third, the auction first-order condition identifies the construction cost net of the continuation value of winning. Construction costs and continuation values are then separated using the dynamic value and linear structure implied by the model.

State Transitions and Policy Functions At the beginning of quarter t, I observe the category vector $\boldsymbol{\kappa}_{it}$ and cumulative awarded Works value W_{it}. During the quarter, I observe the value ΔW_{it} of Works contracts entering the firm's portfolio and construct the Works payment flow Y_{it}^{W} from the VAT records. I then construct backlog recursively using

$$\begin{aligned} W_{i,t+1} &= W_{it} + \Delta W_{it}, \\ B_{i,t+1} &= B_{it} + \Delta W_{it} - Y_{it}^{W}. \end{aligned}$$

The payment flow is capped by the portfolio available to receive payment, so $B_{i,t+1}$ is nonnegative. I treat the resulting measured state $s_{it} = (\boldsymbol{\kappa}_{it}, B_{it}, W_{it})$, and the aggregate state $\bar{s}_t$ (which is observed), as given when estimating the policy functions.

Tender arrivals, industry entry and exit, and the aggregate state transition are assumed to be exogenous transition policies that are observed from their corresponding panels. The promotion policy is identified from category transitions conditional on origin category, sector, and paid experience, all of which are observed in each period. As in De Silva et al. (2023), auction participation is treated as a descriptive matching policy. I observe all eligible active firms and all their bid submissions to all tenders, so the participation policy is identified from match rates by eligible firm-tender pairs. These objects are identified (on the relevant support) under the maintained Markov restriction that, conditional on $(s_{it}, \bar{s}_t)$ and the current tender, earlier history does not directly affect the corresponding policy or transition.

Permanent Execution Types Identification of θ_i relies on firm-level repeated histories, which distinguish permanent heterogeneity from transitory shocks. Specifically, Kasahara and Shimotsu (2009) show that a model with a finite number of permanent types can be identified if current outcomes satisfy a Markov restriction, if there is sufficient variation in observed states that affects types differentially, and if the type-specific choice probabilities satisfy a rank condition ensuring that the types' response profiles are linearly independent. In my context, the Markov restriction is that conditional on $(\theta_i, s_{it}, \bar{s}_t, \chi_\ell)$, the per-period draws generating clearance, participation, and the conditional bid are independent, and draws are independent across quarters and auctions. Thus, the only dependence over time arises from the firm's individual state and the common aggregate state.

Intuitively, identification is based on observing repeated sequences of firm behavior that persistently differ between the two types. For example, consider two firms facing the same individual and aggregate states. A single quarter of fast execution could reflect a favorable transitory draw. But, if one firm repeatedly clears more of its portfolio, participates according to the type-H policy, and submits bids drawn from the type-H distribution, that sequence is more likely to have been generated by $\theta_i = H$ than by $\theta_i = L$. The conditions in Kasahara and Shimotsu (2009) ensure that a finite number of groups of firm types can be distinguished and thus the underlying finite-mixture can be identified.

I assume that execution type is a permanent and immutable firm characteristic; firms cannot choose to exert effort to increase their execution speed after being awarded a contract. The model thus rules out moral hazard. Any persistent effort effect is attributed to θ_i.

Auction Hazards and Cost Inversion Conditional on participation, the first-stage finite-mixture model identifies the parameters of the type-specific bid distribution

$$F^b_\theta\left(\widetilde{b} \mid \chi_\ell, s_{it}, \bar{s}_t\right), \qquad \widetilde{b}_{i\ell t} = \frac{b_{i\ell t}}{z_\ell},$$

under the parametric Gamma specification in (35)–(36), and identification of the type.

For each auction, I observe the participating bidder set and each bidder's public state. Bidders form beliefs about rival types by updating a category-specific prior over rival types using the type-specific probability of participation. Combining this probability with the type-specific bid distributions identifies each rival's bid density and survival function. These objects identify the auction hazard $H_i(b \mid \mathcal{J}_{\ell t})$ on the support of observed bids (Athey and Haile, 2007). The static markup $1/H_i(b \mid \mathcal{J}_{\ell t})$ is well-defined on every row because the bidder's hazard is a sum of Gamma hazards with shapes above one (the smallest being 6.5). For single-bidder lots, the hazard comes from the simulated rival.

Given the discount factor, the estimated clearance, promotion, participation, entry, exit, and aggregate-state transitions, together with the cost parameters, determine the value function. Winning differs from losing because an award increases backlog and cumulative awarded value. These changes affect subsequent payments, participation, promotion, and eligibility. The difference between the winning and losing value functions defines the continuation value $D_{i\ell t}$, which does not depend on the unknown cost parameters conditional on the first-step policies (Hotz and Miller, 1993; De Silva et al., 2023).

Given the auction hazard and continuation value, construction cost is recovered bid by bid using the Guerre et al. (2000)-style inversion adapted to the dynamic auction setting, as in Jofre-Bonet and Pesendorfer (2003):

$$b_{i\ell t} - \frac{1}{H_i(b_{i\ell t} \mid \mathcal{J}_{\ell t})} = c_{i\ell t} - D_{\theta_i, i\ell t}.$$

On the left-hand side, bids are observed and the auction hazard, and therefore the static markup, is identified as described above. Thus, the auction inversion identifies project cost net of the continuation gain from winning, but does not identify the two components separately.

Separating Costs and Continuation Values The auction inversion identifies only the net cost $c_{i\ell t} - D_{\theta,i\ell t}$. Because the continuation gain $D_{\theta,i\ell t}$ in (48) is computed from the first-stage policies and the value vector alone, the net cost separates into its two components bid by bid, without reference to the cost parameters. The construction cost specification (44) then models the recovered cost as $x_{i\ell t}(\theta)'\vartheta + \varepsilon_{i\ell t}$, and Equation (49) identifies ϑ given the other observed or estimated components.

Conditional on the first-step objects, and under the model's cost specification, the cost parameters ϑ are identified in our model under the key assumption that

$$\varepsilon_{i\ell t} = \sigma_{i\ell t} u_{i\ell t}, \qquad u_{i\ell t} \overset{\text{iid}}{\sim} N(0,1),$$

conditional on tender characteristics, public state, and type, where $\sigma_{i\ell t}$ is specified in (45). That is, the structural cost errors are conditionally independent across firm–project matches, firms, and periods. It is also necessary that the cost variables in (49) have full column rank.

Intuitively, the same variables enter current costs and continuation values in different ways. For example, current backlog affects the cost of executing the focal project through its level at the time of bidding. Its effect on the continuation value instead depends on how winning changes the firm's future backlog relative to losing, which varies with project size, execution type, paid experience, and the probability of promotion. As a result, firms with similar current values in the cost covariates can face different future effects from winning. This variation separates current construction cost from the continuation value (De Silva et al., 2023).

Counterfactuals The counterfactuals are identified conditional on four restrictions. First, the project cost distribution and type-specific clearance functions remain fixed. Second, the estimated tender arrival, participation, entry, and exit processes, as well as the category promotion policy, remain the same under the counterfactual rule except where the exercise explicitly changes them. Third, firms use the observed public state and current rival states in each auction, while forecasting future competition from $\bar{s}_t$. Fourth, the counterfactual remains within the support over which the type-specific policies and cost distribution are estimated. Thus, the counterfactual changes the contract award mechanism, competition in auctions, backlog, paid experience, and dynamic markdowns, but holds the environment, government behavior, and project completion speeds fixed.

7 Estimation Results

This section presents the structural model estimates, quantifies the magnitude of the dynamic incentives, and the effectiveness of screening in this market. Section 7.1 presents the results of Stage 1: execution types and the estimated participation and bid policies; Section 7.2 reports the results of Stage 2 cost and continuation-value estimates, and decomposes bids to show how dynamic incentives induce high-type firms to shade their bids; Appendix E.1 assesses model fit.

7.1 First-stage estimates: firm types and participation and bid policies

The EM algorithm estimates type-specific clearance distributions that differ sharply between high- and low-type firms. The algorithm assigns 19.8 percent of firms to the high type; the remaining 80.2 percent is assigned to the low type.

The estimated clearance process differs across types on each of its three margins. At the median backlog stock, high-type firms have zero payment in 36 percent of quarters, compared with 52 percent for low-type firms. They fully clear their outstanding portfolio in 15 percent of quarters, compared with 7 percent for low-type firms. Among quarters with partial clearance, the high type clears 55 percent of its portfolio on average, compared with 39 percent for the low types. Figure 5, Panel B, plots the implied distribution of quarterly clearance at the median backlog. Combining these margins, the high type clears 42.3 percent of its outstanding portfolio per quarter ($[33.1, 49.9]$ percent), compared with 22.4 percent for the low type ($[18.9, 26.3]$ percent). This implies that, with 148.5 million RWF of backlog, a high-type firm finishes half of their outstanding work in 1.3 quarters and 90 percent in 4.2 quarters, while a low-type firm takes approximately 2.7 and 9.1 quarters, respectively.

The clearance distribution varies with the amount of work outstanding. As the post-award stock increases, the probability of full clearance and the share cleared in partial-payment quarters fall, while the probability of no clearance changes little. Larger portfolios therefore clear more slowly (in percentage terms) for both types (Figure 5, Panel A). Appendix Table 14 reports the estimated zero one inflated Beta parameters and confidence intervals.

Figure 5: Estimated zero-one-inflated Beta clearance process by type

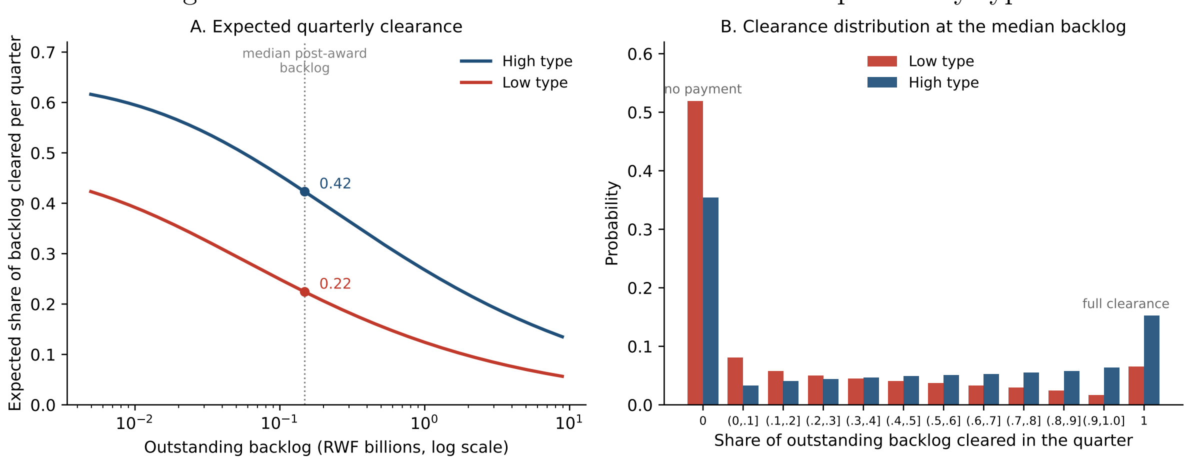

Notes: Panel A: expected share of outstanding backlog cleared per quarter as a function of the backlog stock. The dotted line marks the median post-award backlog: 148.5 million RWF. Panel B: the implied distribution of quarterly clearance at the median backlog. The bar at 0 represents no payment, and the one at 1 represents full backlog clearance. The interior bars are the Beta component integrated over deciles. Parameters are the point estimates reported in Appendix Table 14.

Table 1 reports the main first-stage estimates for the execution, participation, and bid policies. The participation policy includes each firm's mean backlog and paid experience, so the coefficients on the current states are identified from deviations from those means. Firms that operate with greater

Table 1: First-stage estimates

Quantity	Estimate	95% bootstrap CI
Execution and type		
High-type population share (ω_H)	19.8%	[13.4%, 27.5%]
Expected share cleared per quarter, L ($\mathbb{E}[\xi \mid X_{\text{med}}, L]$)	22.4%	[18.9%, 26.3%]
Expected share cleared per quarter, H ($\mathbb{E}[\xi \mid X_{\text{med}}, H]$)	42.3%	[33.1%, 49.9%]
Qualified participation		
Current backlog (δ_B)	−0.242	[−0.339, −0.163]
Paid experience (δ_P)	−0.488	[−0.650, −0.337]
Same-category tender (δ_M)	−0.728	[−1.030, −0.361]
Category gap (δ_R)	−1.204	[−1.431, −0.947]
Firm mean backlog ($\delta_{\bar{B}}$)	0.454	[0.285, 0.691]
Firm mean paid experience ($\delta_{\bar{P}}$)	0.617	[0.349, 0.911]
High type (δ_H)	1.331	[1.054, 1.593]
Conditional normalized-bid policy		
Gamma shape, L (ν_L)	6.84	[6.06, 7.96]
Gamma shape, H (ν_H)	15.44	[12.62, 22.44]
Mean predicted bid / engineering estimate, L ($\overline{\mu}^b_L$)	1.055	[1.027, 1.091]
Mean predicted bid / engineering estimate, H ($\overline{\mu}^b_H$)	1.038	[1.011, 1.079]
Mean predicted bid difference, H minus L ($\overline{\mu}^b_H - \overline{\mu}^b_L$)	−0.017	[−0.068, 0.030]

Notes: First-stage execution, participation and bid-policy point estimates. Intervals are 95 percent percentile intervals from 300 firm-bootstrap draws that re-estimate the first step (execution, participation and bid policy) on each resampled firm set, with the EM started from the point's type posteriors. The predicted bids evaluate the two type-specific conditional Gamma policies at the same observed states and auction cells; the H-minus-L difference is calculated within state. Here $\overline{\mu}^b_\theta$ denotes the sample average of the conditional mean μ^b_θ evaluated at those common states. Additional parameter estimates are reported in the appendix.

backlog and with more paid experience participate more on average. However, firm participates less in quarters when its backlog or paid experience is above its average level. High-type firms participate more often conditional on the same state.

