MARKET ACCESS AND COMPETITIVE ENTRY: EVIDENCE FROM CHINA'S DRUG PROCUREMENT PROGRAM

SHENGMAO CAO, KARAM KANG, ANH NGUYEN, AND CHUAN YU

Preliminary and comments welcome. Please do not cite.
Click here for the most recent version.

Abstract. Entry may remain limited when firms must incur substantial fixed costs without assurance that buyers will recognize their quality or that they can reach sufficient demand. Can governments alleviate these barriers while preserving competition? We study China's centralized drug procurement program, which requires bioequivalence (BE) certification for auction participation and guarantees winners access to a large share of demand. Certification alone raises prices but generates little entry or quantity expansion; once procurement links certification to market access, roughly one-quarter of winners are new entrants and prices fall substantially. We develop and estimate a model of entry, certification, procurement auctions, and post-auction competition that separates the information, market-access, and fixed-cost channels. Comparing centralized procurement with a direct entry subsidy, we show that centralized procurement can promote entry while maintaining competitive discipline: it helps firms recover fixed costs, but allocates valuable market access through competitive bidding.

1. Introduction

Capable producers do not necessarily become effective competitors. Entering a market requires substantial upfront investments in product development, production capacity, and distribution. Recovering these costs is difficult when market access is uncertain: fragmented purchasing and established buyer relationships make demand hard to reach. Consumers may also be unable to verify the quality of unfamiliar products. These frictions can make entry unprofitable even for capable firms: high fixed costs are difficult to recover when access to demand is uncertain and buyers

Date: August 23, 2026.
Cao (email: shengmao.cao@kellogg.northwestern.edu): Northwestern University.
Kang (email: karam.kang@wisc.edu): University of Wisconsin-Madison and NBER.
Nguyen (email: anhnguyen@cmu.edu), Carnegie Mellon University.
Yu (email: chuanyu@unc.edu): University of North Carolina, Chapel Hill.

cannot readily assess product quality. Pharmaceutical markets provide a prominent example, as producers of therapeutically equivalent generics often struggle to displace established and branded firms. Similar barriers arise in food safety, professional services, education, and digital platforms.

This paper studies what government policies can overcome barriers to entry while preserving competition. Standard policy tools address the underlying frictions only partially. Entry subsidies lower upfront costs but do not ensure access to buyers or resolve quality uncertainty. Certification can establish credibility, but it imposes an additional cost and, by itself, provides no assurance of demand. Centralized procurement offers a distinct approach: by aggregating purchases across many buyers, it can provide suppliers with access to a large pool of demand, while an auction can discipline the prices at which that access is granted. When participation is conditioned on certification, centralized procurement can also link market access to verified quality. The overall effects are nevertheless ambiguous. Greater access to demand may encourage entry, but competitive bidding compresses margins and certification raises the cost of becoming eligible. Whether such a policy ultimately expands entry, improves the composition of suppliers, and lowers prices is therefore an empirical question.

We study this question in China's pharmaceutical market, a setting in which each of these frictions has historically been important. Before recent reforms, many off-patent drugs remained expensive despite the presence of domestic generic manufacturers. Consumers and health-care providers had difficulty assessing generic quality, while recurring scandals and weak regulatory oversight reinforced their skepticism, allowing established and branded products to retain a substantial advantage. At the same time, procurement was also fragmented across hospitals and regional purchasing alliances, so suppliers seeking to expand had to reach many buyers individually and incur substantial marketing and distribution costs. The combination of uncertain product credibility and fragmented access to demand made China a natural setting in which to study how certification and centralized procurement reshape firms' incentives to enter and compete.

The Chinese government responded with two closely related reforms. Beginning in 2016, generic manufacturers could obtain bioequivalence (BE) certification, providing an informative signal that their products were therapeutically equivalent to the reference branded drug. China then introduced centralized volume-based procurement (CVBP), first through a pilot in 2018 and nationally in 2019. Under CVBP,

certified generic firms compete in centralized auctions, and winners receive immediate access to a large, guaranteed volume of demand across designated provinces. Together, the reforms pair a quality signal that reduces informational frictions with a procurement mechanism that expands market access.

Our analysis combines several data sources covering 2012 through 2025. We assemble a panel of quarterly prices and quantities for several thousand off-patent drugs across 25 Chinese provinces, records of BE certification applications and outcomes, and procurement data on bids, winners, and province assignments. The sequential introduction of the reforms creates three distinct regimes: markets initially operate without certification or centralized procurement, then transition to certification alone, and, for procured drugs, subsequently combine certification with centralized procurement. Following firms across these regimes allow us to distinguish the informational role of certification from the market-access and cost effects of procurement.

We begin by documenting several empirical facts that help separate the roles of information, market access, and cost reductions. First, certification has market value but only a limited standalone effect on expansion. Obtaining BE approval allows a generic product to command a higher price than uncertified products, indicating that certification conveys meaningful information about quality. Yet certification leads to little expansion in the quantities sold by newly certified firms. This contrast suggests that certification can increase perceived quality and willingness to pay but, on it own, does not overcome barriers to market access.

Second, centralized procurement is associated with a sharp increase in entry and a broad reallocation of demand. Roughly one-quarter of auction winners are new entrants that had not previously sold the drug. Following procurement, certified generics account for more than 80% of quantities, while the branded share falls from about 40% during the certification-only period to about 10%. Because auction participation requires BE approval, these patterns are consistent with guaranteed market access increasing the returns to certification and entry, enabling certified generics, including new entrants, to challenge established brands.

Third, generic firms bid far below their previous prices: the median bid is just 31% of the price by the same firm before procurement and is well below the pre-procurement marginal cost implied by conventional markup estimates. The coexistence of low bids and substantial entry suggests that procurement does not promote entry by shielding firms from competition. Instead, the evidence is consistent with

procurement reducing the costs of reaching and serving buyers through lower marketing expenses and economies of scale, while the auction induces firms to pass these savings through in their bids. Procurement therefore makes entry viable by expanding market access and lowering costs while preserving competitive discipline.

To quantify these mechanisms and evaluate alternative policy designs, we develop and estimate a model of firm entry, quality certification, procurement auctions, and post-auction price competition. Potential entrants differ in privately observed quality and face heterogeneous fixed costs of entry and certification. Firms choose whether to enter and seek certification, while consumers do not observe underlying quality and infer it from certification status. Certification is an informative but imperfect signal of quality. Certified firms then compete for guaranteed market access in split-award auctions, and procurement may further reduce costs through economies of scale. A key equilibrium feedback is that procurement changes the returns to entry and certification, while the resulting entry changes consumer beliefs and the composition of the bidder pool. The model allows us to separate the information, market-access, and scale channels through which procurement affects entry, auction competition, and downstream outcomes.

A key challenge for identification is that product quality is unobserved. We address this challenge by combining information from firms' certification applications and approval outcomes with market data from the periods before certification was introduced and after procurement was implemented. We do not use the intermediate certification-only period for identification because firms may have anticipated future procurement opportunities, and because entry, testing, and certification decisions take time to adjust. Instead, we use this regime later as a counterfactual benchmark.

Identification proceeds in several steps. Pre-certification demand identifies price sensitivity and the distribution of perceived product quality before certification conveys information. Post-procurement market shares identify perceived quality conditional on certification status and the strength of the demand guarantee. Non-winners' downstream pricing decisions identify marginal costs, while procurement bids and province assignments discipline the economies of scale generated by winning additional markets. Together, these steps recover the key demand and cost objects governing downstream competition.

The final step exploits firms' entry and testing decisions as well as certification outcomes to recover the latent composition of firms, the certification technology, and the distributions of entry and certification costs and to separate consumers' valuation

of quality from equilibrium beliefs. Certification applications and approval outcomes provide direct information on firms' incentives to seek certification and on the selectivity of the certification process. We combine these moments with the demand-side perceived-quality measures recovered in the previous steps and with differences in the cost distributions of certified and untested products. The identification strategy closely mirrors the structure of the model, linking observed entry, testing, certification, pricing, and bidding behavior to the underlying primitives governing firm behavior.

We use the estimated model to evaluate the effects of certification and centralized procurement. Certification alone reduces informational frictions but has limited ability to induce entry when firms still face substantial barriers to reaching buyers. Linking certification to procurement access raises the return to entry and attracts firms that had not previously participated in the market. At the same time, allocating access through an auction intensifies competition among both entrants and incumbent producers. Procurement therefore does more than reallocate demand toward a selected set of firms: it changes which firms enter, which firms invest in certification, and which firms compete for the market.

We then compare centralized procurement with a direct entry subsidy calibrated to induce the same amount of entry as procurement. Both policies can increase the return to entering, but they operate through different mechanisms. An entry subsidy lowers firms' fixed costs without directly addressing informational frictions or guaranteeing access to buyers. Procurement instead links verified quality to immediate market access and allocates that access through competition. The comparison therefore allows us to quantify whether government demand can induce entry more effectively than a subsidy, while also generating lower prices by expanding and changing the pool of competing firms. More broadly, the analysis connects procurement to a pro-competitive form of industrial policy. Rather than selecting firms for protection or subsidizing entry directly, the government can reduce barriers to participation and allocate access to demand through competition.

1.1. **Literature Review.** This paper contributes to three strands of literature: certification and quality regulation, procurement and entry, and auction design.

First, we contribute to work on certification and quality regulation in markets with imperfect information. Theoretical work shows that disclosure and certification

can redirect demand toward higher-quality sellers and strengthen incentives to invest in quality, but their effects depend on signal credibility, how consumers interpret certification, and certifier incentives (Dranove and Jin, 2010). Minimum-quality standards and licensing can reduce adverse selection, but they can also raise entry costs or induce excessive investment in credentials (Leland, 1979; Shapiro, 1986).

Empirical work examines how quality signals affect market outcomes in practice. In pharmaceutical markets, Atal et al. (2022) show that bioequivalence regulation in Chile increased demand for certified generics and shifted entry toward higher-quality products, while Ishitani (2025) studies how intermediaries use quality disclosures to discipline product quality in the generic pharmaceutical market. Evidence from other markets similarly shows that public and platform certification can change demand: restaurant hygiene grades shift consumer behavior and improve quality, and eBay certification is more valuable when sellers lack established reputations (Jin and Leslie, 2003; Elfenbein et al., 2015).

In our setting, BE approval was a relatively new and imperfect signal of quality, and consumers may have continued to view certified generics as inferior substitutes for incumbent products. We show that certification has market value, but an imperfect signal need not automatically expand the market or fully resolve information asymmetry. Certified firms may still face limited demand without complementary market-access mechanisms.