Conditional on the same state and auction cell, the two types have nearly identical predicted bids: the H-minus-L difference is small and not statistically significant. The first stage therefore distinguishes the types through execution and participation, not through differences in bid levels.

Figure 6 shows that high-type firms are concentrated at the top of the category ladder. They account for 97 percent of firms observed in category A and 80 percent in category B, compared with 14 percent in category F.

Figure 6: Execution types by registration category

Notes: A firm enters a category if it ever submitted a qualified bid while registered in it, so firms may appear in multiple categories. Bars are means of the classifier posterior over firms in the category, weighing firms equally. Dots show the posterior weighted by bids.

7.2 Second-stage estimates: cost and continuation value

Table 2 reports the main parameter estimates recovered in the second stage. High-type firms have higher estimated costs than low-type firms. The full-sample estimate of the high-type cost coefficient η_H is 0.437 of the ECE, or 43.7 percentage points (bootstrap CI $[-0.207, 1.085]$). The coefficients on backlog, paid experience, and their interaction are -0.104, -0.115, and 0.017; none is statistically different from zero, with 95 percent bootstrap intervals of $[-0.635, 0.768]$, $[-0.462, 0.417]$ and $[-0.442, 0.260]$. Table 2 also reports the firm-category, tender-category and sector effects, the project-scale coefficient and the cost dispersion; the project controls are in Appendix E.

Using the estimated cost and continuation values, I decompose bids for high- vs. low-type firms for otherwise similar firms. Figure 7 presents the model-generated bid decomposition for firms in the same tender and public state. The high type's project cost is estimated to be 43.7pp higher (as a percent of the ECE), while its dynamic markdown is 37.7pp larger, so the dynamic channel offsets most of the cost gap. On average, the total bid is 6.0 points higher for a high-type firm, and it bids below its low-type counterpart in 48.6 percent of bids (holding the firms' states and tender characteristics fixed).

Table 3 measures the dynamic markdown, $\beta\, D_{\theta_i,ilt}$, relative to the ECE. The high-type markdown is much larger than the low-type markdown in every category in percentage point terms, by a factor of up to 10. As a percent of the ECE, the markdown is smallest for category A firms, which have no prospect of promotion to a higher category. It is largest in the middle of the ladder, reaching 76 percent of the ECE for high-type firms in category E and ranging from 50–64 percent between categories D–B.

Table 2: Second-stage cost estimates

Quantity	Estimate	95% bootstrap CI
Current cost: type, states, and scale (share of the engineering estimate)		
High type (η_H)	0.437	[−0.207, 1.085]
Log backlog (ρ_B)	−0.104	[−0.635, 0.768]
Log paid experience (ρ_P)	−0.115	[−0.462, 0.417]
Backlog × paid experience (ρ_{BP})	0.017	[−0.442, 0.260]
Log project scale (γ)	−1.667	[−3.860, 0.583]
Firm-category effects (categories A and B pooled and omitted)		
Category F (α_F^{firm})	−0.495	[−1.082, 0.359]
Category E (α_E^{firm})	−0.474	[−1.114, 0.325]
Category D (α_D^{firm})	−0.291	[−0.826, 0.264]
Category C (α_C^{firm})	−0.185	[−0.626, 0.253]
Tender-category effects (equation intercepts)		
Category F (α_F^{tend})	2.633	[0.827, 4.375]
Category E (α_E^{tend})	2.605	[0.659, 4.349]
Category D (α_D^{tend})	2.629	[0.656, 4.424]
Category C (α_C^{tend})	2.606	[0.560, 4.509]
Category B (α_B^{tend})	2.189	[−0.126, 4.095]
Category A (α_A^{tend})	2.349	[0.244, 4.387]
Sector effects (Buildings omitted)		
Roads ($\alpha_{\text{Roads}}^{\text{sec}}$)	−0.018	[−0.591, 0.178]
Water ($\alpha_{\text{Water}}^{\text{sec}}$)	−0.205	[−0.445, 0.091]
Cost dispersion		
σ	0.607	[0.536, 1.091]
Log-project-scale variance slope (β_z^σ)	−0.090	[−0.197, −0.002]

Notes: Second-stage estimates of the construction cost equation (44). The Estimate column reports full-sample point estimates; intervals are 95 percent percentile intervals from 300 firm-bootstrap draws that hold the first stage fixed. Cost coefficients are expressed as shares of the engineering cost estimate. Backlog and paid experience enter as $\log(1 + \cdot)$ in billions of VAT-exclusive Rwandan francs; project scale enters as $\log z_\ell/25$. The tender-category terms are the intercepts of (44): the equation has no separate constant, so they are not deviations from an omitted tender category. The omitted firm category pools A and B, and the omitted sector is Buildings. Project controls (BOQ item count, advertised duration, procuring-entity type, the lagged material-price shock, and their missing-value indicators) are estimated and reported in Appendix E.

8 Counterfactual policies

I use the model to conduct counterfactuals that address the main research questions. I first quantify the dynamic incentives created by the category ladder by removing the future market access that winning provides. To do this, I replace the existing promotion rule with a random benchmark and rules based on awarded rather than completed work. I then compare bids, prices, and the allocation of contracts. Second, I evaluate the cost of the screening to the government by loosening the eligibility rule and letting firms from categories below the tender's requirement participate.

Each counterfactual replays the same 1,319 auctions in the sample among the set of active firms in the sample. Across replays, the participation policy is held fixed, but a different set of bidders participate in each auction. Firms' states — backlog, paid experience, and category —

Figure 7: Decomposing bids by firm type

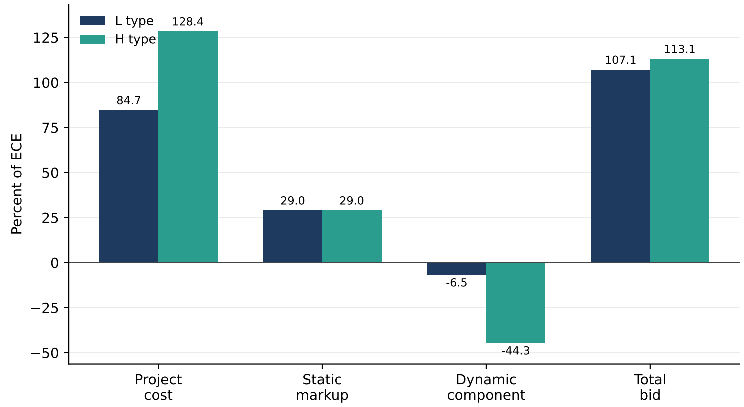

Notes: Baseline specification, all 2,982 supported bids, unweighted means over paired type evaluations, in percent of the engineering estimate. Each type's bid is constructed from the fitted cost equation on the same tender and public state, holding the static markup at its observed-auction value, so rivals' bidding is unchanged. The figure compares otherwise identical firms that differ only in execution type.

Table 3: The dynamic component of bids, by bidder category and type

		High type		Low type	
Bidder category	Bids	Gain (M RWF)	% of ECE	Gain (M RWF)	% of ECE
A	215	324.0	12.4	22.6	1.6
B	93	505.9	64.2	30.4	3.2
C	531	174.1	51.1	8.6	2.5
D	636	104.6	49.8	8.5	8.9
E	210	56.8	75.7	3.0	6.5
F	1,297	9.1	38.8	1.9	8.4
All	2,982	107.2	46.7	2.5	7.7

Notes: The gain columns report the median continuation gain from winning, ΔV, in million RWF, by bidder category and type. Median of the discounted continuation-value wedge $\beta\,\Delta V$ divided by the engineering cost estimate, by the bidder's registration category. Sample includes 2,982 bids, with each evaluated as both types at the point estimates.

evolve with the simulated awards, clearance, and promotion. I re-solve each firm's value function under the counterfactual transition or eligibility policy, and shift bids by the implied change in the continuation value of winning (holding the estimated cost distributions and the participation policy fixed). I simulate 20 replications of the sample and compare each counterfactual replication with the status-quo replication drawn with the same random numbers, so that differences reflect the policy rather than simulation noise.[25]

The main outcomes of interest are procurement prices, the allocation of contracts across firm types, and expected completion time. Prices are measured by the value-weighted ratio of winning bids to the ECE, which equals 0.933 under the status quo, and by the total cost of the lots awarded both in the counterfactual and in the status quo. Allocation is measured by the share of awarded value won by high-type firms, 80.8 percent under the status quo. For a project of value z awarded

[25] More draws are in progress.

to a type-θ firm with backlog B, I measure expected completion time by

$$\mathcal{T} = \sum_{t=0}^{\infty} \varrho_t, \qquad \varrho_{t+1} = \varrho_t[1 - \bar{\xi}(X_t, \theta)], \qquad X_{t+1} = [1 - \bar{\xi}(X_t, \theta)]X_t,$$

where ϱ_t is the remaining payment share, $X_0 = B + z$, and $\bar{\xi}(X, \theta)$ is the estimated mean portfolio clearing rate. Thus $\mathcal{T}$ is the area under the expected remaining payment path; it depends on the winner's type and on her backlog at award. Under the status quo it equals 6.38 quarters on average.

8.1 Removing future market access gains

The dynamic incentive exists because promotion rewards completed and paid work. To quantify it, I replace the RPPA's promotion rule with three alternatives that leave the ladder in place but change what promotion rewards. Under the first, promotion is random within each category, at the category's status-quo promotion rate, so that winning no longer moves a firm up the ladder. Under the second and third, promotion depends on awarded rather than paid value: the rate-matched rule promotes on cumulative awarded value through a quantile map that reproduces the status quo's promotion rate, and the legal-threshold rule promotes a firm once her awarded value crosses the thresholds that the RPPA applies to paid experience. The latter two rules are the ones a government that could not track execution would plausibly adopt, since awards are observed at the time of contracting and payments are not.

Random promotion answers directly how much of the value of winning comes from the ladder. Under that rule, winning still adds a contract to the firm's backlog and paid experience, but no longer moves her up the ladder, so the gain from winning that survives is the value of the work itself, and the remainder is the value of climbing. For high-type firms, the median continuation gain is 45.8 percent of the bid under the status quo, of which only 2.1 points survive under random promotion; the remaining 43.7 points, or 95 percent of the gain, come from the prospect of moving up a category. For low-type firms, the gain is 6.7 percent of the bid, of which 0.8 points is the work itself and 5.9 points the ladder. The split is exact rather than a decomposition, because random promotion is the one rule that severs winning from climbing while leaving everything else in the firm's problem unchanged. Figure 8 reports the same markdown under each rule. Promotion on awarded value at the status quo's rate preserves most of the incentive (34.1 percent of the bid for the high type), since awarded and paid value are closely related for the firms that win. Promotion at the legal thresholds on awarded value raises the markdown above the status quo (59.3 percent), because awarded value crosses the thresholds sooner than paid value does, so that promotion becomes more frequent: 140 promotions per replication against 113 under the status quo.