Second, we contribute to work on procurement and entry. Evidence from procurement across various industries shows that bidder participation is endogenous and responds to entry costs, bidder preferences, set-asides, publicity rules, buyer discretion, and contract complexity (Li and Zheng, 2009; Krasnokutskaya and Seim, 2011; Athey et al., 2013; Coviello and Mariniello, 2014; Coviello et al., 2018; Bajari et al., 2008; Kang and Miller, 2022). This literature motivates our treatment of procurement design as a policy that affects both firm participation and prices conditional on participation.

The pharmaceutical procurement literature shows that centralized purchasing and competitive bidding can lower drug prices while also affecting supply, selection, and market structure. Studies of procurement systems that use competitive bidding or aggregate demand across buyers show that these effects depend on competition, supplier concentration, buyer heterogeneity, and supply responses (Danzon et al., 2015; Dubois et al., 2021; Brugués, 2020; Allende et al., 2024). Related work on China's centralized drug procurement program documents large price reductions and

broader market responses, including changes in firms' sales and marketing incentives (Cao et al., 2024; Zhang et al., 2025; Liu et al., 2025). We contribute to this literature by showing that procurement can operate as a market-access policy, not only as a purchasing mechanism. Because auction eligibility requires BE certification, procurement changes firms' incentives to become certified as well as their downstream bidding and post-auction competition.

Importantly, we study the complementarity of quality certification and market access through procurement in addressing information asymmetry and market power. Work on quality thresholds and qualification rules shows that screening affects bidder participation and procurement costs (Cripps and Ireland, 1994; Estache and Iimi, 2012). In our setting, BE approval creates an informative but incomplete quality signal. Procurement makes this signal payoff-relevant by using certification as an eligibility requirement for guaranteed demand, while certification determines which firms can compete for the procurement prize. This interaction is central to our analysis.

Third, we connect to a large body of literature on auction design. Classic mechanism-design work studies procurement with private costs and non-price attributes (Che, 1993; Manelli and Vincent, 1995). Scoring and multidimensional auctions can improve outcomes when price and quality both matter (Asker and Cantillon, 2010). China's design instead screens quality through BE certification before running a price-focused auction. This separation can be attractive when certification is informative, but it also shifts participation toward firms with stronger incentives to become certified.

An important feature of China's CVBP program is its use of split-award auctions. Split awards can preserve multiple suppliers, but their desirability depends on coordination, investment incentives, and the efficient number of firms (Anton and Yao, 1992; Chaturvedi et al., 2014; Gong et al., 2012). Closest to our paper, Zhang (2023) studies the consequences of split-award procurement in China's pharmaceutical market. We extend this literature by linking procurement design to firms' incentives to invest in quality certification and enter the market, and by analyzing how procurement-induced market access affects subsequent competition and prices.

2. Institution and Data

This section provides the institutional and empirical background for our analysis of certification, procurement, and competition in China's pharmaceutical market.

The setting is particularly useful because a firm's access to centralized procurement depends on obtaining bioequivalence (BE) certification, creating a direct link between quality regulation and market access. After describing the industry, policy environment, and data, we document several patterns that guide the construction of our model, including firms' certification decisions, participation in procurement auctions, and the resulting effects on entry and competition.

2.1. **Institutional Background.** China's pharmaceutical market is the largest in the world by consumption volume and the second largest by revenue. A defining institutional feature is the central role of hospitals in drug distribution: approximately 85% of prescription drugs are dispensed through hospital pharmacies rather than retail outlets, reflecting a system in which prescribing and dispensing are integrated within the same institution. Patients receive drugs following physician consultations, and their choices are largely determined by the products stocked in the hospital pharmacy. Public health insurance covers drugs listed in national and provincial formularies, with patients' out-of-pocket share ranging from approximately 20% to 60% depending on the region, insurance plan, and formulary tier.[1]

Drug prices in China have been consistently high relative to international benchmarks. In 2018, pharmaceutical expenditure accounted for approximately 40% of total healthcare spending, more than double the OECD average of 17%. A primary contributor is spending on off-patent drugs: despite a large industry with thousands of domestic generic producers, prices for many off-patent drugs remained elevated — in some cases far exceeding those in the United States after adjusting for purchasing power parity (Zhang et al., 2025).

Two structural features of the industry have historically sustained high drug prices. First, widespread public skepticism about the safety and efficacy of domestic generics — fueled by lax regulatory oversight and recurring quality scandals — gave branded originator products substantial pricing power long after patent expiry. Second, drug procurement was historically conducted by individual hospitals or regional alliances. Limited bargaining leverage under a fragmented procurement system, coupled with strong brand preferences and widespread detailing efforts, allowed branded

[1]The catalog distinguishes two coverage tiers. For tier A drugs, the full cost is eligible for reimbursement at plan-specific rates. For tier B drugs, patients first pay 10% to 20% out of pocket before coverage applies to the remainder.

drugs to command large market shares in many therapeutic markets. Generic aversion and incumbents' institutional ties also deterred entry of high-quality generics, entrenching market power among established firms.

In response to these challenges, the Chinese government implemented two landmark reforms over the past decade. In 2016, it introduced bioequivalence (BE) testing for quality certification of generic products. In late 2018, it launched the centralized volume-based procurement (CVBP) program. These two reforms are closely linked by design: BE certification is a prerequisite for participation in CVBP auctions.

To obtain BE certification, a generic manufacturer conducts a clinical pharmacokinetic study — typically involving around 50 healthy human subjects — comparing the absorption profile of its product against the branded originator. A product passes if key pharmacokinetic parameters, primarily the area under the concentration-time curve and peak plasma concentration, fall within an 80–125% equivalence window relative to the reference product. The test typically costs between 3 and 10 million CNY and takes one to two years from study initiation to approval. Approved products carry a BE certification mark on their packaging.

The government launched CVBP in late 2018, beginning with a pilot in eleven major cities before expanding to a nationwide rollout across all provinces the following year. There have been eleven national rounds between 2019 and 2026. At the start of each round, the government announces a list of drugs to be procured — defined at the molecule-dosage-form level — along with quantity guarantees by province and a reserve price set at approximately the prior-year average market price. The guaranteed quantity is a fixed fraction of total prior-year procurement, ranging from 50% to 80% depending on the number of bidders. Manufacturers of branded and bioequivalent generic products are invited to submit bids. All national rounds follow a split-award format: the lowest bidder selects a province first, and all bidders then take turns selecting one province each in ascending bid order until all provinces are allocated. Contract duration ranges from one to three years, depending on the number of bidders. The eleven rounds cover a total of 487 molecule-dosage-forms (Appendix Table A1).

Once a contract is in place, winning firms are required to supply the contracted quantity; failure to do so bars them from participating in future procurement rounds. Individual hospitals are allocated purchase quotas from the winning firms; meeting these targets is a central criterion in hospital management performance reviews and in determining future hospital budgets. This mandatory compliance structure

is consequential for firms' entry decisions: winning procurement provides immediate, contractually secured access to a substantial share of the market in assigned provinces, without the need to establish a brand reputation or hospital relationships that incumbent firms have accumulated over years.

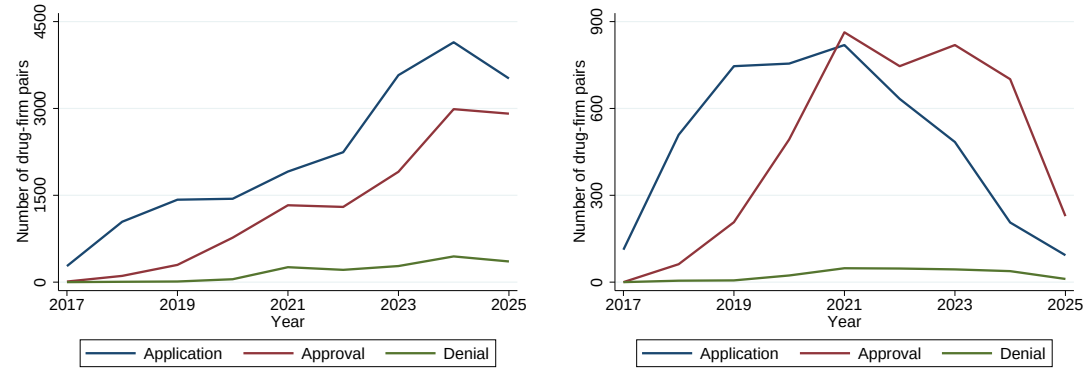
2.2. **Source of Data.** We obtain quarterly drug price and sales data at the product-by-province level from MENET, a leading pharmaceutical market research company in China. The data covers 25 provinces from the first quarter of 2012 to the second quarter of 2026. Each product is a stock-keeping unit (SKU) defined at the firm-molecule-form-strength level. We aggregate the sample to the firm-molecule-form-province-quarter level, converting quantities into standardized units using the same conversion factors used by the Chinese government to determine total procurement quantities. The data are drawn from a panel of over 900 hospitals in partnership with MENET, which over-represents tertiary hospitals in major cities; our sample therefore chiefly reflects drugs dispensed at these facilities.

MENET also provides the BE status and evaluation timeline for each generic product, compiled from public records of the National Medical Product Administration (NMPA) and the Center for Drug Evaluation (CDE). For each application, we observe the submission date, the outcome, and the decision date; for a small subset of products that underwent a revision process, we observe additional intermediate review milestones. For approved products, we additionally obtain published trial reports, including the number of subjects enrolled and key pharmacokinetic parameters for both the test and reference products.

We collect detailed information on each auction round from the official CVBP model, including the procured molecules, quantity guarantees, auction rules, winner identities, and the provincial assignments of each winner. For the first nine rounds, we additionally observe all submitted bids, including those of non-winning participants. For a subset of cases involving outsourced productions, we observe both the bidding firm and the contracted manufacturing partner.

Finally, we compile several auxiliary datasets. We collect all versions of China's Basic Insurance Drug Catalog from 2005 to 2026, recording the national and provincial formulary status of each drug over time. Formulary status is set at the molecule-dosage-form level and applies uniformly to all products in that category. Prior to 2017, individual provinces administered their own supplementary tier B catalogs;

FIGURE 1. BE Applications and Review Outcomes



(A) Full Sample (B) Procurement Sample

Notes: Panel A uses the full BE dataset. Panel B restricts to the procurement sample. Both panels plot annual counts of applications, approvals, and denials.

from 2019 onwards, the formulary was unified at the national level. We additionally obtain province-level average copay rates from Barwick et al. (2025).

2.3. **Certification Participation.** Figure 1 shows that BE applications rise gradually over time in the full set of drugs after the government launched the BE certification program in 2016. The increase is more pronounced among procurement drugs, for which applications and approvals surge around the procurement schedule. As BE certification is required for procurement eligibility, this surge suggests that firms respond to procurement incentives when deciding whether and when to apply.

Not all firms apply for BE certification, suggesting that certification is costly and that some firms self-select out of the process, especially when they know they do not meet the standard. This interpretation is consistent with anecdotal evidence that certification requires substantial investment and imposes a real hurdle on firms. At the same time, denials are not uncommon: in the procurement sample, denied applications account for 5.1% of resolved applications over 2018–2025.

This pattern suggests that certification is selective but imperfect. If certification perfectly distinguished between firms that meet and do not meet the standard, firms with little chance of approval would rarely apply, and the approval rate among applicants would be close to 100%. Instead, the observed approval rate of roughly 95% implies that some firms with a nontrivial risk of failure still choose to apply. We therefore model certification as an informative but imperfect screening mechanism:

firms possess private information about their likelihood of meeting the standard, yet certification outcomes remain uncertain.

2.4. **Market Structure and Procurement Outcomes.** We conclude the institutional discussion with data patterns that provide context for the economic mechanisms studied in the paper. Specifically, we first summarize how market structure evolves across the three institutional regimes observed in our data, and characterize participation, bidding, and allocations in the procurement auctions.

2.4.1. *Market Outcomes Across Different Regimes.* Panel A of Table 1 provides a descriptive overview of the market outcomes for drugs included in centralized procurement. We organize each drug's observed history into three periods. The pre-certification period extends from the beginning of the sample through the quarter before the first generic firm obtains BE certification. The post-certification, pre-auction period begins with the first certification and ends in the quarter before procurement implementation. The post-auction period begins when procurement is implemented. A generic firm is classified as certified in a given period if it has obtained certification by the end of that period.

Panel A reveals a marked change in market composition across the three periods. Before certification, branded products account for 45.7% of quantities. By the post-certification, pre-auction period, the market contains an average of 4.4 certified generic firms, which together account for 40.4% of quantities. Nevertheless, branded products retain nearly 40% of the market, even though their average price is about 80% higher than that of certified generics. Market composition changes much more sharply after procurement: certified generics account for 86.2% of quantities, and their average number rises to 9.2 firms. There price also fall sharply. Meanwhile, the branded share falls to 11.0%, and uncertified generics account for only 2.8% of quantities.

The best- and worst-ranked winner rows follow the firms that eventually submit the lowest and highest successful bids, respectively. These firms have relatively small market shares before procurement. After procurement, both gain substantial market share: the best-ranked winner accounts for 31.9% of drug quantity on average, while even the worst-ranked winner accounts for 16.7%.

These descriptive patterns provide an overview of the margins studied in the remainder of the paper. Certification is followed by a reallocation toward certified generics before procurement, while the largest changes in market shares and prices occur after procurement links certification to auction-based market access.

TABLE 1. Market and Auction Outcomes

Panel A: Market Outcomes

	Pre-Certification			Post-Certification, Pre-Auction			Post-Auction		
	Price / Reserve	Market Share (%)	# Firms	Price / Reserve	Market Share (%)	# Firms	Price / Reserve	Market Share (%)	# Firms
Branded	3.35	45.7	0.8	3.04	39.8	0.7	2.46	11.0	0.7
	(6.44)	(40.8)	(0.5)	(6.00)	(37.6)	(0.5)	(5.29)	(13.3)	(0.5)
Generic, certified	–	0.0	0.0	1.66	40.4	4.4	0.45	86.2	9.2
	–	(0.0)	(0.0)	(2.50)	(33.5)	(3.2)	(0.63)	(15.6)	(5.8)
Generic, uncertified	1.76	54.3	14.6	1.77	19.7	7.0	1.61	2.8	3.7
	(2.23)	(40.8)	(18.9)	(2.50)	(26.0)	(10.0)	(2.00)	(8.9)	(6.3)
Best-ranked winner	1.40	9.4	0.6	1.46	10.6	0.8	0.33	31.9	1.0
	(1.29)	(19.1)	(0.5)	(2.78)	(19.6)	(0.4)	(0.30)	(17.1)	(0.1)
Worst-ranked winner	1.76	5.8	0.5	1.69	7.5	0.7	0.56	16.7	0.9
	(2.69)	(13.2)	(0.5)	(2.72)	(15.3)	(0.5)	(1.09)	(14.4)	(0.2)

Panel B: Auction Outcomes

	N	Mean	SD	P25	Median	P75
Number of bidders	333	6.7	3.8	4.0	6.0	8.0
Number of winners	333	4.6	2.3	3.0	4.0	6.0
Bid/reserve price, all bidders	2,077	0.398	0.270	0.186	0.337	0.501
Bid/reserve price, best-ranked winner	333	0.310	0.243	0.117	0.239	0.438
Bid/reserve price, worst-ranked winner	333	0.455	0.245	0.292	0.415	0.599
Quantity share, all winners (%)	1,543	16.1	10.8	8.3	13.2	21.3
Quantity share, best-ranked winner (%)	333	31.2	9.8	24.8	31.0	35.8
Quantity share, worst-ranked winner (%)	333	13.7	8.5	7.4	12.9	17.9

Notes: Panel A reports outcomes for branded firms, certified and uncertified generic firms, and the best- and worst-ranked auction winners in each period. Pre-certification runs from the beginning of the sample through the quarter before the drug's first generic certification; post-certification, pre-auction runs from that certification quarter through the quarter before auction implementation; and post-auction begins with implementation. Generic firms are classified by their certification status at the end of each period. Best and worst winners are the firms with the lowest and highest winning bids, respectively. Standard deviations are in parentheses. Panel B reports the number of bidders and winners, bid/reserve ratios, and realized winner quantity shares. A winner's quantity share equals the procurement coverage rate multiplied by the share of national quantity in its assigned provinces.

2.4.2. *Auction Participation and Outcomes.* Panel B of Table 1 summarizes participation, bidding, and allocation in the centralized procurement auctions. We focus on rounds 2–9, for which we observe all submitted bids, excluding round 6, which covers only insulin products for diabetes treatment and follows different auction rules. The median auction has six bidders and four winners, so the program generally combines a certification requirement with a split-award auction rather than a single-winner rule. Across all submitted bids, the average bid is 39.8% of the reserve price; the best and worst successful bids average 31.0% and 45.5%, respectively. The procurement quantity allocated to the average winner equals 16.1% of the national drug quantity. For the best- and worst-ranked winners, the corresponding shares are 31.2% and 13.7%.

Appendix Figure A2 provides additional statistics across procurement waves. Panel A shows that bidder participation varies substantially across auctions. The number of bidders increases in later waves, indicating that more firms become certified and actively compete for procurement access. Panel B shows the lowest winning bids are consistently far below reserve prices across all auction waves.

3. Effects of Certification and Procurement on Competition

This section presents three sets of empirical results that illustrate how certification and procurement interact to shape competition in China's pharmaceutical market. We first show that BE certification affects perceived product quality but has only a limited effect on market access when firms compete in decentralized hospital markets. We then document substantial entry around procurement implementation, consistent with guaranteed market access increasing the returns to certification and entry. Finally, we show evidence that procurement lowers firms' cost of serving the market and allows for steep price reductions. Together, these findings suggest that certification and procurement are complementary policies: certification determines eligibility, while procurement creates the market access that makes certification valuable. These patterns motivate the structural framework developed in the next section.

3.1. **The Limited Role of Certification.** We begin by isolating the effect of BE certification before centralized procurement was introduced. This certification-only regime allows us to ask whether providing an informative signal of generic quality is sufficient, on its own, to expand a firm's presence in the market. If consumers and health-care providers value certification, approval may increase demand for the product, its price, or both. The relative responses of quantities and prices therefore

reveal how much certification alone changes firms' ability to compete in decentralized hospital markets.

We estimate how BE approval affects prices and quantities before procurement. The sample includes all procured drugs in waves 2 through 9. We aggregate sales and quantities to the national drug(d)-firm(i)-quarter(t) level and focus on the pre-procurement period for each drug. The treatment variable is BE approval for generic products. Branded products and generic drug-firm pairs that never obtain BE approval during the sample period serve as controls.

We estimate a stacked event-study specification. Let c index BE approval cohorts, defined by the quarter in which a generic product obtains approval. For each cohort, we stack drug-firm pairs first approved in quarter c ($\text{ApprovedInCohort}_{dic} = 1$) with the control products described above, and define event time as $t - c$. We estimate:

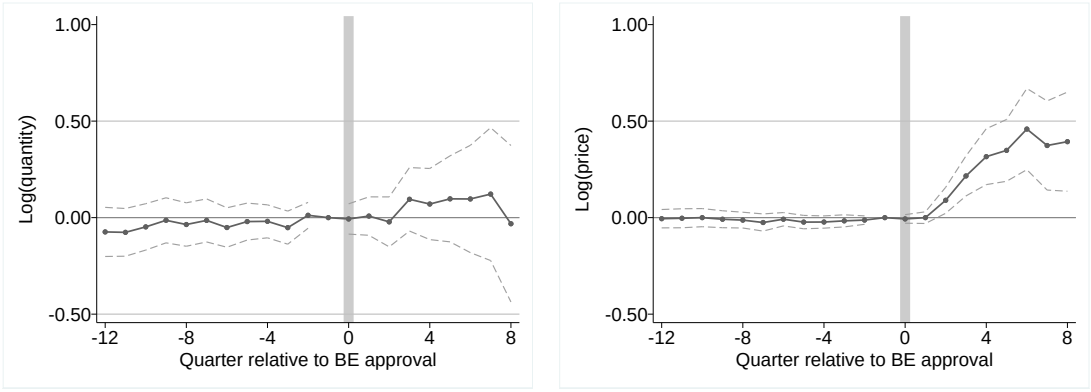
$$Y_{ditc} = \sum_{\ell=-12}^{8} \beta_\ell \mathbf{1}\{t - c = \ell\} \times \text{ApprovedInCohort}_{dic} + \lambda_{dic} + \lambda_{dtc} + \epsilon_{ditc}, \qquad (1)$$

where we normalize β_{-1} to zero and include endpoint bins for event times below -12 and above 8. The outcome Y_{ditc} is either the log quantity or log price. Drug-firm-cohort fixed effects (λ_{dic}) absorb permanent differences across products, and drug-quarter-cohort fixed effects (λ_{dtc}) absorb shocks common to all products of a drug in a quarter. The coefficients β_ℓ therefore compare an approved product's trajectory around its approval date with the contemporaneous trajectories of other products of the same drug. Standard errors are clustered at the drug level.