Figure 9 reports the effects of each rule on prices and on the allocation of contracts. Removing the incentive raises prices: under random promotion, the value-weighted price index rises by 8.2 percent, and the total cost of the lots awarded both under the rule and under the status quo rises by 8.0 percent. Thus the current regime is, at least in part, self-financing: the government pays lower prices today precisely because the prospect of future market access makes firms mark down their bids. The allocation of current contracts, by contrast, moves little. The high-type share of awarded value falls by 1.0 percentage points, and its share of the value awarded in categories A and B, 92.2 percent under the status quo, falls by 1.5 points. Where the ladder's screening shows up is in who gets promoted. Promotions fall from 113 per replication to 91, and the share of them that go to high-type firms falls from 49 to 32 percent, so that the composition of the upper categories deteriorates over time: at the end of the sample, the high-type share of firms registered in category A falls from 71 to 64 percent, and in category B from 46 to 38 percent (Figure 10). The two awarded-value rules differ. The rate-matched rule is close to a wash on every outcome, which is consistent with awarded value tracking paid value closely at the same promotion rate. The legal-threshold rule promotes more firms and, because the incentive it creates is larger, lowers prices by 4.3 percent; but the additional promotions go disproportionately to low-type firms, the high-type share of awarded value falls by 0.9 points, and the high-type share of category B firms falls from 46 to 38 percent (Figure 11). The change in expected completion time under these rules is small, and its sign depends on how lots are weighted, so I do not emphasize it.[26]

Figure 8: The dynamic markdown in bids under each promotion rule

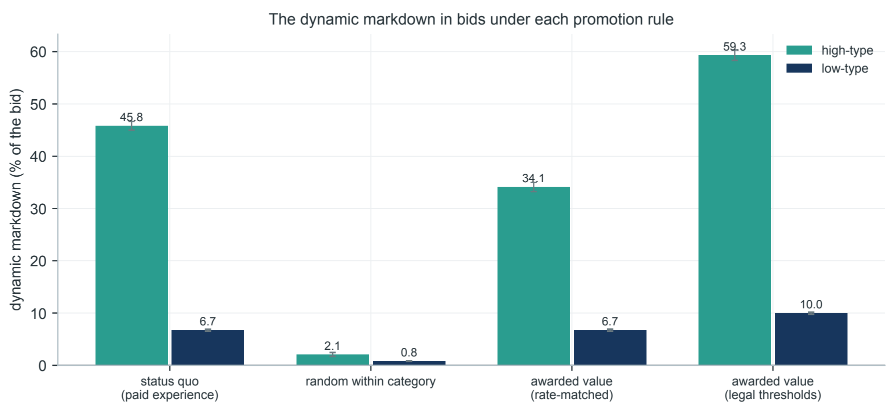

Notes: Median across bids of the discounted continuation gain from winning as a percentage of the bid, by firm type, under the status quo (promotion on paid experience), random promotion within category at the status quo's category-level rate, promotion on awarded value at the status quo's promotion rate, and promotion on awarded value at the legal thresholds. Whiskers are two Monte Carlo standard errors across 20 replications.

8.2 The cost of the screening: loosening eligibility

I now discuss the costs of the eligibility ladder, by counterfactually varying the eligibility requirement while keeping the award rule unchanged. I use the government's own definitions for firm categories to simulate allowing firms below the tender requirement to become eligible for projects above their category. I consider three degrees of loosening. In the first, I admit firms one category below to bid on projects at or below the within-category median engineering estimate of projects one category above; for instance, a category D firm becomes eligible for the smaller half of category C projects.

[26]Weighting lots by the winning bid, random promotion lowers expected completion time by 0.29 quarters; weighting them equally it raises it by 0.15 quarters. The rule moves lots toward slower firms but also toward firms with emptier backlogs, and the two effects offset.

Figure 9: Effects of each promotion rule relative to the status quo

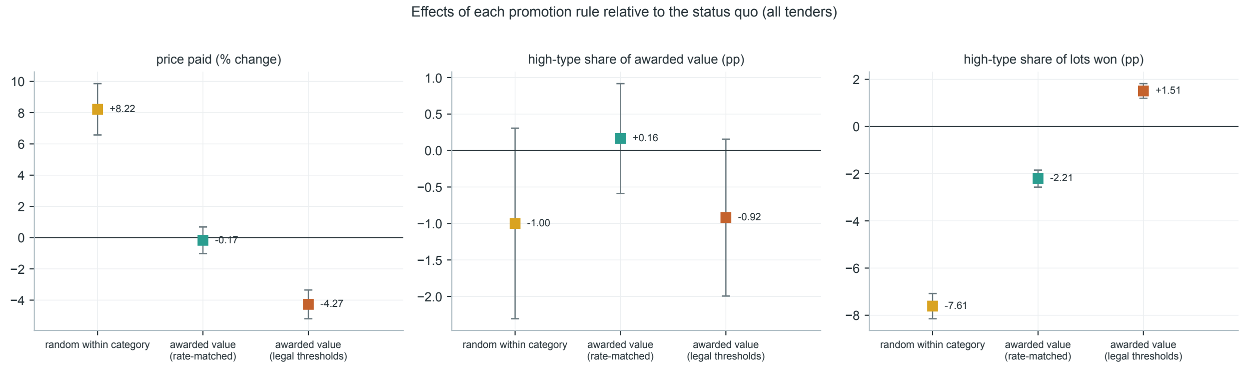

Notes: All tenders. Price paid: percentage change in the value-weighted index of winning bids over the ECE. High-type shares of awarded value and of lots won, in percentage points. Whiskers are two Monte Carlo standard errors across 20 replications.

Figure 10: Composition of the ladder under random promotion

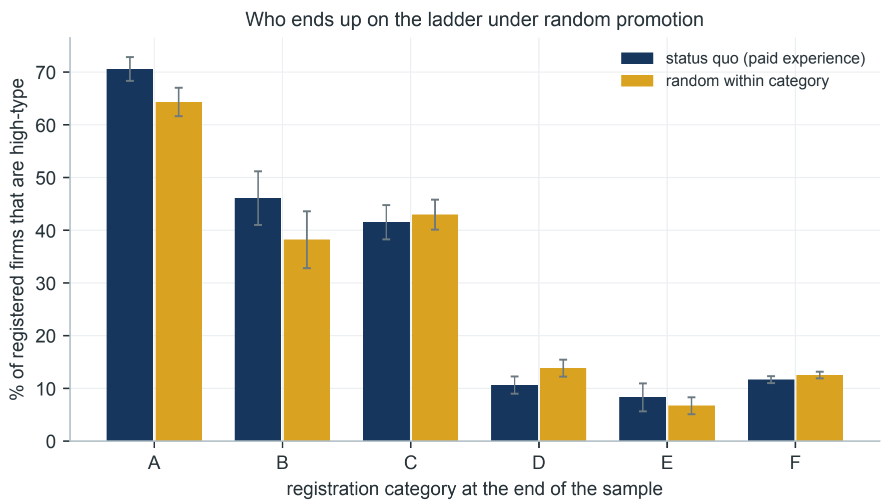

Notes: Share of registered firms that are high-type, by registration category at the end of the sample, under the status quo and under random promotion within category. Whiskers are two Monte Carlo standard errors across 20 replications.

In the second, I allow all firms to bid for projects one category above theirs, and in the third, for projects up to two categories above theirs. Since a category-F tender has no category below it, these counterfactuals apply to tenders of categories E to A, on which the status-quo price index is 0.943, the high-type share of awarded value is 83.5 percent, and expected completion time is 6.47 quarters. Table 4 reports the results, and Appendix Figure 16 plots them along the three degrees of loosening.

Loosening eligibility lowers prices, and the more it is loosened, the more prices fall. Admitting firms one category below to the smaller half of projects lowers the total cost of the lots awarded in both regimes by 2.4 percent, and the price of the median lot by 1.9 percent; the newly eligible firms win 8.5 percent of awarded value. Admitting them to all projects one category above lowers total cost by 6.3 percent and the median lot's price by 8.2 percent, with newly eligible firms winning 21.4 percent of awarded value; admitting firms up to two categories above lowers total cost by 7.9 percent

Figure 11: Composition of the ladder under promotion on awarded value at the legal thresholds

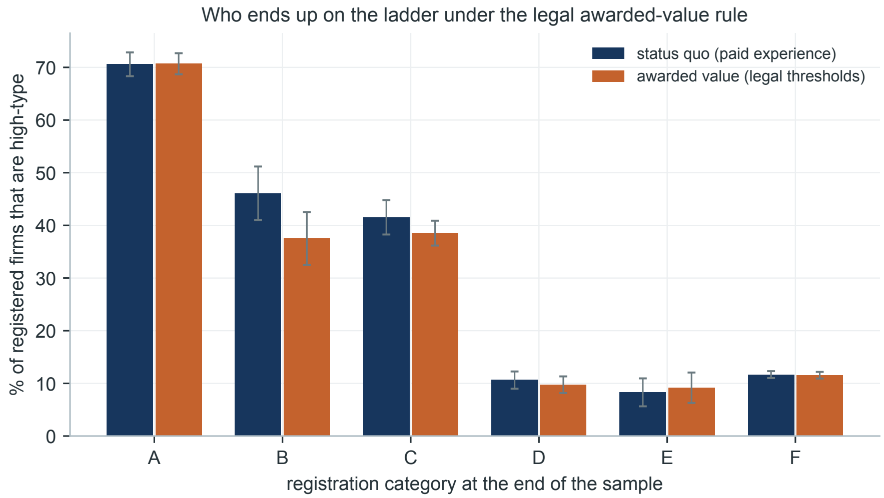

Notes: Share of registered firms that are high-type, by registration category at the end of the sample, under the status quo and under promotion on awarded value at the legal thresholds. Whiskers are two Monte Carlo standard errors across 20 replications.

and the median lot's price by 20.4 percent. The price reductions come from added competition on the projects opened: submissions per lot rise from 1.65 under the status quo to 2.4 when one category is opened, and the newly eligible firms bid aggressively because they carry light backlogs (0.3 billion RWF at award, against 2.8 billion for the incumbents they displace).

The cost of loosening is the composition of winners. The share of awarded value won by high-type firms falls by 1.8, 4.6 and 8.1 percentage points at the three degrees of loosening, and the share of the value awarded in categories A and B that high-type firms win, 92 percent under the status quo, falls by 1.3, 2.5 and 7.9 points. Expected completion time rises accordingly, by 0.08, 0.18 and 0.30 quarters, or roughly one, two and four weeks; the first of these is not distinguishable from zero, and the differences across degrees of loosening are within simulation error, so the evidence supports the direction of the effect more than its exact profile. The newly eligible winners are slower on average (their high-type share is 42 percent, against 60 percent among the incumbents they replace), but they also carry emptier backlogs, which partly offsets their slower execution under the estimated clearing process. Thus the two sides of the screen line up as the model predicts: a modest, targeted loosening of the eligibility rule is nearly free, lowering prices with no measurable change in completion time, while opening tenders to all firms from lower categories buys larger price reductions at the cost of slower execution and a shift of large projects toward low-type firms.

9 Conclusion

This paper studies Rwanda's public procurement market for Works contracts, which imposes eligibility rules that condition access to larger contracts on the completion of smaller ones. Payment records reveal persistent differences in the rate at which firms finish awarded work. These differences

Table 4: Opening tenders to lower categories: effects relative to the status quo

Categories opened	Total cost (%)	Median lot (%)	High-type share (pp)	Newly eligible share (%)	Delay (quarters)
0.5	−2.38 (0.28)	−1.93	−1.81	8.5	+0.083 (0.064)
1	−6.30 (0.47)	−8.23	−4.56	21.4	+0.176 (0.067)
1.5	−6.02 (0.64)	−12.92	−4.97	26.7	+0.119 (0.102)
2	−7.94 (0.78)	−20.38	−8.11	36.8	+0.297 (0.193)

Notes: Tenders of categories E to A. 0.5 opens the smaller half of each category's tenders (estimate at or below the median) to firms one category below; 1.5 opens all tenders one category below and the smaller half two below. Total cost: change in the sum of winning bids on lots awarded in both regimes. Median lot: median change in the winning bid over the ECE on those lots. High-type share: change in the high-type share of awarded value (status quo 83.5%). Newly eligible share: share of awarded value won by newly eligible firms. Delay: change in bid-weighted expected completion time (status quo 6.47 quarters). Monte Carlo standard errors over 20 replications in parentheses.

do not enter bid evaluation directly. Instead, faster-clearing firms accumulate paid experience more quickly and move up the category ladder sooner. Therefore, the ladder selects faster firms into larger tenders without requiring the government to observe their type. The model estimates indicate that high-type firms clear their portfolios about twice as fast and place far more value on winning. This larger continuation value produces a larger dynamic bid markdown, which absorbs most, but not all, of high-type firms' higher costs. On net, high-type firms bid slightly above otherwise similar low-type firms facing the same tender.

This screening comes at the cost of competition. Loosening eligibility restrictions moderately, by allowing firms from one category below the tender requirement to participate, reduces procurement prices by 6 percent, but shifts about 5 percent of awarded value toward slower firms and increases expected completion time by 0.2 quarters. A more limited expansion, limited to smaller tenders in the next category, reduces prices by 2 percent without significantly increasing delay.

References

Victor Aguirregabiria and Pedro Mira. Sequential estimation of dynamic discrete games. *Econometrica*, 75(1):1–53, 2007.

Victor Aguirregabiria and Junichi Suzuki. Identification and counterfactuals in dynamic models of market entry and exit. *Quantitative Marketing and Economics*, 12(3):267–304, 2014. doi: 10.1007/s11129-014-9147-5.

Pasha Andreyanov, Francesco Decarolis, Riccardo Pacini, and Giancarlo Spagnolo. Past performance and procurement outcomes. Working paper, SSRN No. 4929595, 2024.

Peter Arcidiacono and John Bailey Jones. Finite mixture distributions, sequential likelihood and the EM algorithm. *Econometrica*, 71(3):933–946, 2003.