Figure 2 presents the event-study results. In both panels, the pre-approval coefficients are stable and close to zero, supporting the parallel trends assumption. Panel A shows little evidence that BE approval expands quantities: the coefficients remain small and statistically insignificant throughout the two years after approval. This suggests that certification by itself does not solve the market-access problem faced by generic firms.

Panel B shows a different pattern for prices. The price response is negligible in the approval quarter but builds steadily thereafter, reaching an increase of roughly 50 percent two years after approval. This gradual build-up is consistent with BE certification raising perceived product quality and consumers' willingness to pay as information about certification diffuses. Certification nevertheless closes only part of

FIGURE 2. Event Studies of BE Certification Effects



(A) Log(Quantity) (B) Log(Price)

Notes: Event studies are stacked by BE approval cohort and include drug-firm-cohort and drug-quarter-cohort fixed effects. The coefficient for event quarter -1 is normalized to zero; endpoint bins are included for event times before -12 and after 8. Standard errors are clustered at the drug level.

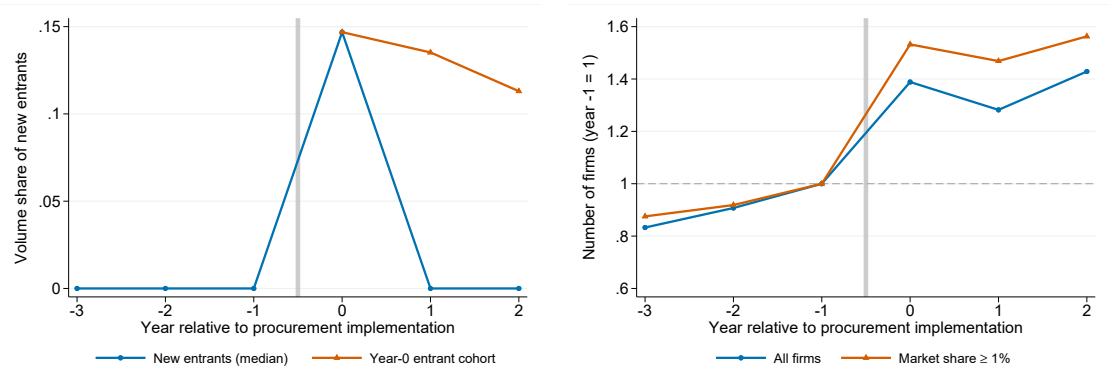
the generic-branded price gap: in the four quarters before procurement implementation, the median BE-certified generic sells at about 55% of the branded price of the same drug.

Together, these patterns indicate that BE certification has demand value, but that this value is expressed mainly through higher prices rather than through a large standalone increase in quantities in decentralized hospital markets.

3.2. **Centralized Procurement and Firm Entry.** Our second result is that the guaranteed market access created by centralized procurement is associated with substantial entry. For each procured drug, we define year 0 as the first year its centralized auction takes effect and, in each year relative to year 0, measure the market share of new entrants and the number of firms selling the drug. An entrant is a generic firm observed selling the drug for the first time in that year.

Panel A of Figure 3 plots the median volume share of new entrants across drugs by event year, together with the median share held by the year-0 entrant cohort in post-auction years. Entry is negligible before procurement: the median entrant volume share is zero in every pre-implementation year, and the mean is only about 1.4%. Entry jumps sharply at procurement implementation. New entrants capture 14.7% of the median drug's volume in year 0, and these entrants largely retains their market shares over the next two years. Entrants also account for a substantial

FIGURE 3. Entry Around Procurement Implementation



(A) Volume Share of New Entrants (B) Number of Firms

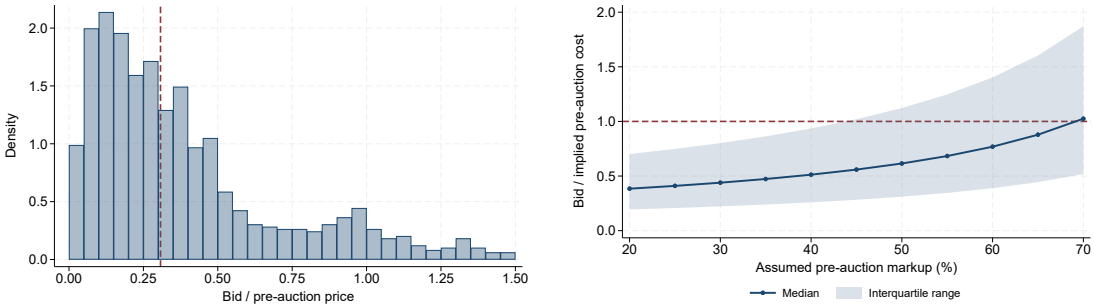
Notes: In both panels, year 0 marks the year of centralized procurement implementation. An entrant is a non-branded firm-product observed selling the drug for the first time; firms first appearing in the quarter before implementation are assigned to the year-0 cohort. Panel A plots the median across drugs of the volume share of that year's new entrants, and the median volume share of the year-0 entrant cohort in subsequent years. Panel B plots the average across drugs of the number of firms selling the drug, expressed relative to the same drug's count in year -1, for all firms and for firms with at least a 1% share of drug-year volume.

share of auction winners: 24.3% of winners first appear at or after procurement implementation, having never sold the product before. Because auction participation requires BE certification, these entrants enter already certified—procurement provides the market access that makes certification and entry worthwhile.

Panel B examines net entry. For each drug, we count the firms selling it in each event year and calculate the ratio relative to the firm count in year -1, the last year before auction implementation. The figure plots the average of this ratio across drugs, with and without excluding fringe firms that account for less than a 1% share of drug-year volume. The number of firms rises by 39% in year 0, and by 53% among firms above the 1% threshold. The increase persists through the following two years, suggesting that procurement is associated with net entry of firms. Firm counts also rise gradually before implementation even though entrant volume shares remain negligible (Panel A), consistent with new entrants facing substantial frictions in gaining market share before the auctions.

3.3. **Suggestive Evidence on Procurement-induced Cost Savings.** Our third result compares firms' auction bids with their own pre-auction prices and with a range

FIGURE 4. Bids Relative to Pre-Auction Prices and Implied Costs



(A) Bids Relative to Pre-Auction Prices (B) Bids Relative to Implied Costs

Notes: The sample is incumbent generic bidders (1,052 firm-drug pairs): non-branded firms that submit a bid and sell the drug before the auction. The pre-auction price is the firm's median price across provinces in the two years preceding procurement implementation. Panel A shows the distribution of the bid divided by the pre-auction price; the dashed line marks the median (0.31), and ratios above 1.5 (5.7% of bidders) are omitted from the display. Panel B plots the median and interquartile range across bidders of the bid divided by the implied pre-auction marginal cost based on a range of mark-up assumptions.

of marginal costs implied by markup estimates from the prior literature. The sample is incumbent generic bidders: generic firms that submit a bid and sell the drug before the auction (1,052 firm-drug pairs). For each firm-drug pair, we measure the pre-auction price as the firm's median price across provinces in the two years preceding implementation.

Panel A of Figure 4 plots the distribution of the ratio of a firm's bid to its own pre-auction price. Bids sit far below the prices the same firms charged before the auction: the median ratio is 0.31, and three quarters of firms bid less than 56% of their pre-auction price.

The low bid-price ratios are difficult to reconcile with unchanged marginal costs, under reasonable assumptions on markups. Panel B plots the median and interquartile range across products of the ratio of bids to the marginal costs implied by a range of assumptions on markups. Markup estimates for generic drugs in the prior literature typically range from 10% to 50%.[2] Over this range, the median firm bids

[2]Using provincial procurement prices for lipid-lowering drugs in China, Huang et al. (2026) estimate a markup of 13% for generic products. Dubois and Lasio (2018) estimate markups of 23–47% for generic anti-ulcer drugs in France on ex-manufacturer prices. Estimates based on retail prices, which embed distribution margins, are higher: Chaudhuri et al. (2006) report upper-bound markups of 34–59% for domestic fluoroquinolones in India.

well below its implied pre-auction marginal cost. The median bid would reach parity with implied cost only if the markup exceeded 70%.

The large gap between auction bids and marginal costs implied by pre-auction prices suggests that firms' cost of serving the market changes substantially under centralized procurement. If pre-auction costs remained unchanged, the observed bids would imply implausibly low or negative margins for many firms under conventional markup estimates. A more natural interpretation is that winning procurement reduces costs that firms previously incurred to reach and serve demand. In particular, centralized procurement provides direct access to hospitals and patients, reducing expenditures associated with marketing to individual hospitals and physicians, distribution, and market development; the resulting certainty of demand may also lower financing and inventory costs. We allow for these procurement-induced cost savings in the structural model rather than interpreting low bids as evidence of persistent below-cost pricing.

4. A Model of Quality Certification, Entry, and Procurement

We develop a model of generic drug entry, certification, centralized procurement auctions, and post-auction price competition. The model captures three key features of the institutional environment. First, firms differ in unobserved product quality, and bioequivalence (BE) certification provides an imperfect signal of quality to consumers. Second, centralized procurement auctions allocate market access across certified firms and may restrict consumer choice through demand guarantees. Third, procurement outcomes affect downstream product-market competition because winning firms must supply consumers at the procurement price and may benefit from cost reductions associated with larger market scale. The model links firms' entry, certification, bidding, and pricing decisions in a unified framework and highlights how procurement incentives interact with downstream competition and quality signaling.

4.1. **Setup.**

There is one branded incumbent firm and J potential entrants producing generic drugs. Each potential entrant has a predetermined quality type $\tau \in \{H, L\}$, where a fraction $\phi \in (0, 1)$ of potential entrants are high quality, and the remaining fraction $1 - \phi$ are low quality. Quality is privately observed by the firm but not by consumers, and affects both the demand side, through consumer beliefs, and the supply side, through firms' cost distributions and certification outcomes.

Each potential entrant draws an entry cost e from distribution F_e and decides whether to enter the market. Firms that enter produce a drug of their predetermined quality type τ.

Entrants may choose whether to take a bioequivalence (BE) certification test. The branded incumbent does not require certification by definition, since the purpose of the BE test is to verify equivalence to the branded drug. Firm i incurs a testing cost η_i drawn from distribution F_η. The outcome of this stage is summarized by an observable status variable

$$v_i \in \{B, C, U\},$$

where B denotes the branded incumbent, C denotes a certified generic firm, and U denotes an uncertified generic firm. Consumers observe v_i but cannot distinguish between firms that did not take the test and firms that took the test but failed.[3]

Let $T = 1$ denote taking the BE test. We assume that the BE test is informative but imperfect:

$$\Pr(v = C \mid T = 1, H) = 1, \quad \Pr(v = C \mid T = 1, L) = \psi \in [0, 1). \tag{2}$$

Thus, high-quality firms pass the test with certainty, while low-quality firms pass with probability ψ. The parameter ψ captures the imperfection of the certification technology: when $\psi = 0$, certification perfectly separates high- and low-quality firms; when $\psi > 0$, some low-quality firms may obtain certification.