Peter Arcidiacono and Robert A. Miller. Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica*, 79(6):1823–1867, 2011.

Susan Athey and Philip A. Haile. Nonparametric approaches to auctions. In James J. Heckman and Edward E. Leamer, editors, *Handbook of Econometrics*, volume 6A, chapter 60, pages 3847–3965. Elsevier, 2007.

Matthew Backus and Greg Lewis. Dynamic demand estimation in auction markets. *Review of Economic Studies*, 92(2):837–872, 2025. doi: 10.1093/restud/rdae031.

Patrick Bajari and Steven Tadelis. Incentives versus transaction costs: A theory of procurement contracts. *The RAND Journal of Economics*, 32(3):387–407, 2001.

Patrick Bajari, C. Lanier Benkard, and Jonathan Levin. Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370, 2007.

Patrick Bajari, Stephanie Houghton, and Steven Tadelis. Bidding for incomplete contracts: An empirical analysis of adaptation costs. *American Economic Review*, 104(4):1288–1319, 2014.

Jorge Balat. Highway procurement and the stimulus package: Identification and estimation of dynamic auctions with unobserved heterogeneity. Working paper, University of Texas at Austin, 2015.

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

Rodrigo Carril, Andres Gonzalez-Lira, and Michael S. Walker. Competition under incomplete contracts and the design of procurement policies. *American Economic Review*, 116(2):535–581, 2026.

Matias D. Cattaneo, Richard K. Crump, Max H. Farrell, and Yingjie Feng. On binscatter. *American Economic Review*, 114(5):1488–1514, 2024.

Decio Coviello, Luigi Moretti, Giancarlo Spagnolo, and Paola Valbonesi. Court efficiency and procurement performance. *Scandinavian Journal of Economics*, 120(3):826–858, 2018. doi: 10.1111/sjoe.12225.

Jingyi Cui. Bidding for reputation. Job market paper, Yale University, 2025.

Dakshina G. De Silva, Matthew Gentry, and Pasquale Schiraldi. Dynamic oligopoly in procurement auction markets. Working paper, 2023.

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

Francesco Decarolis, Raymond Fisman, Paolo Pinotti, and Silvia Vannutelli. Rules, discretion, and corruption in procurement: Evidence from Italian government contracting. *Journal of Political Economy Microeconomics*, 3(2):213–254, 2025. doi: 10.1086/732654.

Raphael A. Espinoza and Andrea F. Presbitero. Delays in public investment projects. *International Economics*, 172:297–310, 2022.

Matthew Gentry, Tatiana Komarova, and Pasquale Schiraldi. Preferences and performance in simultaneous first-price auctions: A structural analysis. *The Review of Economic Studies*, 90(2): 852–878, 2023.

Joachim R. Groeger. A study of participation in dynamic auctions. *International Economic Review*, 55(4):1129–1154, 2014.

Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574, 2000.

Peter Hall and Xiao-Hua Zhou. Nonparametric estimation of component distributions in a multivariate mixture. *The Annals of Statistics*, 31(1):201–224, 2003. doi: 10.1214/aos/1046294462.

Oliver Hart, Andrei Shleifer, and Robert W. Vishny. The proper scope of government: Theory and an application to prisons. *Quarterly Journal of Economics*, 112(4):1127–1161, 1997. doi: 10.1162/003355300555448.

V. Joseph Hotz and Robert A. Miller. Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, 60(3):497–529, 1993.

Bar Ifrach and Gabriel Y. Weintraub. A framework for dynamic oligopoly in concentrated industries. *The Review of Economic Studies*, 84(3):1106–1150, 2017.

Przemysław Jeziorski and Elena Krasnokutskaya. Dynamic auction environment with subcontracting. *The RAND Journal of Economics*, 47(4):751–791, 2016.

Mireia Jofre-Bonet and Martin Pesendorfer. Estimation of a dynamic auction game. *Econometrica*, 71(5):1443–1489, 2003.

Myrto Kalouptsidi, Paul T. Scott, and Eduardo Souza-Rodrigues. On the non-identification of counterfactuals in dynamic discrete games. *International Journal of Industrial Organization*, 50: 362–371, 2017.

Myrto Kalouptsidi, Paul T. Scott, and Eduardo Souza-Rodrigues. Identification of counterfactuals in dynamic discrete choice models. *Quantitative Economics*, 12(2):351–403, 2021. doi: 10.3982/QE1253.

Hiroyuki Kasahara and Katsumi Shimotsu. Nonparametric identification of finite mixture models of dynamic discrete choices. *Econometrica*, 77(1):135–175, 2009.

Gregory Lewis and Patrick Bajari. Procurement contracting with time incentives: Theory and evidence. *The Quarterly Journal of Economics*, 126(3):1173–1211, 2011.

Giuseppe Lopomo, Nicola Persico, and Alessandro T. Villa. Optimal procurement with quality concerns. *American Economic Review*, 113(6):1505–1529, 2023. doi: 10.1257/aer.20211437.

Alejandro M. Manelli and Daniel R. Vincent. Optimal procurement mechanisms. *Econometrica*, 63 (3):591–620, 1995.

Yair Mundlak. On the pooling of time series and cross section data. *Econometrica*, 46(1):69–85, 1978.

Aroon Narayanan. Governance of supply relationships: Evidence from Indian manufacturing. Job market paper, MIT, 2025.

Rwanda Public Procurement Authority. *Categorization of Companies Operating in the Field of Building and Civil Engineering Works*. Rwanda Public Procurement Authority, Kigali, Rwanda, 2014.

Rwanda Public Procurement Authority. RPPA annual activity report 2022–2023. Technical report, Rwanda Public Procurement Authority, Kigali, September 2023.

Rwanda Revenue Authority. *RRA Tax Handbook*. Rwanda Revenue Authority, Kigali, Rwanda, 2 edition, 2025. First published 2018; second published 2025.

Nicholas Ryan. Contract enforcement and productive efficiency: Evidence from the bidding and renegotiation of power contracts in India. *Econometrica*, 88(2):383–424, 2020. doi: 10.3982/ECTA17041.

Nicholas Ryan. Holding up green energy: Counterparty risk in the Indian solar power market. *Econometrica*, 94(3):767–810, 2026. doi: 10.3982/ECTA21084.

Viplav Saini. Endogenous asymmetry in a dynamic procurement auction. *The RAND Journal of Economics*, 43(4):726–760, 2012.

Gabriel Y. Weintraub, C. Lanier Benkard, and Benjamin Van Roy. Markov perfect industry dynamics with many firms. *Econometrica*, 76(6):1375–1411, 2008.

Gabriel Y. Weintraub, C. Lanier Benkard, and Benjamin Van Roy. Computational methods for oblivious equilibrium. *Operations Research*, 58:1247–1265, 2010.

World Bank. Rwanda: Assessment of the public procurement system, volume I: Main report. Methodology for assessing procurement systems (maps) assessment, World Bank, Washington, DC, 2020. URL https://www.mapsinitiative.org/content/dam/maps-initiative/en/assessments/rwanda/maps-assessment-rwanda-main-report.pdf.

A Payment Measurement and Validation

This appendix presents the details behind the construction of the payment-based variables defined in Section 2.5: the component of firm payments attributed to contracts outside the Works portfolio, the contract schedules used for the scheduled-flow measures in the empirical evidence, the persistent firm clearing component, and the normalizations that enter the state in the structural model.

A.1 Observed contract portfolios

Let j index public contracts held by firm i. I focus on the set of contracts $\mathcal{W}_i$, which contains signed Umucyo Works contracts. The OCDS data supplies another contract set containing goods, services, and consulting contracts held by the same firms. OCDS contracts are used only to partition firm-level payments; they are not added to the Works backlog studied in the paper.

For each observed contract, let z_j^c denote its VAT-exclusive signed value, and let s_j and e_j denote the quarter containing its recorded start and end dates. I use the signed value to construct the scheduled duration and quarterly flow,

$$d_j^c = e_j - s_j + 1, \qquad r_j = \frac{z_j^c}{d_j^c}.$$

When an end date is missing or precedes the start date, I impute duration using the median valid duration within contract type (sector-category combination). This affects no Works contract and only 64 of the 2,283 OCDS (non-Works) contracts.

A.2 Allocating payments to the Works portfolio

The government withholdings variable in VAT records reports the total payments from public entities, Y_{it}^{tot}, at the firm-quarter level. Because of this, I do not observe contract-level payments. This subsection defines how I attribute payments to Works contracts and portfolios.

I first allocate payments to the firm's observed OCDS non-Works contracts using their values and contract dates. Let Y_{it}^{rem} denote the payment remaining unassigned in this pass, and let t_i^W denote the quarter in which the firm's first observed Works contract begins. Let

$$U_{it} = \min\left\{Y_{it}^{\text{tot}}, \text{median}_{\tau<t_i^W}\, Y_{i\tau}^{\text{rem}}\right\},$$

defining the *unaccounted component* to equal to the "typical" payment the firm received before it held any observed Works contract. For firms with fewer than four such quarters, I set $U_{it} = 0$, attributing their public payments entirely to Works.

The payment that can be attributed to Works is therefore

$$Y_{it}^{\text{tot}} - U_{it}.$$

I attribute this payment to the firm's aggregate Works portfolio rather than to particular contracts. Let B_{it} be the VAT-exclusive signed-contract value of previously awarded Works that remains to be cleared at the beginning of quarter t. For each contract ℓ, let z_ℓ^c denote its signed value net of VAT, and let

$$\Delta W_{it} = \sum_{\ell \in \mathcal{W}_{it}} z_\ell^c$$

be the VAT-exclusive signed-contract value of Works awarded to firm i in that quarter, where $\mathcal{W}_{it}$ is the set of those awards.The portfolio available to receive payment is

$$X_{it} = B_{it} + \Delta W_{it}.$$

I define Works payments as

$$Y_{it}^W = \min\{Y_{it}^{\text{tot}} - U_{it}, X_{it}\}.$$

Thus, Works payments cannot exceed either the firm's payment available for Works or the value remaining in its Works portfolio. Any available payment above X_{it} remains unallocated and is not carried forward to future awards.

The corresponding backlog evolves as

$$B_{i,t+1} = X_{it} - Y_{it}^W = B_{it} + \Delta W_{it} - Y_{it}^W.$$

Thus, realized commitments and payments are measured on a common VAT-exclusive contractual basis.

Portfolio clearance and paid experience. For firm-quarters with a positive Works portfolio, I measure quarterly portfolio clearance as

$$\xi_{it} = \frac{Y_{it}^W}{X_{it}}, \qquad X_{it} > 0.$$

Because Works payments are capped by the available portfolio, $\xi_{it} \in [0, 1]$. A value of zero indicates that none of the portfolio is cleared in quarter t, while a value of one indicates that the entire portfolio is cleared. Firm-quarters with $X_{it} = 0$ are outside the execution risk set and ξ_{it} is set to missing.

As firms clear their Works portfolios, they accumulate *paid experience.* Let W_{i0} be the cumulative VAT-exclusive signed-contract value of awards before the accounting panel, so that $W_{it} = W_{i0} + \sum_{\tau=0}^{t-1} \Delta W_{i\tau}$, and set $P_{i0} = W_{i0} - B_{i0}$. At the beginning of quarter t, paid experience is

$$P_{it} = P_{i0} + \sum_{\tau=0}^{t-1} Y_{i\tau}^W = W_{it} - B_{it}.$$

For a firm observed from its first Works award, $P_{i0} = 0$.

A.3 Reduced-form execution measures

All firm-level execution measures in Section 3.2 are built from the quarterly portfolio clearing rate of Section 2.5,

$$\xi_{it} = \frac{Y_{it}^W}{X_{it}},$$

Works payments received in quarter t divided by the portfolio available to receive payment, for firm-quarters with $X_{it} > 0$. Contract schedules enter these measures only as sample filters and workload controls, never as a denominator.

Predetermined clearing score. I residualize ξ_{it} on $\log(1 + X_{it})$, an indicator for a new award in t, indicators for the number of active Works contracts, and calendar-quarter fixed effects. For each firm I average the residual over its first three quarters with a positive portfolio and standardize the average across firms; firms with fewer than three such quarters are dropped, leaving 319 firms. The score is fixed at the end of a firm's third quarter with positive portfolio.

Explaining variation in clearing. Table 9 uses ξ_{it} itself on firm-quarters with a positive portfolio and at least 0.005 billion RWF of scheduled Works execution, retaining firms observed in at least three such quarters (1,784 firm-quarters, 233 firms).

Execution quality and tender evaluation. Table 10 uses the predetermined clearing score.

A.4 Normalizations entering the state

For the reduced-form regressions the backlog level is divided by the scale of the firm's current category and capped at six; I report both an indicator for any positive backlog and the standardized continuous state. Scheduled backlog, constructed from the administrative contract windows, serves as a comparison measure.

Paid experience enters the promotion problem relative to the requirement stated in the 2014 RPPA categorization manual,

$$\widetilde{P}_{it}(m, \kappa) = \frac{P_{it}}{T_{m,\kappa}},$$

where $T_{m,\kappa}$ is the stated valule requirement for sector m and origin category κ.

A.5 Cross-validating payment-measured Works execution

The measure Y_{it}^W assigns the portion of firm i's government payments in quarter t attributable to its Works portfolio observed with the *Umucyo* data. I assess it using independently collected project reports, which record physical completion (as a percent of the project) and the date the report was produced. Project reports are produced by third-party contractors hired to monitor the project execution and submitted to the relevant procuring entity. I use the date to measure delay in completion relative to scheduled contract start and end dates. I compare portfolio and directly attributable project payments, where I can observe (linking them using the project-specific code).

A.5.1 Payments around independently recorded final reports

The date of the Final Report provides an external measure of physical completion. Panel A of Figure 12 uses the cleanest sample: 12 contracts for which the firm's payment flow can be attributed directly to one contract. (There are no contracts immediately before or after). In the final-report quarter, the 25th percentile, median, and 75th percentile of cumulative payment are 8, 24, and 65 percent of contract value. One quarter later they are 73, 100, and 100 percent. Thus, the distribution shifts sharply after the final report: in the median project, the remaining amount is recorded as paid in the following quarter.

Panel B expands the sample to all contracts with a matched final report and allocates the exact firm-quarter Y_{it}^{W} flow across the firm's open Works contracts in proportion to contract value, subject to each contract's unpaid balance. The average cumulative paid share rises from 33.8 percent two quarters before the report to 56.9 percent in the report quarter, 65.5 percent one quarter later, and 69.0 percent two quarters later. This panel confirms that there is a rapid increase in payments around the date of the final report completion.

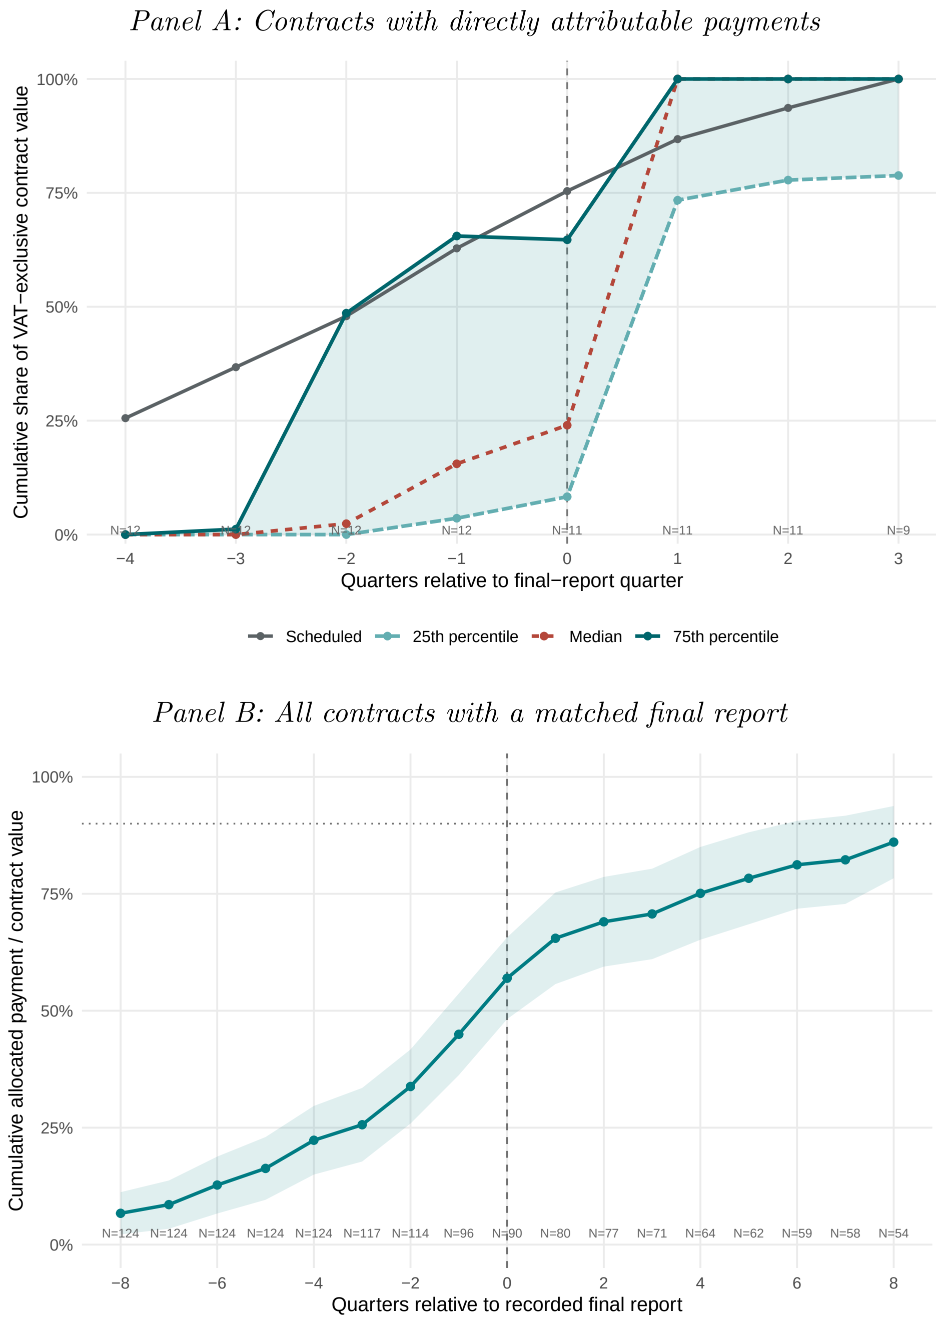

Figure 12: Payment clearing around a recorded final report

Notes: Event time is measured in quarters relative to the date of the Final Report. Panel A shows cross-sectional percentiles of cumulative project-attributable VAT payments divided by VAT-exclusive contract value, capped at one. Eleven of the 12 contracts are observed in the report quarter because one report occurs after the end of the VAT panel. Panel B shows mean cumulative contract-level Y^W after proportionally allocating the exact firm-quarter portfolio flow across started and unpaid Works contracts. In total, 124 contracts have a Final Report, but 34 have reports that end after the end of the VAT panel and so are excluded. Shaded bands are 95% confidence intervals based on standard errors clustered by firm.

Table 5: Reported project completion and payment clearing

Payment at report	All reports		Non-final only				
Payment after signed end				Unpaid share		Not 90% paid	
Specification	Raw	Controls		Raw	Controls	Raw	Controls
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Reported completion share	0.641***	0.510**	1.04***				
	(0.157)	(0.221)	(0.152)				
Report-coded delay				-0.022	0.016	0.139	0.133
				(0.070)	(0.071)	(0.098)	(0.095)
Mean dependent variable	0.551	0.551	0.509	0.397	0.397	0.702	0.702
Contract value and duration controls	No	Yes	No	No	Yes	No	Yes
SE clustered by	Firm	Firm	Firm	Firm	Firm	Firm	Firm
Sector fixed effects		✓			✓		✓
Report-year fixed effects		✓					
Start-year fixed effects					✓		✓
Observations	108	108	26	121	121	121	121
Firms	62	62	20	65	65	65	65

Notes: Each contract uses its last recorded report, which may be Final or non-Final (i.e., progress report). Payment at report is cumulative Y^W divided by contract value. End-point outcomes are measured one quarter after the signed contract end. Controlled specifications include contract value, scheduled duration, sector fixed effects, and report- or start-year fixed effects. Standard errors are clustered by firm.

A.5.2 Reported completion and payment clearing

Table 5 uses the last available report for each matched contract. A 10 percentage point (pp) increase in reported physical completion is associated with a 5.1pp increase in cumulative Y^W, conditional on contract value, scheduled duration, sector, and report year. Contracts coded as delayed are 13.3pp more likely not to have reached 90 percent payment one quarter after the scheduled end, although this estimate is not statistically significant.

A.5.3 Consistency across project and portfolio measures

As a final check, Figure 13 compares the portfolio clearing rate with an on-time indicator constructed for the subset of contracts whose payment flows can be attributed to individual projects. Across 264 firms, the correlation between mean portfolio clearing and the share of projects paid at least 90 percent by scheduled end plus one quarter is 0.675. The correlation is 0.651 among the set of 55 firms with at least two attributable projects. Thus, the portfolio measure preserves the project-level ranking where both are observed.

In summary, payments increase around independent reports of completion, physical progress predicts cumulative payments, and the portfolio measure agrees with project outcomes where direct attribution is possible. Together, these validation checks support interpreting Y^W as the Works projects execution measure.

Figure 13: Project-level on-time payment and portfolio clearing

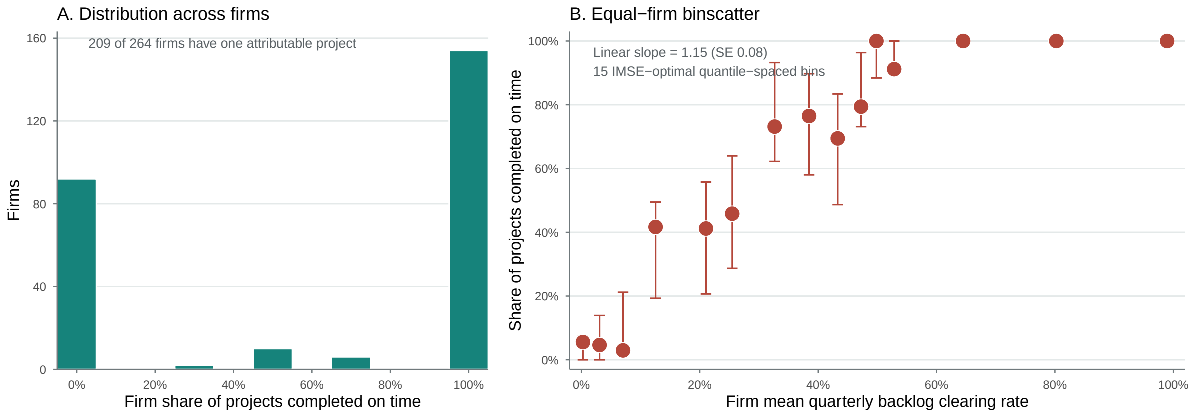

Notes: The sample contains 264 firms, 332 attributable projects, and 2,427 positive-backlog firm-quarters. A project is on time if cumulative Y^W reaches at least 90 percent of VAT-exclusive contract value by scheduled end plus one quarter. Portfolio clearing is the firm's mean quarterly clearing rate across positive-backlog quarters. Each firm receives equal weight. Because 209 of 264 firms have a single attributable project, the firm-level outcome is essentially binary (Panel A); Panel B's bin means are therefore the share of firms in each bin whose projects are on time. Panel B reports the IMSE-optimal number of quantile-spaced bins (15) in the firm mean portfolio clearing rate using (Cattaneo et al., 2024) in R; whiskers are 95% CIs.

B Supporting Empirical Evidence

B.1 Firm categorization requirements

The RPPA category assigned to a firm reflects four documented forms of capacity: average construction turnover over the previous five years, executed-contract references, owned equipment, and permanent senior staff. Table 6 reports the turnover and past executed-tender requirements for the three sectors studied in the paper: Buildings, Roads and Bridges, and Drinking Water Supply. These requirements applied to firms requesting a category during the sample period.

Table 6: Partial requirements for firm categorization

Category	Tender value	Five-year average turnover	Executed tenders: # and value of each
A	$>$ 2 bn	$>$ 1.5 bn	2; each $>$ 1.2 bn
B	1.5–2 bn	1.0–1.5 bn*	2; each 0.9–1.2 bn*
C	0.8–1.5 bn	0.6–1.0 bn*	2; each 0.48–0.90 bn*
D	0.3–0.8 bn	0.25–0.60 bn	1; 0.18–0.48 bn
E	0.1–0.3 bn	0.075–0.25 bn	1; 0.06–0.18 bn
F	$<$ 0.1 bn	None	None

Notes: Values are in billions of RWF. Stars identify the exceptions summarized here. For Drinking Water Supply, the turnover requirement is 1.1–1.5 billion RWF in category B and 0.6–1.1 billion RWF in category C; the category-C executed-tender requirement is 0.50–0.90 billion RWF. Executed tenders must fall within the previous five years. Roads and Bridges divides categories A and B into A1/A2 and B1/B2. A1 and B1 require asphalt-road project references, whereas A2 and B2 require road-construction project references; their numerical thresholds are the same. The category-B reference cell is marked to flag this distinction. The manual also specifies asphalt-specific equipment and staff experience for A1 and B1, and owned equipment and permanent senior staff in all sectors. Source: Rwanda Public Procurement Authority (2014).

B.2 Category eligibility binds in the bidding data

Table 7 shows results whether the tender eligibility requirements are enforced. The sample includes all National and International Competitive Bidding, NCB and ICB respectively, bids in the three main sectors, Buildings, Roads and Bridges, and Drinking Water Supply, before October 2023. In the combined sample, 99.5 percent of bids come from firms in the tender category or no more than two categories above it, and the same is true of winning bids.