Certified generic firms participate in the centralized procurement auction.[4] For simplicity, we assume that participation is costless. Auction participants privately observe a signal c for their marginal production cost drawn from the type-specific distribution $F_{c|\tau}$ and submit bids in the procurement auction. We assume that the density $f_{c|\tau}$ is log-concave.

Once procurement outcomes are realized, firms then observe their marginal production costs and the realization of the product-level demand shocks, after which they compete in prices within each market. The procurement winner's price is fixed at its auction bid, while all other firms choose prices in a Bertrand–Nash equilibrium. Profits are subsequently realized.

[3]Equivalently, because test outcomes can be credibly disclosed, firms reveal their test result in equilibrium only when they pass.

[4]We assume that only certified generic firms are eligible to participate in the procurement auction. While branded firms are formally eligible, reserve prices are typically set sufficiently low that branded firms rarely participate or submit competitive bids. We therefore abstract from branded-firm participation in the auction.

4.1.1. *Consumer Demand.* We model drug choice as a joint decision of the patient and the physician: conditional on a molecule, the patient chooses among available products under physician supervision. Substitution toward other molecules or no treatment is absorbed into the outside option.

Consumer j's indirect utility from product i in market m is

$$u_{ijm} = \alpha \log(\kappa_m p_{im}) + \omega \mu_{v_i} + \xi_i + \varepsilon_{ijm}, \tag{3}$$

where ε_{ijm} is an i.i.d. Type I extreme value taste shock. The outside option, whose mean utility is normalized to zero, includes other molecules or no treatment. The utility depends on the consumer's out-of-pocket payment, the list price p_{im} multiplied by the copay rate κ_m.[5] The term μ_{v_i} is consumers' equilibrium belief that product i is high quality:

$$\mu_{v_i} = \Pr(\tau_i = H \mid v_i).$$

The branded product is known to be high quality, so $\mu_B = 1$. For generic products, the posterior beliefs μ_C and μ_U are determined in equilibrium by firms' entry and testing decisions and the precision of the certification test.

The term ξ_i captures product appeal that is unrelated to quality. For generic products, ξ_i is drawn independently from a common distribution F_ξ, regardless of quality type or certification status. For the branded incumbent, we leave the demand shock ξ_B unrestricted, allowing it to capture brand equity — brand reputation, prescribing habit, and detailing networks accumulated before the reforms — that is distinct from perceived quality. A generic product that consumers believe to be high quality with certainty therefore attains the branded product's perceived quality, but not its brand advantage. Demand shocks are drawn after entry and procurement outcomes but before downstream price competition, and their realizations are observed by all firms prior to pricing. This timing preserves symmetry among auction participants while allowing for ex-post vertical differentiation across products.

We model the quantity guarantee in market m as restricting the choice set for a fraction $\lambda_m \in (0, 1)$ of consumers. These consumers may choose only between the procurement winner and the outside option, while the remaining consumers choose freely among all available products. This formulation reflects how quantity guarantees

[5]The copay rate combines the drug's formulary tier with province-level coinsurance: uncovered drugs are paid fully out of pocket, tier B drugs carry a 10–20% copay before coinsurance applies, and tier A drugs are subject only to coinsurance. For covered drugs, the average out-of-pocket share is approximately 25%. In estimation, we measure the copay rate at the drug-province-year level.

are implemented in practice: hospitals receive purchase quotas for the winning product, some hospitals stock only the winner, and physicians are strongly incentivized to prescribe it. Because quota enforcement varies across provinces and contracts, we do not impose the statutory guarantee rates; instead, we treat λ_m as a market-specific primitive and estimate it from post-procurement demand.

Let $Q_{im}^R(p_m^w, \xi_i)$ denote the market share from restricted consumers for the procurement winner i, where p_m^w is the winning price. Let $Q_{im}^U(p_{im}, v_i, \xi_i, \mathbf{p}_{m,-i}, \mathbf{v}_{-i}, \boldsymbol{\xi}_{-i})$ denote the market share for product i — winner or not — from unrestricted consumers, which depends on product i's own price, certification status, and demand shock, as well as those of competing products in market m.[6]

4.1.2. *Firm Marginal Costs.* Let mc_{im} denote the marginal cost of producing product i to serve in market m. At the bidding stage, firms do not yet observe realized marginal costs. Instead, product i's firm forms expectations according to

$$mc_{im} = \frac{c_i}{1 + \sum_{m'} \gamma_{m'} \mathbf{1}\{w_{m'} = i\}} + \delta_m + \epsilon_{im}, \tag{4}$$

where c_i is the firm's private cost signal, observed before bidding, $w_{m'}$ denotes the winner of market m', and ϵ_{im} is a cost shock, i.i.d. across products and markets, drawn from F_ϵ with $\mathbb{E}(\epsilon_{im}) = 0$. After the procurement auction and before the pricing stage, realized marginal costs become publicly observed.[7]

The parameters $\gamma_m \geq 0$ and δ_m are publicly known and common across firms. The term $\sum_{m'} \gamma_{m'} \mathbf{1}\{w_{m'} = i\}$ captures the effect of procurement outcomes on production efficiency: winning additional provinces may reduce marginal costs through economies of scale, production complementarities, or distribution efficiencies, and such effects may depend on province attributes like market size. The parameter δ_m captures market-specific cost determinants common to all firms.

4.1.3. *Centralized Procurement Auction.* The procurement auction is conducted in a split-award format. Each participating firm submits a bid b_i, which specifies the price the firm commits to charge in any assigned market. Before bidding, the auctioneer announces a ceiling price $\bar{p}$ and a market assignment rule $\{r_m\}_{m \in M}$, where r_m denotes

[6]Restricted consumers choose between the winner and the outside option. Because every procurement winner is a certified generic product, demand from restricted consumers does not depend separately on certification status.

[7]This informational assumption is standard in Bertrand–Nash pricing games, although the analysis can be extended to settings with incomplete information.

the bid rank (in ascending order) assigned to market m. Firms with bids above the ceiling price are excluded from assignment.

The split-award structure allows different markets to be assigned to firms with different bid rankings.[8] For example, larger or strategically important markets may be allocated to the lowest bidder, while smaller markets may be assigned to lower-ranked firms. We assume that the assignment rule satisfies a monotonicity condition: firms with better bid rankings weakly prefer their assigned set of markets to the set assigned to firms with worse rankings. Intuitively, lower bids are rewarded with more desirable market allocations.

A single bid simultaneously determines assignments across all markets. Because procurement winners subsequently compete in downstream product markets, bidding incentives depend not only on the probability of winning markets, but also on how procurement outcomes affect downstream pricing incentives and production costs.

4.2. **Equilibrium.** We characterize equilibrium by solving the game backwards.

4.2.1. *Pricing Strategies.* In the final stage, firms compete in prices according to a Bertrand–Nash equilibrium. This pricing game takes as given the realized marginal costs $\mathbf{mc}$ and demand shocks $\boldsymbol{\xi}$, certification statuses $\mathbf{v}$, and procurement outcomes in each market, including the winner identity w_m and the winning price p_m^w. To simplify notation, define the post-auction state in market m for product i as

$$\Omega_{im} = (v_i, mc_{im}, \xi_i, \mathbf{v}_{-i}, \mathbf{mc}_{m,-i}, \boldsymbol{\xi}_{-i}, w_m, p_m^w).$$

Given rivals' prices $\mathbf{p}_{m,-i}$, the post-auction profit in market m for product i is

$$\pi_m(\mathbf{p}_{m,-i}, \Omega_{im}) = \begin{cases} (p_m^w - mc_{im})\Big[\lambda_m Q_{im}^R(p_m^w, \xi_i) \\ \qquad + (1-\lambda_m) Q_{im}^U(p_m^w, v_i, \xi_i, \mathbf{p}_{m,-i}, \mathbf{v}_{-i}, \boldsymbol{\xi}_{-i})\Big], & \text{if } i = w_m, \\ \max\limits_{p_i} (p_i - mc_{im})(1-\lambda_m) \\ \qquad \times Q_{im}^U(p_i, v_i, \xi_i, \mathbf{p}_{m,-i}, \mathbf{v}_{-i}, \boldsymbol{\xi}_{-i}), & \text{if } i \neq w_m. \end{cases} \tag{5}$$

[8] In practice, winners choose markets sequentially in ascending order of their bids. We abstract from this choice by assuming a predetermined assignment rule, which is equivalent to assuming that firms share a common ranking of market attractiveness. This rules out firm-specific market preferences, which we do not view as central to our analysis.

Procurement winners serve both restricted and unrestricted consumers, whereas non-winning firms serve only unrestricted consumers. Standard logit-demand assumptions imply existence and uniqueness of equilibrium prices as long as $\alpha < -1$.[9]

4.2.2. *Bidding Strategies.* We focus on symmetric monotone equilibria in which firms follow a bidding strategy

$$\beta(c \mid n^C, n^U),$$

where c denotes the firm's private cost signal, n^C denotes the number of auction participants, and n^U denotes the number of uncertified firms.[10] Since auction participation is costless, the number of participants equals the number of certified generic firms. The number of uncertified firms also affects bidding strategies because it influences the pricing equilibrium in the unrestricted market. Certification outcomes are publicly observed prior to bidding, so firms observe n^C and n^U before submitting bids.

Bids affect continuation payoffs through two channels. First, lowering the bid changes the probability of obtaining desirable market assignments. Second, conditional on winning, the firm's bid becomes its downstream market price, thereby affecting future product-market profits. Moreover, procurement outcomes also affect firm-wide production efficiency through (4).

Consider product i with cost signal c_i. Suppose rival firms follow equilibrium strategy $\beta(\cdot \mid \cdot, \cdot)$. Let $\mathbf{c}_{-i}$ denote rivals' cost signals, and let $\mathbf{c}^{(k)}$ denote the kth order statistic of rival costs. Let

$$\pi_m^w(b, c_i, \mathbf{b}_{-i}, n^C, n^U)$$

denote the expected continuation payoff for product i in market m conditional on winning market m with bid b. This continuation value is derived from the post-auction profit function in (5), integrating over realized marginal costs, product-level demand shocks, and subsequently equilibrium pricing outcomes. Since the procurement winner must charge its bid price in the downstream market, the continuation profit depends directly on the submitted bid. The bid also affects winner allocations across all provinces and, through scale economies, the winner's effective cost.