Table 7: Category eligibility among observed bidders

Sector	Bids	Lots	Firms	Eligible bids	Same cat.	One above	Two above
All main sectors	7,034	1,704	766	99.5	67.0	20.1	12.4
Buildings	5,190	1,172	696	99.6	71.1	16.9	11.5
Roads and Bridges	850	287	152	98.7	53.3	27.3	18.1
Drinking Water Supply	994	245	184	99.6	56.8	30.8	12.0

Notes: NCB and ICB bids in the three main Works sectors before October 2023, restricting to observations with classified tender and firm categories. "Eligible bids" is the share of bids in which the firm is in the tender category or one or two categories above it. The final three columns decompose the bidder–tender category relationship.

B.3 Category progression in the structural sample

Figure 14 compares the first and last observed category of each firm-sector pair in the structural sample. Most pairs remain in their initial category, while movement through the ladder is almost entirely upward.

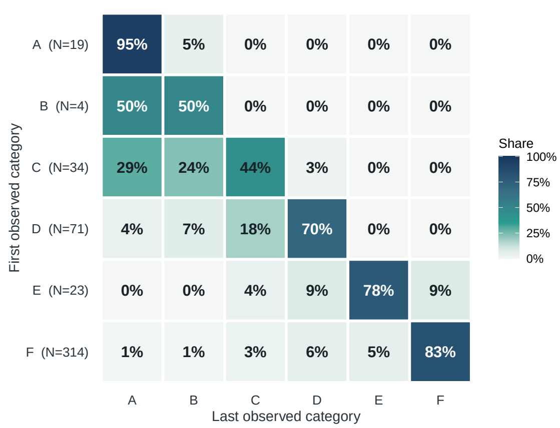

Figure 14: Category progression in the structural sample

Notes: The unit is a firm-sector pair in the structural sample. Rows show the first observed category and columns show the last observed category through October 2023. Cells report row shares; row labels report the number of firm-sector pairs. The sample contains 465 firm-sector pairs held by 394 firms. Categories A and B combine the Roads and Bridges subcategories A1/A2 and B1/B2.

B.4 Paid experience and approval of category-upgrade requests

Table 8 compares firms requesting the same sector-specific category transition in the same quarter. Applicants with more cumulative paid Works experience are more likely to be approved. Meeting the stated value requirement for the requested category is associated with a 10.3 percentage point higher approval probability, relative to an 88.5 percent mean.

Table 8: Paid experience and approval of category-upgrade requests

Outcome	Category-upgrade request approved			
	(1)	(2)	(3)	(4)
Log(1 + paid experience / target requirement)	0.056**	0.114**		
	(0.028)	(0.054)		
Paid experience meets target requirement			0.103***	0.151*
			(0.036)	(0.084)
Firm experience controls	No	Yes	No	Yes
Request-quarter × sector × origin × target FE	Yes	Yes	Yes	Yes
SE clustered by	Firm	Firm	Firm	Firm
Observations	338	338	338	338
Mean dependent variable	0.885	0.885	0.885	0.885
Request events	201	201	201	201
Firms	164	164	164	164

Notes: The unit of observation is a sector-specific category-upgrade request. Paid experience is cumulative Works payments, P_{it}, measured by the quarter before the request. The target requirement is the stated value requirement for the requested sector and category. Firm experience controls include firm age, prior observed Works contracts, and prior category requests. Standard errors are clustered by firm.

B.5 Sources of variation in portfolio clearing

Table 9 reports the explanatory power of successively richer sets of covariates for the quarterly portfolio clearing rate. I begin with only including quarter fixed effects. I add "Portfolio workload" controls, which include the value of the firm's existing portfolio, the value of new awards, and the number of active contracts. This substantially increases the Adjusted R^2 to 0.145. I then add Procuring entity (PE) composition covariates, which include the share of the firm's portfolio value that is contracted with the PE, the share of contracts located in the firm's own district, indicators for the number of procuring entities that the firm is contracting with, and a HHI index of value across procuring entities. These covariates aim to summarize the newness of the firm's relationships with the PEs and how concentrated the firm's contracts are. Next, I add prior PE payment controls, which include the value-weighted mean of each PE's earlier clearing rates for contracts between the PE and *other* firms. These are payer-side controls that aim to control for PE payment delays to firms. Despite the richness in these controls, the adjusted R^2 only increases to 0.149. Adding firm-level covariates, including category, prior-year VAT sales, paid experience, prior Works contracts, and number of quarters in the market, increases the Adjusted R^2 modestly to 0.193. Finally, adding firm fixed effects more than doubles the adjusted R^2 to 0.431, providing suggestive evidence that the clearing rate contains an important unobserved firm-specific component.

Table 9: Explaining variation in portfolio clearing

Specification	Adjusted R^2
Quarter fixed effects	0.034
+ Portfolio workload	0.145
+ PE composition	0.148
+ Prior PE payment	0.149
+ Firm size and experience	0.193
+ Firm fixed effects	0.431
Observations	1,784
Firms	233

Notes: The outcome is the quarterly portfolio clearing rate $\xi_{it} = Y_{it}^W / X_{it}$. The sample contains firm-quarters with a positive Works portfolio and at least 5 million RWF of scheduled Works execution, retaining firms observed in at least three such quarters.

B.6 Payment-measured execution quality and tender evaluation

Table 10 tests whether the portfolio-clearing component predicts disqualification process and bid award. I use bid-level Linear Probability Model regressions to test whether the predetermined portfolio clearing score predicts disqualification for any reason, disqualification for technical reasons, and contract award (lowest bid among qualified). For award, I also include bid-rank Fixed Effects. The coefficient for disqualification for any reason is negative and marginally statistically significant, but all other regressions show no evidence that the clearing score (the proxy for firm quality) predicts disqualification or bid award. Conditional on the price ranking that determines award, firms with higher clearing rates are no more likely to win.

Table 10: Payment-measured execution quality and tender evaluation

	Any disq. (1)	Technical disq. (2)	Award (3)	Award (4)
Predetermined clearing score	-0.031* (0.017)	-0.014 (0.012)	-0.027 (0.030)	-0.001 (0.006)
Bid-category controls	Yes	Yes	Yes	Yes
Current scheduled backlog	Yes	Yes	Yes	Yes
Tender-lot FE	Yes	Yes	Yes	Yes
Bid-rank FE	No	No	No	Yes
Sample	Entrants	Entrants	Qualified	Qualified
Clustered SEs	Firm	Firm	Firm	Firm
N	3,216	3,216	1,742	1,742
N firms	183	183	161	161

Notes: Bid-level LPMs. The quality proxy is the standardized predetermined portfolio-clearing score of Appendix A.3, fixed at the end of a firm's third positive-portfolio quarter. The regression sample includes bids submitted at or after the quarter in which the bidder's score is calculated. Columns 1 and 2 use all bids, while columns 3 and 4 use bids were not disqualified. Relative to column 3, column 4 includes bid rank Fixed Effects. All specifications absorb tender-lot fixed effects and control for bidder category, the bidder–tender category gap, and current scheduled backlog. Standard errors are clustered by firm.

C Estimation Details

C.1 Detailed State Accounting

This appendix describes the state panel used by the model. Let ΔW_{it} be the total VAT-exclusive signed value of Works contracts awarded to firm i in quarter t, so $\Delta W_{it} = \sum_{\ell \in \mathcal{W}_{it}} z^c_\ell$, and let $\widehat{Y}^W_{it}$ be RRA payments allocated to those Works contracts. Starting from the observed stock B_{it} at the beginning of the quarter, I construct

$$\begin{aligned}
X_{it} &= B_{it} + \Delta W_{it}, \\
Y^W_{it} &= \min\{\widehat{Y}^W_{it}, X_{it}\}, \\
B_{i,t+1} &= X_{it} - Y^W_{it}, \\
W_{i,t+1} &= W_{it} + \Delta W_{it}, \\
P_{i,t+1} &= W_{i,t+1} - B_{i,t+1}.
\end{aligned}$$

The initial cumulative-award stock is the VAT-exclusive signed value of observed awards before the estimation window, subject to $W_{i0} \geq B_{i0}$. I set $P_{i0} = W_{i0} - B_{i0}$; hence the recursion gives

$$P_{it} = P_{i0} + \sum_{\tau=0}^{t-1} Y^W_{i\tau} = W_{it} - B_{it}.$$

Thereafter the recursion uses awards and payments only. I verify the recursion against the independently constructed backlog panel in every overlapping firm-quarter.

Constructing the payments this way has two important timing properties. First, an award enters neither B_{it} nor W_{it} before its award quarter. Second, its full VAT-exclusive signed value enters X_{it} before the quarter's clearance draw. Thus, realized commitments and payments use the same VAT-exclusive signed contract value. The continuation value's win–lose calculation is different: because z^c_ℓ is not observed when bids are submitted, it uses the engineering estimate z_ℓ as a predetermined proxy for the prospective award increment, similar to (De Silva et al., 2023). Under this approximation, the continuation gain is fixed with respect to the bid.

C.2 Execution-type likelihood

The three components of (32) are logits in the standardized log outstanding work,

$$x^\xi_{it} = \frac{\log X_{it} - \overline{\log X}}{\mathrm{sd}(\log X)},$$

where the mean and standard deviation are taken over the risk-set quarters, those with $X_{it} > 0$. Execution type enters each component through the single common loading λ^ξ, as in (33). Standardizing $\log X_{it}$ absorbs the mechanical dependence of the cleared share on the size of the stock, so λ^ξ is not identified from firm scale. Restricting the three type effects to a single λ^ξ with this sign pattern is what delivers the stochastic ordering (8), and fixes which mixture component is denoted H.

For each observation in the risk set, i.e., has outstanding work $X_{it} > 5$mill. RWF, the likelihood uses the clearance rate $\xi_{it} \in [0,1]$ from (6). Let $f_{\mathrm{B}}(\cdot; a, b)$ denote the Beta density. Under (32) and (33), the execution component of the firm likelihood is

$$\mathcal{L}_i^{\xi}(\theta) = \prod_{t:X_{it}>0} \begin{cases} p_{0\theta}(X_{it}), & \xi_{it} = 0, \\ [1 - p_{0\theta}(X_{it})] p_{1\theta}(X_{it}), & \xi_{it} = 1, \\ [1 - p_{0\theta}(X_{it})][1 - p_{1\theta}(X_{it})] f_{\mathrm{B}}\Big(\xi_{it}; \mu_\theta^{\xi}(X_{it})\phi, [1 - \mu_\theta^{\xi}(X_{it})]\phi\Big), & 0 < \xi_{it} < 1. \end{cases}$$

The common positive loading λ^{ξ} implies

$$p_{0H}(X) < p_{0L}(X), \qquad p_{1H}(X) > p_{1L}(X), \qquad \mu_H^{\xi}(X) > \mu_L^{\xi}(X).$$

The Beta precision ϕ is common across types. Repeated zero payments, partial payments, and complete clearance, conditional on the total remaining portfolio, therefore all contribute to type classification.

C.3 Sector-specific categories and sector entry

Firm-level total backlog and payment are shared across sectors, while registration is sector-specific. So, experience gained in Buildings can contribute to a higher probability of upgrading in Roads & Bridges. For every active sector m, the RPPA category upgrade policy is

$$\frac{P_{it}}{T_{m,\kappa_{im,t}}}$$

and updates the firm's category in that sector, $\kappa_{im,t}$ using the sector-specific payment threshold $T_{m,\kappa_{im,t}}$. While the thresholds are allowed to be sector-specific, in practice they are almost all equal across sectors (see Table B.1).

C.4 Exogenous policy estimates

This appendix collects the policies the model takes as given rather than recovers from optimizing behavior: the RPPA category transition kernel, the industry entry and exit hazards, and the aggregate state and its law of motion. Each is estimated once, on its own sample, and then held fixed for the remainder of estimation. Bootstrap intervals do not account for the uncertainty in the estimation of these policies. The estimated primitives — the clearance process, the participation policy, and the cost equation — are reported in Appendix E.

C.4.1 RPPA category transitions

The promotion hazard (30) and destination rule (31) are estimated on observations at the firm–sector–quarter–level, excluding category A, which cannot be promoted. The sample used for estimation contains 5,081 firm–sector–quarters covering 357 firms.

Table 11 reports the fitted promotion hazard by origin category and paid-experience ratio. The hazard rises steeply in $\widetilde{P}$ throughout, and it does so smoothly. Crossing the stated requirement raises the promotion probability substantially, but not to probability 1. This is consistent with the legal criteria also requiring contract counts, equipment, and professional staff that I do not observe.