[9]In the data, not all firms participate in every market. The model rationalizes non-participation through equilibrium demand approaching zero for some products.

[10]The privately observed quality type τ affects the distribution of cost signals through $F_{c|\tau}$, but conditional on the realized signal c, it does not directly affect downstream marginal costs. In addition, all bidders have $v_i = C$. Hence, bidding strategies depend on the realized signal rather than directly on τ.

Similarly, let

$$\pi_m^{\ell}(b, c_i, \mathbf{b}_{-i}, n^C, n^U)$$

denote the expected continuation payoff from product i conditional on not being assigned market m. This payoff also depends on the submitted bid through scale economies. Given rivals' strategies, the expected payoff from bid b for the firm that produces product i is

$$\begin{aligned}\pi^a(b, c_i, n^C, n^U; \beta) = \sum_{m \in M} N_m \Bigg\{ & \Pr\left(\beta^{-1}(b \mid n^C, n^U) \in \left[c^{(r_m-1)}, c^{(r_m)}\right]\right) \\ & \times \mathbb{E}_{\mathbf{c}_{-i}}\Big[\pi_m^w\big(b, c_i, \beta(\mathbf{c}_{-i} \mid n^C, n^U), n^C, n^U\big) \\ & \qquad - \pi_m^{\ell}\big(b, c_i, \beta(\mathbf{c}_{-i} \mid n^C, n^U), n^C, n^U\big) \\ & \qquad \Big| \beta^{-1}(b \mid n^C, n^U) \in \left[c^{(r_m-1)}, c^{(r_m)}\right]\Big] \\ & + \mathbb{E}_{\mathbf{c}_{-i}}\Big[\pi_m^{\ell}\big(b, c_i, \beta(\mathbf{c}_{-i} \mid n^C, n^U), n^C, n^U\big)\Big] \Bigg\}, \end{aligned} \tag{6}$$

where $\beta^{-1}(b) \in \left[c^{(r_m-1)}, c^{(r_m)}\right]$ corresponds to bid b being the r_m-th lowest bid and thus getting assigned market m, and N_m is the market size of province m. We adopt the convention $c^{(0)} = -\infty$ and $c^{(n^C)} = \infty$. In equilibrium,

$$\beta(c_i \mid n^C, n^U) \in \arg\max_{b \leq \bar{p}} \pi^a(b, c_i, n^C, n^U; \beta). \tag{7}$$

The equilibrium bidding condition captures a tradeoff between the winning probability and profit conditional on winning. A lower bid increases the likelihood of receiving a more valuable bundle of markets and associated guaranteed demand, but it also reduces the price and per-unit margin if the firm wins. When the market guarantee is sufficiently large, firms are willing to accept smaller margins to improve their rank, potentially generating large price reductions. This incentive may be especially strong when winning also produces scale economies, because losing entails both forgone guaranteed demand and forgone cost reductions. In the Appendix, we provide sufficient conditions under which the symmetric monotonic equilibrium bidding strategy is unique.

4.2.3. *Test-Taking Strategies.* Let $\pi_{v,\tau}$ denote the expected payoff for a type-τ product with certification outcome v. For a product that passes the BE test ($v_i = C$), expected profit is

$$\pi_{C,\tau} = \mathbb{E}_{(c_i, n^C, n^U)}\left[\pi^a(\beta(c_i \mid n^C, n^U), c_i, n^C, n^U; \beta) \mid \tau\right]. \tag{8}$$

For products that either fail the BE test or do not have test results ($v_i = U$),

$$\pi_{U,\tau} = \sum_{m \in M} N_m \mathbb{E}\left[\pi_m(\mathbf{p}_{m,-i}, v_i = U, mc_{im}, \xi_i, \mathbf{v}_{-i}, \mathbf{mc}_{m,-i}, \boldsymbol{\xi}_{-i}, w_m, p_m^w) \mid \tau\right], \tag{9}$$

where expectations are taken over certification outcomes of rival products, procurement outcomes, realized marginal costs, and product-level demand shocks.

A firm with a type-τ product takes the BE test if and only if expected gains exceed testing cost:

$$\pi_{C,\tau} \Pr(v = C \mid T = 1, \tau) + \pi_{U,\tau} \Pr(v = U \mid T = 1, \tau) - \eta_i \geq \pi_{U,\tau}.$$

Using (2), the equilibrium probability that a firm with a type-τ product takes the test is

$$\rho_\tau = \begin{cases} F_\eta(\pi_{C,H} - \pi_{U,H}) & \text{if } \tau = H, \\ F_\eta\big(\psi(\pi_{C,L} - \pi_{U,L})\big) & \text{if } \tau = L. \end{cases} \tag{10}$$

Firms with a low-quality product take the test with positive probability only when $\psi > 0$, since otherwise passing is impossible. Therefore,

$$\Pr(v_i = C | \tau_i = H) = \rho_H, \ \Pr(v_i = C | \tau_i = L) = \rho_L \psi. \tag{11}$$

4.2.4. *Entry Decisions.* The expected profit from entering for a firm with a high-quality product is

$$\pi_H = \mathbb{E}_\eta\big(\pi_{C,H} - \eta \mid \eta \leq \pi_{C,H} - \pi_{U,H}\big)\rho_H + \pi_{U,H}\big(1 - \rho_H\big). \tag{12}$$

The expected profit from entering for a firm with a low-quality product is

$$\pi_L = \mathbb{E}_\eta\Big[\psi\pi_{C,L} - \eta \mid \eta \leq \psi(\pi_{C,L} - \pi_{U,L})\Big]\rho_L + \pi_{U,L}(1 - \rho_L\psi). \tag{13}$$

The probability of entering for a firm with a type-τ product is

$$q_\tau = F_e(\pi_\tau). \tag{14}$$

Therefore, the probability that an entrant is of high type is $\phi q_H/(\phi q_H + (1-\phi)q_L)$, and equilibrium posterior beliefs are

$$\mu_C = \frac{\rho_H \phi q_H}{\rho_H \phi q_H + \rho_L \psi (1-\phi) q_L}, \tag{15}$$

$$\mu_U = \frac{(1-\rho_H)\phi q_H}{(1-\rho_H)\phi q_H + \big[1 - \rho_L + \rho_L(1-\psi)\big](1-\phi)q_L}, \tag{16}$$

which enter demand through (3).

Taken together, an equilibrium consists of entry probabilities by quality type, certification probabilities, posterior beliefs, a symmetric monotone bidding strategy, and post-auction pricing strategies such that all decisions are sequentially optimal and beliefs are consistent with firms' equilibrium behavior. Given posterior beliefs about certified and uncertified generics, downstream prices solve the Bertrand–Nash pricing game induced by certification statuses, realized costs, demand shocks, and procurement outcomes. Anticipating these continuation payoffs, certified firms choose bids that trade off a lower procurement price against a higher probability of securing market access and realizing scale-related cost advantages. These bidding incentives determine the expected value of certification, which in turn governs firms' test-taking decisions. Finally, expected payoffs from certification and non-certification determine the profitability of entry by high- and low-quality firms, and the resulting composition of entrants pins down consumers' posterior beliefs through Bayes' rule.

The equilibrium therefore captures a feedback loop between information and market access. Certification affects demand by changing consumer beliefs, but it also affects supply-side incentives by determining eligibility for procurement. Procurement, in turn, changes both the value of certification and the nature of downstream competition: winners obtain guaranteed demand and potential cost reductions, but must commit to the procurement price, while non-winners compete only for unrestricted consumers. Through this interaction, the model allows procurement design to shape not only prices conditional on entry, but also the equilibrium composition of entrants and the informational content of certification itself.

4.3. **The Role of Certification and Procurement.** We now describe the no certification and no procurement benchmark and use it to clarify the roles of BE certification and centralized procurement in the model. In this benchmark, generic firms cannot signal product quality, and consumers cannot distinguish high- and low-quality generic drugs. Generic quality is therefore pooled from consumers' perspective. Let μ^0 denote the probability that an entering generic firm is high quality. While the branded incumbent is perceived as high quality, all generic firms share the common belief μ^0.

Consumer utility in market m is therefore

$$u^0_{ijm} = \alpha \log(\kappa_m p_{im}) + \omega \mu^0_i + \xi_i + \varepsilon_{ijm},$$

where $\mu^0_i = 1$ for the branded drug and $\mu^0_i = \mu^0$ for generic drugs. Since there is no procurement, all consumers choose among all available products and the outside

option. Given realized marginal costs, firms compete in a Bertrand–Nash pricing game.

Let π_τ^0 denote the expected operating profit of a type-τ entrant in this regime. Entry satisfies

$$q_\tau^0 = F_e(\pi_\tau^0).$$

The equilibrium belief μ^0 is determined by Bayes' rule:

$$\mu^0 = \frac{\phi F_e(\pi_H^0)}{\phi F_e(\pi_H^0) + (1-\phi) F_e(\pi_L^0)}.$$

Because consumers cannot condition their beliefs on any observable signal of quality, high-quality generic firms receive no direct demand advantage relative to low-quality generic firms. As a result, the incentive to enter as a high-quality producer is limited by the fact that quality is not separately rewarded in the downstream market.

BE certification changes this environment by introducing an observable signal of quality. Once certification is available, consumers can condition beliefs on certification status. Certified generics are perceived to be higher quality than uncertified generics, with posterior beliefs determined by the testing technology, the composition of entrants, and type-specific testing probabilities. Certification therefore creates an information channel: it allows high-quality firms to separate themselves from low-quality firms, though imperfectly, raises demand for certified products, and changes entry and testing incentives.

Centralized procurement introduces a second channel. In addition to shaping consumer beliefs, certification determines eligibility for procurement auctions and subsequent access to the downstream market. Without procurement, the market power of incumbent branded products may limit the returns to entering with a high-quality product, particularly because such products typically have higher marginal costs. Procurement places certified firms on a more level playing field. Winning a procurement contract can provide guaranteed demand and generate scale-related cost advantages, thereby affecting firms' incentives to enter, obtain certification, and compete in downstream markets. Thus, certification and procurement interact in the model: certification provides information about product quality, while procurement changes the returns to certification by linking it to market access.

5. Identification and Estimation

This section describes how the model primitives are identified from the data and how the parameters are estimated. The data consist of a panel of off-patent drug markets observed before and after the introduction of BE certification and centralized procurement. For each drug, we observe the identities of firms that enter and remain active, the identities of firms that take the BE test, and the identities of firms that obtain certification. For drugs included in centralized procurement, we observe all bids submitted in the auction, the identity of the winners, the winning prices, and the market assignment induced by the procurement rule. We also observe firm-level prices and quantities across markets over time, allowing us to trace how certification and procurement affect downstream market shares, prices, and entry.

The data span multiple policy regimes. All drugs begin in a regime with neither BE certification nor centralized procurement. This pre-certification period, covering 2013 to around 2016, provides information on market outcomes when generic quality is pooled from consumers' perspective. Starting in 2017, most off-patent drugs gradually became subject to the BE certification requirement. A subset of these drugs were subsequently included in centralized procurement beginning in 2019, across multiple procurement waves. Thus, some drugs transition from no certification to certification and then to procurement, while others transition only from no certification to certification.

Because the model is static, we interpret each regime as a long-run equilibrium under a fixed policy environment. We therefore use the pre-certification and post-procurement periods as the primary empirical counterparts to the model's long-run equilibria. The intermediate certification-only period is less suitable as a direct empirical counterpart to the certification-only regime, because firms may have anticipated future procurement waves and because entry, testing, and certification decisions take time.[11]

5.1. **Observables.** For each drug, we observe annual product-level market shares by province during both the pre-certification and post-procurement periods. Let s_{imt} denote the observed market share of product i in province m and year t. We allow observed market shares to be measured with multiplicative error. In particular, let

[11]We instead treat the certification-only regime primarily as a model-based counterfactual that isolates the information role of BE certification absent procurement-induced market access.

s^*_{imt} denote the true market share and suppose that

$$s_{imt} = \frac{s^*_{imt} \exp(e_{imt})}{\sum_{j=0}^{N_{mt}} s^*_{jmt} \exp(e_{jmt})}, \tag{17}$$

where $j = 0$ denotes the outside option and N_{mt} is the number of products available in province m and year t.

Defining the measurement error in the log-share ratio relative to the outside option as

$$u_{imt} \equiv e_{imt} - e_{0mt},$$

we obtain

$$\log\left(\frac{s_{imt}}{s_{0mt}}\right) = \log\left(\frac{s^*_{imt}}{s^*_{0mt}}\right) + u_{imt}. \tag{18}$$

We assume that u_{imt} has conditional mean zero given the observed product and market characteristics.[12] Because the log-share-ratio errors share the common component e_{0mt}, we allow u_{imt} to be correlated across products within the same province-year.

In addition to market shares, we observe the price p_{imt}, the copay rates κ_{mt}, and certification status v_{it} of each product. We also observe the procurement assignment rule $\{r_m\}_{m\in M}$ and the procurement auction outcomes, including the identity of the winning product in each province, w_m, and the corresponding winning bid, p^{w}_m, as well as the bid b_i submitted by each participating product. For each product, we observe whether it underwent bioequivalence certification testing and, conditional on being tested, whether it passed or failed. Finally, we assume that the set of potential entrants J is observed, allowing us to construct the sample entry probability.

5.2. **Model Primitives.** The primitives of the model consist of: (i) the entry environment, including the entry-cost distribution F_e and the fraction of high-quality firms among potential entrants, ϕ; (ii) the certification technology, including the testing-cost distribution F_η and the probability that a low-quality firm passes the BE test, ψ; (iii) demand-side parameters, including price sensitivity α, quality sensitivity ω, the distribution of product-level demand shocks F_ξ, and the market-specific demand-guarantee rates λ_m; and (iv) supply-side parameters, including the type-specific cost-signal distributions $F_{c|\tau}$, the scale-economy parameters $\{\gamma_m\}_m$, market-specific cost shifters $\{\delta_m\}_m$, and the distribution of idiosyncratic cost shocks F_ϵ.

[12]The primitive measurement shocks need not individually have conditional mean zero. Their conditional means need only be common across products within a province-year.

We parameterize F_ξ by its mean and variance, and assume that the idiosyncratic quality component has the same conditional mean across certification statuses and the pre-certification regime

$$\mathbb{E}(\xi_i|v_i = C) = \mathbb{E}(\xi_i|v_i = U) = \mathbb{E}(\xi_i^0) = \bar{\xi}, \tag{19}$$

and the same conditional variance σ_ξ^2. In addition, we parameterize the testing- and entry-cost distributions by their respective means and variances, $\theta_\eta \in \mathbb{R}^2$ and $\theta_e \in \mathbb{R}^2$. Lastly, we also parametrize province-specific scale effects, γ_m, by:

$$\gamma_m = \theta_\gamma \text{size}_m,$$

where size_m represents province size tiers and we require $\theta_\gamma \geq 0$.

5.3. **Identification Strategy.** The main challenge is that we do not observe the quality type of each generic firm. As a result, the consumer's belief μ_v is not observed and is in equilibrium determined by testing error ψ and two other equilibrium objects, the type-specific entry probabilities and testing probabilities, neither of which are observed. Our identification strategy exploits our auxiliary data on BE testing and the fact that we observe the market equilibria before and after the introduction of the centralized procurement auctions. Our strategy proceeds in four steps, and the estimation closely follows each step. Except for the common price coefficient α, which is identified from pooled pre-procurement variation across drugs in Step 1, the remaining identification argument does not rely on variation across drugs. We therefore present Steps 2–4 for a single drug.

5.3.1. *Step 1: Price elasticity.* We identify the price coefficient α from pre-procurement data. Although procurement induces large price changes, they cannot be used to estimate the price elasticity: procurement price cuts coincide with the imposition of the demand guarantee, so the demand response to prices cannot be separated from the effect of restricted choice sets. We therefore rely on pre-procurement variation, using changes in insurance coverage to instrument for out-of-pocket prices.[13, 14]

[13]We instrument the out-of-pocket price $\kappa_{mt}p_{imt}$ with the drug's insurance-formulary tier in China's Basic Insurance Drug Catalog, and we exclude the years in which a drug's tier changes. Before 2017, provinces supplemented the national catalog with their own tier B lists; from 2019 onward, the formulary is unified at the national level. Both national revisions and provincial supplements generate tier changes during our sample period.

[14]Controlling for product-province fixed effects ensures that α is identified only from changes in a product's insurance tier within the same province over time.

5.3.2. *Step 2: Demand guarantee enforcement and product quality.* Given α, we invert market shares in the post-procurement regime to recover product qualities and the demand-guarantee rates. Define the composite product quality

$$\widetilde{\xi}_i \equiv \omega\mu_{v_i} + \xi_i,$$

for the post-auction period, and

$$\widetilde{\xi}_i^0 \equiv \omega\mu_i^0 + \xi_i,$$

for the pre-certification period. Using Equation (18), the observed log share ratios in the post-auction period satisfy

$$\log(s_{imt}/s_{0mt}) = \alpha\log(\kappa_{mt}p_{imt}) + \widetilde{\xi}_i + \mathbf{1}\{i \neq w_m\}\, z_{mt} + u_{imt}, \tag{20}$$

where

$$z_{mt} = -\log\left(\frac{s^*_{0mt} + s^*_{w_m mt}}{s^*_{0mt} + s^*_{w_m mt} - \lambda_{mt}}\right). \tag{21}$$

Non-winners sell only to unrestricted consumers, so their observed shares fall short of the logit benchmark, and the common shortfall is the wedge z_{mt}. The wedge thus measures the severity of the restriction: when $\lambda_{mt} = 0$, $z_{mt} = 0$ and non-winners face no disadvantage; as $\lambda_{mt} \to 1$, $z_{mt} \to -\infty$, as the entire market is restricted and non-winners retain no demand.

Since u_{imt} represents measurement error in observed shares rather than an economic demand shock observed by firms, we assume that it is mean-independent of the equilibrium price vector and the variables entering the demand-guarantee wedge. Under this assumption, Equation (20) identifies $\{\widetilde{\xi}_i\}_i$ and $\{z_{mt}\}_{m,t}$, provided that the product–province network is sufficiently connected. Winner observations anchor product quality because their demand inversions do not contain the wedge z_{mt}.

Having identified α, $\{\widetilde{\xi}_i\}_i$, and $\{z_{mt}\}_{m,t}$, we can recover the true outside-option share s^*_{0m} and winner share $s^*_{w_m m}$ in each province from the corresponding market share equations for the post-auction period. In particular, the identified demand parameters determine the winner's relative attractiveness, while the province-specific wedge adjusts the relative attractiveness of all non-winning products. Substituting these recovered shares and the identified wedge z_{mt} into (21) then identifies each λ_{mt}.

Lastly, using the identified composite product qualities, $\{\widetilde{\xi}_i^0, \widetilde{\xi}_i\}_i$, we recover their mean and variance in the pre-certification period and conditional on certification status in the post-auction period. The conditional variances identify σ^2_ξ, the variance of the idiosyncratic quality component ξ_i.

5.3.3. *Step 3: Marginal costs.* We next recover realized marginal costs, mc_{im}, from the downstream pricing first-order conditions of products that do not win the procurement contract in a given province. Winners are required to charge their winning bids, whereas non-winners remain free to choose their downstream prices. The non-winners' pricing first-order conditions therefore provide an inversion from observed prices and market shares to marginal costs.

We then decompose marginal cost into a product-level component, a province-specific cost shifter, and a product–province-specific cost shock. Define

$$\widetilde{c}_i \equiv \frac{c_i}{1 + \sum_{m'} \gamma_{m'} \mathbf{1}\{w_{m'} = i\}},$$

where c_i is product i's cost signal and the denominator captures the reduction in marginal cost generated by the set of provinces won by product i. Marginal cost in province m and year t can then be written as the sum of $\widetilde{c}_i$, the province-specific cost shifter δ_m, and the idiosyncratic cost shock ϵ_{im}. The pricing inversions for non-winners identify $\widetilde{c}_i$ and δ_m, subject to a normalization that $\sum_m \delta_m = 0$.

To identify the scale-economy parameters θ_γ, we use the optimality of firms' procurement bids. A firm's bid affects not only the price it receives if it wins, but also the set of provinces it is awarded and, through the scale-economy term, its marginal costs in all downstream markets. For any candidate value of θ_γ, we can therefore compute the continuation profit associated with each possible procurement outcome and solve for the corresponding equilibrium bidding strategy. The scale-economy parameters are identified by requiring the observed bids to satisfy these bidding optimality conditions.

Having recovered θ_γ, we can invert the scale-economy adjustment and recover each product's underlying cost signal c_i from $\widetilde{c}_i$ and its observed set of winning provinces. We then identify the cost-signal distributions of certified and uncertified firms, $G_{c|v}$ for $v \in \{C, U\}$. Because a failed certification test reveals that the firm is a low type, and because we observe post-test cost signals or pricing information for firms that fail, these observations identify $F_{c|L}$.[15] In addition, the residuals from the cost decomposition identify F_ϵ.