Table 11: Fitted quarterly promotion hazard, by origin category and paid experience

	Paid experience relative to requirement, $\widetilde{P}$						
Origin	0	$(0, 0.5]$	$(0.5, 1]$	$(1, 2]$	$(2, 4]$	$(4, 8]$	> 8
F	0.005	0.006	0.009	0.015	0.029	0.062	0.110
E	0.006	0.007	0.011	0.018	0.034	0.072	0.127
D	0.012	0.015	0.023	0.038	0.071	0.143	0.238
C	0.019	0.023	0.037	0.059	0.108	0.208	0.330
B	0.025	0.030	0.047	0.076	0.136	0.256	0.392

Notes: Quarterly probability of moving to any higher category. Entries are the sector-pooled policy, which is used in the value computation. Sector-specific hazards are also estimated and are used where a sector-specific transition is required. Category A is omitted because it cannot be promoted.

Conditional on promotion, Table 12 reports where firms go. Firms primarily move up to the adjacent category from E and D, while promotions out of F and C are more dispersed. Upgrades from categories E and D are restricted to at most two ranks. Moves from F are unrestricted.

Table 12: Destination shares conditional on promotion

	Destination category				
Origin	E	D	C	B	A
F	0.599	0.237	0.101	0.055	0.009
E		0.928	0.072		
D			0.892	0.108	
C				0.589	0.411
B					1.000

Notes: Shares average the sector-pooled destination distribution across paid-experience bins. Rows sum to one. Blank cells are destinations the policy does not assign positive probability.

C.4.2 Industry entry and exit

Table 13 reports the entry and exit fitted hazards. Exit falls with backlog and, more steeply, with paid experience. Entry is a function of the aggregate state (new firms have no individual state).

These hazards are assumed to be exogenous policy functions and used in the value computation. The model measures value before the survival draw: only surviving firms receive the current auction

Table 13: Industry entry and exit hazards

	Entry	Exit
Constant	−2.452	−1.898
Category E		−0.636
Category D		−0.703
Category C		−0.724
Category B		−0.469
Category A		−0.102
$\log(1 + B_{it})$		−0.371
$\log(1 + P_{it})$		−0.760
Aggregate state, A	−0.064	−0.096
Aggregate state, B	−0.054	−0.019
Aggregate state, C	0.002	0.070
Aggregate state, D	0.376	0.073
Aggregate state, E	−0.225	−0.198
Firm-quarters	7,916	4,726
Event rate	0.081	0.100
Log loss	0.274	0.310
Brier score	0.073	0.087
Holdout firm-quarters	579	1,083
Holdout log loss	0.444	0.358
Holdout Brier score	0.130	0.104

Notes: Table reports logit coefficients for the entry and exit hazards as defined above. Category F is the omitted category in the exit model; the entry model conditions on the aggregate state alone, since the risk set is defined before a firm's first Works bid. Aggregate-state components are the standardized four-quarter qualified-bidder moments of (27). Values are measured in billions of Rwandan francs.

flow and proceed to the next state. The firm's final bid date is used to to measure the firm's exit. The sample for structural estimation ends in 2023Q3, but the full sample ends in 2025Q3, so I use the quarters after the structural estimation to determine whether the firm bids again. A firm who bid twice by 2023Q3 and at least once between 2023Q3 and 2025Q3 is not coded as having exited.

C.4.3 The aggregate state and its law of motion

The aggregate state summarizes the distribution of potential competitors by a small number of moments that firms use to forecast future auction competition. I use the average number of qualified bidders in each tender category, for categories E–A (those that firms can upgrade to).

Let $\mathcal{A}_\tau(k)$ be the set of Works lots in category $k \in \{A, B, C, D, E\}$ auctioned in quarter τ, and let

$$N_\ell^q = \#\{j : \ j \text{ submits a bid on } \ell \text{ and is not disqualified}\}.$$

At the beginning of quarter t, define

$$\bar{n}_t^k = \frac{\displaystyle\sum_{\tau=t-4}^{t-1} \sum_{\ell \in \mathcal{A}_\tau(k)} N_\ell^q}{\displaystyle\sum_{\tau=t-4}^{t-1} |\mathcal{A}_\tau(k)|}, \qquad k \in \{A, B, C, D, E\},$$

and collect these five numbers in $\bar{s}_t = (\bar{n}_t^A, \ldots, \bar{n}_t^E)$. The four-quarter window ends before t, so neither a firm's current participation choice nor the current auction's realized bidder count enters its own state. A window with no category-k lot is assigned the mean over all earlier category-k lots, and zero only before the first category-k lot appears in the data.

The participation, bid, entry, and exit specifications use standardized $\log(1 + \bar{n}_t^k)$. For the conditional-choice-probability step, each log component is discretized at into terciles into low, medium, and high values (across the time series). I use only the 22 distinct combinations that ever occur during the 24 sample quarters, instead of the possible $3^5 = 243$ possible combinations. I fit a multinomial logit for the probability that each category bin moves to low, medium, or high next quarter, as a function of all 5 current bins. The joint transition is the (normalized) product of the five component probabilities. Figure 15 shows how the five components evolve over the estimation window. The series are persistent within category but not across categories, illustrating varying degrees of competition across categories.

C.5 Type posterior and bid policy

C.5.1 Observed-data likelihood

For firm i, let $\mathcal{D}_i^\xi$, $\mathcal{D}_i^a$, and $\mathcal{D}_i^b$ denote its execution, participation, and bid histories. Let $\mathcal{C}_i$ collect the initial measured state and the opportunity and tender characteristics treated as conditioning

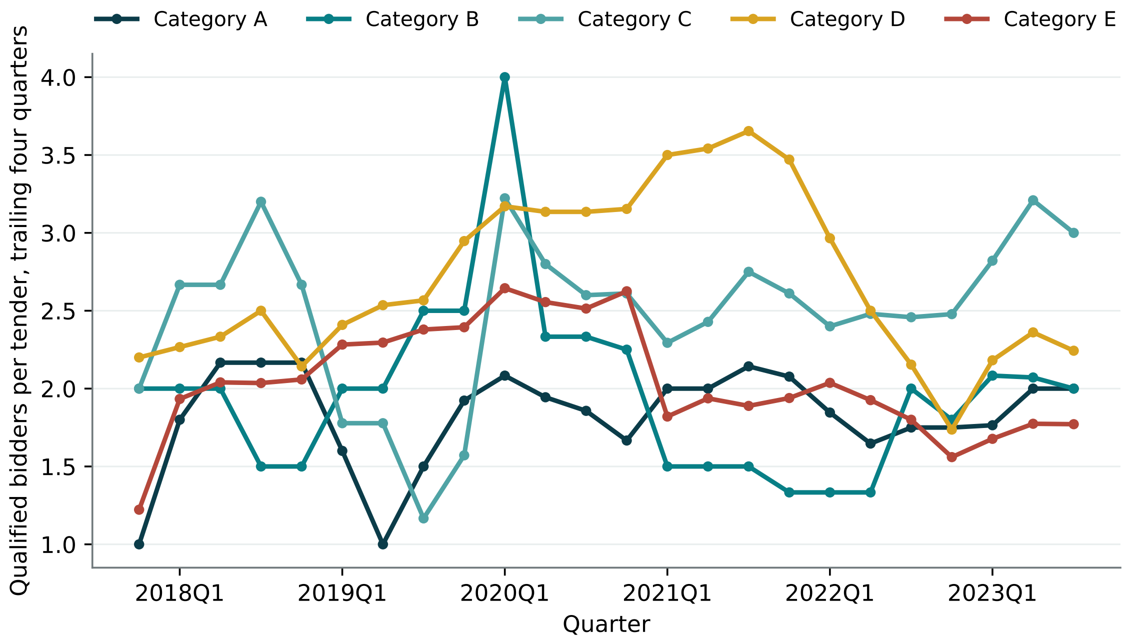

Figure 15: Evolution of the aggregate state

Notes: Each series is a component of $\bar{s}_t$: the mean number of qualified bidders per tender in that category over quarters $t-4,\ldots,t-1$. Category B is the category with least number of auctions, averaging 3.9 tender-lots per four-quarter window compared with 39.5 for category E, so its component is the most volatile.

variables. The state s_{it} appearing in each contribution is predetermined relative to that outcome: it is constructed recursively from the initial state and outcomes observed before the current draw.

Conditional on permanent type, the chain rule gives

$$
\begin{aligned}
&p\Big(\mathcal{D}_i^{\xi}, \mathcal{D}_i^{a}, \mathcal{D}_i^{b} \mid \theta_i = \theta, \mathcal{C}_i\Big) \\
&\quad = p\Big(\mathcal{D}_i^{\xi} \mid \theta_i = \theta, \mathcal{C}_i\Big)\, p\Big(\mathcal{D}_i^{a} \mid \mathcal{D}_i^{\xi}, \theta_i = \theta, \mathcal{C}_i\Big) \\
&\qquad \times p\Big(\mathcal{D}_i^{b} \mid \mathcal{D}_i^{a}, \mathcal{D}_i^{\xi}, \theta_i = \theta, \mathcal{C}_i\Big). \qquad (50)
\end{aligned}
$$

The maintained conditional-independence restriction states that, after conditioning on the current state, tender characteristics, and permanent type, a transitory execution draw contains no additional information about the participation or bid draw. Transitory draws are also independent across quarters and auctions. A bid is observed only when the firm submits a qualified bid, i.e.,

$a_{i\ell t} = 1$. Disqualified bids are excluded. It follows that

$$
\begin{aligned}
& p\Big(\mathcal{D}_i^{\xi}, \mathcal{D}_i^{a}, \mathcal{D}_i^{b} \mid \theta_i = \theta, \mathcal{C}_i\Big) \\
&\quad = \prod_{t \in \mathcal{D}_i^{\xi}} f_\theta^{\xi}(\xi_{it} \mid X_{it}) \\
&\qquad \times \prod_{(\ell,t)\in\mathcal{D}_i^{a}} \left[p^a_{\theta,i\ell t}\right]^{a_{i\ell t}} \left[1 - p^a_{\theta,i\ell t}\right]^{1-a_{i\ell t}} \\
&\qquad \times \prod_{(\ell,t)\in\mathcal{D}_i^{b}} f_\theta^{b}\Big(\widetilde{b}_{i\ell t} \mid a_{i\ell t} = 1, \chi_\ell, s_{it}, \bar{s}_t\Big) \\
&\quad = \mathcal{L}_i^{\xi}(\theta)\mathcal{L}_i^{a}(\theta)\mathcal{L}_i^{b}(\theta),
\end{aligned} \tag{51}
$$

where

$$p^a_{\theta,i\ell t} = a_\theta(\chi_\ell, s_{it}, \bar{s}_t).$$

Equation (50) follows from the chain rule. The product representation in (51) follows from the maintained Markov and conditional-independence restrictions.

The baseline specification uses a constant mixing share,

$$\mathbb{P}(\theta_i = \theta \mid \mathcal{C}_i) = \omega_\theta.$$

Summing over permanent types therefore gives

$$\mathcal{L}_i = \sum_{\theta \in \{L,H\}} \omega_\theta \mathcal{L}_i^{\xi}(\theta)\mathcal{L}_i^{a}(\theta)\mathcal{L}_i^{b}(\theta). \tag{52}$$

Backlog evolves mechanically from the initial state, new awards, and observed clearance. The exogenous policies – category promotion, entry, exit, tender-arrival, and aggregate-state – are estimated separately.

C.5.2 Econometrician's type posterior

The participation outcome is = 1 only when the firm's submitted bid is not disqualified. The econometrician's posterior is

$$w_i = \mathbb{P}(\theta_i = H \mid \mathcal{D}_i^{\xi}, \mathcal{D}_i^{b}, \mathcal{D}_i^{a}). \tag{53}$$

D Auction Inversion and Value Computation

D.1 Category distribution beliefs and the bid hazard

I parametrize normalized bids to follow a type-specific conditional Gamma function whose conditional mean depends on normalized backlog, $\log(1 + \tilde{P})$, the aggregate state, project scale, bidder count, sector, quarter, and firm-category–tender-category cells. Firm i knows its own bid function (because it knows θ_i), but not rivals' bid functions. Let $f_{\theta,j}(b \mid \chi_\ell, s_{jt}, \bar{s}_t)$ denote the type-specific bid density and let $\bar{F}_{\theta,j}(b \mid \chi_\ell, s_{jt}, \bar{s}_t) = 1 - F_{\theta,j}(b \mid \chi_\ell, s_{jt}, \bar{s}_t)$ denote its survival function.

For lot ℓ, set $m = m_\ell$, so rival j's eligibility category is $\kappa_{jm,t}$. Let $\omega_{H,\kappa_{jm,t}}$ be the estimated share of type-H firms in category κ. These six category shares are common prior beliefs and are constructed from the estimated posterior masses in the relevant eligibility categories. Once rival j is observed in the entrant set, the prior is reweighted by the type-specific participation policy as in (16). Write the resulting probability as $\tilde{\omega}_{j\ell t}$. For rival j, define the perceived density and survival function

$$
\begin{aligned}
f_j^{\text{mix}}(b \mid \chi_\ell, s_{jt}, \bar{s}_t) &= \tilde{\omega}_{j\ell t} f_{H,j}(b \mid \chi_\ell, s_{jt}, \bar{s}_t) \\
&\quad + [1 - \tilde{\omega}_{j\ell t}] f_{L,j}(b \mid \chi_\ell, s_{jt}, \bar{s}_t), \\
\bar{F}_j^{\text{mix}}(b \mid \chi_\ell, s_{jt}, \bar{s}_t) &= \tilde{\omega}_{j\ell t} \bar{F}_{H,j}(b \mid \chi_\ell, s_{jt}, \bar{s}_t) \\
&\quad + [1 - \tilde{\omega}_{j\ell t}] \bar{F}_{L,j}(b \mid \chi_\ell, s_{jt}, \bar{s}_t).
\end{aligned} \tag{54}
$$

The perceived contribution of rival j to the procurement hazard is the hazard of this mixture. Thus, bidder i's hazard is

$$
H_i(b \mid \mathcal{J}_{\ell t}) = \sum_{j \in \mathcal{J}_{\ell t} \setminus \{i\}} \frac{f_j^{\text{mix}}(b \mid \chi_\ell, s_{jt}, \bar{s}_t)}{\bar{F}_j^{\text{mix}}(b \mid \chi_\ell, s_{jt}, \bar{s}_t)}. \tag{55}
$$

Writing $h_{\theta,j} = f_{\theta,j}/\bar{F}_{\theta,j}$ for the type-specific hazards,

$$
\frac{f_j^{\text{mix}}(b)}{\bar{F}_j^{\text{mix}}(b)} = \tilde{\omega}_{j\ell t} \frac{\bar{F}_{H,j}(b)}{\bar{F}_j^{\text{mix}}(b)} h_{H,j}(b) + \big(1 - \tilde{\omega}_{j\ell t}\big) \frac{\bar{F}_{L,j}(b)}{\bar{F}_j^{\text{mix}}(b)} h_{L,j}(b), \tag{56}
$$

a survival-function-weighted average of the type hazards whose weights depend on the bid.

D.2 Derivation of the bid first-order condition

From the value function to the bid objective. Fix a project ℓ in which firm i participates. The bid is chosen after the firm observes its cost, the entrant set, and its rivals' public states. Hold the firm's actions and outcomes in its other quarter-t tenders fixed, and split the expectation in (21) by the outcome of the focal auction: with probability $G_i(b \mid \mathcal{J}_{\ell t})$ the firm wins, collects $b - c_{i\ell t}$, and continues from $s_{i,t+1}^{\text{win}}$; otherwise it continues from $s_{i,t+1}^{\text{lose}}$. The terms involving b are

$$
\begin{aligned}
G_i(b \mid \mathcal{J}_{\ell t}) \Big[b - c_{i\ell t} + \beta\, \mathbb{E}\, V_{\theta_i}\big(s_{i,t+1}^{\text{win}}, \bar{s}_{t+1}\big) \Big] &+ \big[1 - G_i(b \mid \mathcal{J}_{\ell t})\big]\, \beta\, \mathbb{E}\, V_{\theta_i}\big(s_{i,t+1}^{\text{lose}}, \bar{s}_{t+1}\big) \\
= G_i(b \mid \mathcal{J}_{\ell t}) \big[b - c_{i\ell t} + D_{i\ell t} \big] &+ \beta\, \mathbb{E}\, V_{\theta_i}\big(s_{i,t+1}^{\text{lose}}, \bar{s}_{t+1}\big),
\end{aligned}
$$

where the equality adds and subtracts the losing continuation value and applies the definition of $D_{i\ell t}$ in (22). The losing continuation does not depend on b, flow profits from the firm's other quarter-t tenders enter additively, and the survival probability multiplies the entire bracket in (21); none of these affect the optimal bid. This yields the objective (23).

First- and second-order conditions. Suppress indices and write the expected payoff from bid b as

$$u(b) = G(b)(b - c + D),$$

where $G'(b) < 0$ in a low-price auction. The engineering cost estimate increment and clearance transition process are both fixed before bidding, so the continuation gain satisfies $\partial D/\partial b = 0$. The first-order condition is

$$G(b) + G'(b)(b - c + D) = 0.$$

Therefore,

$$b - c + D = -\frac{G(b)}{G'(b)} = \frac{1}{-G'(b)/G(b)} = \frac{1}{H(b)},$$

which yields (24). The second-order condition is

$$2G'(b) + G''(b)(b - c + D) < 0.$$

The empirical support restricts attention to bids for which the first-order inversion is well defined and the estimated rival hazard is positive.

D.3 Value computation

The value computation follows the procedure of (De Silva et al., 2023), adapted for discrete time (their model is in continuous time). Let M^0_θ denote the "no-award" transition matrix that simulates that firms win no tenders in the future for type θ under the estimated policies. This transition matrix integrates over clearance, payments, paid experience, category promotion, exit, and the aggregate state law of motion, and its row sums incorporate the estimated probability of survival, but it sets new prospective awards to zero. Let $\mathbf{\Pi}^A_\theta$ denote the vector, over the empirical state grid, of the survival-weighted auction surplus $\Pi^A_\theta(s) = E[\, G_\theta(B \mid \mathcal{J})/H_\theta(B \mid \mathcal{J}) \mid s, \theta\,]$, where the expectation is over participation, project arrivals and characteristics, the firm's own bid, the realized rival set under the estimated policies.

This expectation is estimated from the observed bids. The per-bid ratio, $\widehat{G}/\widehat{H}$, is regressed on project characteristics, the firm's state, the aggregate state, and type, and the fitted conditional mean, multiplied by the participation probability and the arrival intensities, is evaluated on the grid. By the first-order condition above, $G(b)/H(b) = G(b)(b - c + D)$ bid by bid, so Π^A_θ equals the expected current profit inclusive of the continuation value of winning. This step does not use the construction cost parameters. Let $\mathbf{V}_\theta$ stack the state-specific value function $V_\theta(s, \bar{s})$ over the same grid. The stacked value function is

$$\mathbf{V}_\theta = (I - \beta M^0_\theta)^{-1} \mathbf{\Pi}^A_\theta. \tag{57}$$

The no-award matrix is required because $\mathbf{\Pi}^A_\theta$ already carries the value of prospective awards. I solve the linear system once per type on the empirical state support, then calculate the continuation gain of each observed award from the resulting $\mathbf{V}_\theta$ by applying the discount factor β. I then recover the construction cost parameters from the bid first-order condition.

E Additional Estimates and Model Fit

This appendix reports additional first-stage and cost-equation estimates of the estimated specification to the ones reported in the main text. Appendix coefficients point estimates, with 95% percentile intervals from 300 firm-bootstrap draws. Table 2 in the main text reports the cost equation's type, state, scale, category and sector coefficients. Table 14 below reports the project controls and the variance slope.

Table 14: Execution: ZOIB clearance process

Parameter	Estimate	95% CI
No-clearance intercept, α_0^ξ	0.077	[−0.156, 0.290]
No-clearance standardized log backlog, β_0^ξ	−0.001	[−0.188, 0.173]
Full-clearance intercept, α_1^ξ	−1.851	[−2.114, −1.610]
Full-clearance standardized log backlog, β_1^ξ	−1.331	[−1.619, −1.029]
Partial-mean intercept, α_μ^ξ	−0.480	[−0.639, −0.327]
Partial-mean standardized log backlog, β_μ^ξ	−0.806	[−0.942, −0.663]
Log type loading, $\log \lambda^\xi$	−0.392	[−1.252, −0.027]
Log Beta precision, $\log \phi$	0.738	[0.625, 0.891]
Logit H share, α^ω	−1.396	[−1.865, −0.971]

Notes: The nine-parameter finite-mixture clearance model of Appendix C.2. The single type loading $\lambda^\xi = \exp(\log \lambda^\xi)$ shifts all three components toward faster clearance for the High type. The backlog variable is standardized as in (33). At the median post-award backlog the implied expected clearance per quarter is 0.224 for the Low type and 0.423 for the High type. Confidence intervals are 95% percentile intervals from the 300 firm bootstrap draws.

Table 15: Participation logit

Parameter	Estimate	95% CI
Constant, δ_A	−3.450	[−4.083, −2.624]
Current backlog, δ_B	−0.242	[−0.339, −0.163]
Paid experience, δ_P	−0.488	[−0.650, −0.337]
Same category, δ_M	−0.728	[−1.030, −0.361]
Category gap, δ_R	−1.204	[−1.431, −0.947]
Mean backlog, $\delta_{\bar{B}}$	0.454	[0.285, 0.691]
Mean paid experience, $\delta_{\bar{P}}$	0.617	[0.349, 0.911]
Mean same category, $\delta_{\bar{M}}$	0.086	[−0.753, 0.729]
Mean gap, $\delta_{\bar{R}}$	0.870	[0.373, 1.301]
High-type intercept, δ_H	1.331	[1.054, 1.593]

Notes: A positive outcome is a submitted bid that is not disqualified before evaluation. Aggregate-state controls are estimated but omitted here. Confidence intervals are 95% percentile intervals from the 300 firm bootstrap draws.

Table 16: Additional second-stage coefficients: project costs

Parameter	Estimate	95% CI
Project controls		
BOQ log item count (standardized)	0.109	[0.005, 0.196]
BOQ item count missing	−0.163	[−0.824, 0.406]
Log advertised duration (standardized)	−0.017	[−0.093, 0.068]
Advertised duration missing	0.034	[−0.967, 1.237]
Procuring entity: district	−0.128	[−0.457, 0.249]
Procuring entity: government agency	−0.068	[−0.409, 0.285]
Procuring entity: hospital	0.009	[−0.338, 0.417]
Procuring entity: ministry	−0.019	[−1.211, 0.311]
Lagged material-price shock (standardized)	0.018	[−0.033, 0.060]
Material-price shock missing	−0.182	[−0.611, 0.420]

Notes: Coefficients are in units of the engineering cost estimate. The type, state, scale, firm-category, tender-category and sector coefficients of the same equation, and the cost dispersion σ, are reported in Table 2. Confidence intervals are 95% percentile intervals from the 300 firm bootstrap draws.

E.1 Model fit

Table 17 compares selected statistics from the model's predictions compared with the data at observed states with qualified bidders. Overall, the model's predictions are close to the data. The model slightly over-predicts mean quarterly backlog clearance (0.280 against 0.260) and the median submitted bid. It slightly under-predicts participation and the median winning bid.

Table 17: Model fit statistics

Moment	Data	Model
Auction participation and bids		
Mean participating firms per tender-lot	2.13	1.89
Median bid / engineering cost estimate	0.981	1.023
Median winning bid / engineering cost estimate	0.923	0.800
Quarterly backlog clearance		
Mean share of backlog cleared	26.0%	28.0%

Notes: Simulated at observed states using the estimated participation and bid policies, on the 1,319 lots of the estimation sample; the winning bid is the lowest qualified bid, with the simulated rival standing in on one-qualified lots.

E.2 Counterfactual exhibits

Figure 16: Opening tenders to lower categories: outcomes by degree of loosening

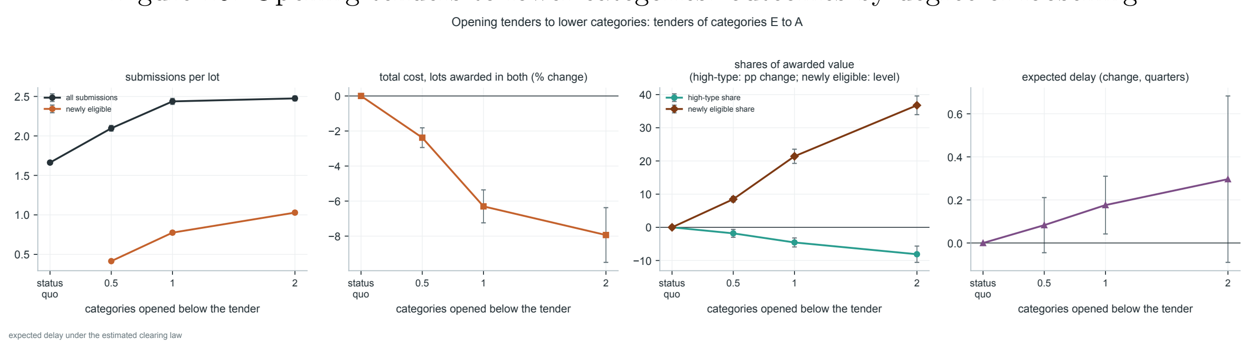

Notes: Tenders of categories E to A; status quo and 0.5, 1 and 2 categories opened below the tender (0.5 admits firms one category below to the smaller half of the category, at or below its median estimate). Submissions per lot, in total and by newly eligible firms; percentage change in the sum of winning bids on the lots awarded both in the counterfactual and in the status quo; change in the share of awarded value won by high-type firms (percentage points) and the level of the share won by newly eligible firms; change in expected completion time, weighting lots by the winning bid. Whiskers are two Monte Carlo standard errors across 20 replications.