5.3.4. *Step 4: Testing, entry, and type distributions.* The remaining objects to be identified are: (i) the parameters governing entry and testing, θ_e, θ_η, ψ, and ϕ; (ii) the

[15]This interpretation relies on the model timing under which c_i is realized after the testing decision, so that test participation does not select firms on their realized cost signals.

demand-side parameters, $\bar{\xi}$ and ω; and (iii) the cost-signal distribution of high-quality firms, $F_{c|H}$. The last object is pinned down once the finite-dimensional parameters in (i) and (ii) are identified. In particular,

$$F_{c|H}(c) = F_{c|L}(c) + \frac{G_{c|C}(c) - F_{c|L}(c)}{\mu_C}, \tag{22}$$

for any c in the support of the cost signal. Step 3 identifies $F_{c|L}(\cdot)$ and $G_{c|C}(\cdot)$, while μ_C is the equilibrium object determined by the model parameters. Thus, it remains to identify eight parameters.

We have eight corresponding moment restrictions: the entry probabilities in the pre-certification and post-procurement regimes; the testing probability conditional on entry; the certification probability conditional on testing; the relationship between the cost-signal distributions of certified and uncertified firms; and the three mean-quality moments given by the mean of $\widetilde{\xi}_i$ in the pre-certification period and its post-procurement means conditional on certification and non-certification.[16] The equality between the number of parameters and the number of moment restrictions is only a necessary condition for identification. Our identification argument is therefore conditional on the corresponding moment Jacobian having full rank at the true parameter values.

A practical complication is that the type-specific entry and testing probabilities and consumer belief are equilibrium fixed points: expected profits depend on these objects through the distribution of competitors, while the probabilities and belief themselves are determined by expected profits. To implement the identification argument, we therefore use a nested procedure. For each candidate value of the structural parameters, we first recover the type-specific equilibrium probabilities and quality parameters from the observed moments, then compute the corresponding expected profits and model-implied probabilities, and finally choose the parameters so that the recovered and model-implied probabilities coincide.

5.4. **Estimation.** We estimate the model sequentially, following the identification argument above. We first estimate the price coefficient using pooled pre-procurement data and then, separately for each drug, recover composite product qualities, demand-guarantee rates, marginal costs, and scale-economy parameters from the demand, pricing, and bidding conditions. In the final step, we estimate the remaining parameters using a nested minimum-distance procedure. We then obtain standard errors by

[16]The cost distribution restriction is functional, so we summarize it by a single moment.

bootstrapping the full multi-step estimator. Appendix A provides the detailed moment conditions, inversion formulas, equilibrium-profit calculations, and numerical implementation.

6. Estimation Results

This section reports estimates of the demand-side objects: the price coefficient from Step 1 and the demand-guarantee rates from Step 2. Estimation of the remaining supply-side objects — marginal costs, bidding, and the entry and testing primitives — is in progress.

6.1. **Price Sensitivity.**

Table 2 reports estimates of Equation (A.1). In the baseline sample of procured drugs, the price coefficient is -1.418. Demand becomes slightly less elastic when we include never-procured drugs, and the estimates are robust to the inclusion of therapeutic class-province-year fixed effects, which capture disease-specific time trends. The instruments are strong, with first-stage F statistics above 1,200 in all specifications. The implied own-price elasticities average slightly above one in absolute value, consistent with prior estimates of prescription-drug demand in China (Barwick et al., 2025; Cao et al., 2024).

Obtaining BE certification raises a product's mean utility by the equivalent of a 15% price reduction. This result is consistent with the reduced-form evidence in Section 3 that certification carries market value.

6.2. **Demand-Guarantee Enforcement.**

We next apply Step 2 to the post-procurement annual panel. Figure 5 shows the distribution of the estimated demand-guarantee rates. The market-by-market inversion is noisy: non-winner shares are small and measured with error, and the naive inversion can yield values outside the unit interval — a negative λ_m, for example, arises when a non-winner's observed share exceeds the level implied by its estimated quality in a market without restriction. In our sample, 12% of the inverted values fall outside the unit interval. We therefore discipline the estimates with the logistic projection on market observables described in Step 2 and use the fitted values in the remainder of the analysis. The fitted distribution tracks the raw inversions while shrinking the estimates into the range of sensible values.

The estimates imply that demand guarantees are strongly enforced: the average fitted λ_m is 0.85, so in a typical post-procurement market, the large majority of demand is effectively restricted to the procurement winner. Enforcement also varies

TABLE 2. Demand Estimates: Price Coefficient

	(1)	(2)	(3)	(4)
Log OOP price	-1.410	-1.418	-1.361	-1.384
	(0.074)	(0.073)	(0.034)	(0.034)
Passed BE	0.211	0.175	0.273	0.254
	(0.027)	(0.027)	(0.021)	(0.021)
N	167,170	167,168	518,318	518,318
KP rk Wald F	1331.1	1279.1	5617.1	5235.1
Product-province FE	Yes	Yes	Yes	Yes
Province-year FE	Yes		Yes	
Class-province-year FE		Yes		Yes
Sample	Procured	Procured	Full	Full

Notes: Observations are at the drug-firm-province-year level, restricted to pre-procurement years; years in which a drug's insurance tier changes are excluded. The out-of-pocket price is instrumented with insurance-formulary tier dummies. Columns (1)–(2) use procured drugs in waves 2–9; columns (3)–(4) add never-procured drugs. All columns include product-province fixed effects and a BE-approval indicator. Standard errors are clustered at the drug-province level.

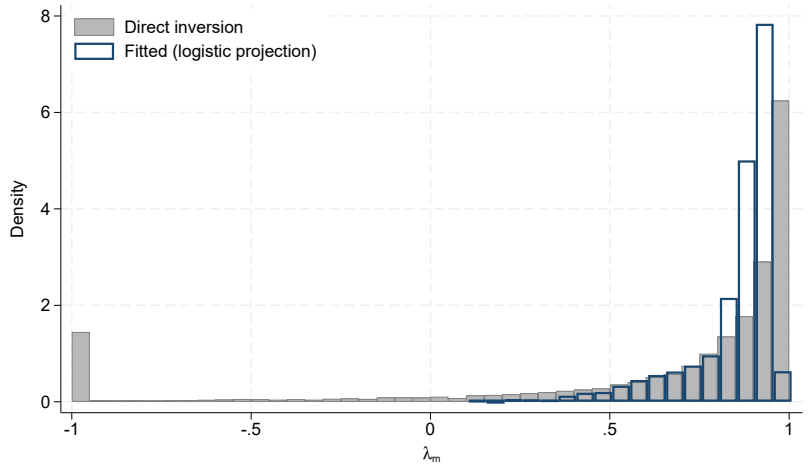
systematically across markets (Appendix Table A2). Most notably, λ_m rises steeply with the official guarantee rate announced at auction: markets with more generous official guarantees exhibit stricter enforcement. At the mean of the fitted values, the estimates imply that a ten-percentage-point increase in the official guarantee rate raises λ_m by roughly 2.6 percentage points. The estimated relationship is stable as auction-round, province, and therapeutic-class controls are added sequentially. Enforcement also strengthens modestly over the contract cycle and varies substantially across provinces and auction rounds, consistent with heterogeneous implementation of procurement quotas.

These estimates matter for the model because λ_m governs the size of the procurement prize: a higher λ_m raises the value of winning and hence the returns to certification and entry. The same estimation recovers the product qualities $\widetilde{\xi}_i$, anchored by winner observations; their distributions conditional on certification status enter the equilibrium restrictions in Step 4.

7. Counterfactuals and Conclusion

We are currently completing the remaining estimation steps — marginal costs, bidding, and the entry and testing primitives — and will report the corresponding counterfactual results in the next version. The counterfactual analysis is organized

FIGURE 5. Estimated Demand-Guarantee Rates



Notes: The gray bars show the distribution of the demand-guarantee rates λ_m obtained by direct inversion of the estimated market wedges; values below -1 are shown at -1, and 12% of the inverted values fall outside the unit interval. The navy outline shows the distribution of fitted values from the logistic projection on market observables. Markets are drug-province-year cells served by a single procurement winner.

around three questions. First, is quality certification alone sufficient to induce entry when firms continue to face barriers to market access? Second, how much of procurement's effect operates through guaranteed market access, as opposed to scale-related cost reductions? Third, how does procurement compare with a direct entry subsidy as a policy for encouraging entry while preserving competition?

The first set of counterfactuals considers markets with BE certification but no centralized procurement. Although the data include a brief period in which certification was available before procurement was implemented, entry during this period may have reflected firms' anticipation of the upcoming procurement program. The observed outcomes therefore need not correspond to the equilibrium that would arise under certification alone. This counterfactual isolates the information role of BE certification by asking whether improving consumers' information about generic quality is sufficient to induce entry when certified firms must still compete for demand against the branded incumbent and other products. We also consider a counterfactual that removes the branded incumbent's residual demand advantage, captured by ξ_B, to quantify the extent to which brand loyalty limits the private returns to certification.

The second set of counterfactuals introduces procurement. We compare certification only, certification with procurement but no scale economies ($\gamma = 0$), and the baseline procurement regime. This decomposition separates two channels. The first is the market-access channel: procurement guarantees demand to certified winners and therefore raises the private return to certification. The second is the scale channel: guaranteed demand lowers marginal costs, which can further increase entry and intensify downstream competition. Comparing these regimes shows whether procurement matters primarily because it shifts demand to certified firms, or because it allows entrants to reach the scale needed to compete at lower cost.

The final set of counterfactuals compares centralized procurement with a direct entry subsidy. An entry subsidy increases participation by lowering firms' fixed cost of entering, but it does not directly reduce informational frictions, provide access to buyers, or require firms to compete for the resulting support. Procurement instead conditions eligibility on quality certification and allocates access to demand through an auction. We compare the two policies under common fiscal-cost and entry-equivalent benchmarks, examining their effects on the number and composition of entrants, certification decisions, procurement and downstream prices, government expenditure, and consumer welfare. This comparison allows us to assess whether reducing barriers to market participation can encourage entry more effectively than subsidizing it directly, while retaining the competitive discipline generated by the procurement auction.

Together, these counterfactuals clarify the distinct roles of information, market access, and scale in shaping entry and competition. They also allow us to evaluate the central policy argument of the paper: procurement may promote entry not by protecting firms from competition, but by reducing barriers to participation and requiring firms to compete for access to a large market.

References

Allende, Claudia, Juan Pablo Atal, Rodrigo Carril, Jose Ignacio Cuesta, and Andres Gonzalez-Lira, "Centralizing Procurement: The Roles of Scale, Selection, and Variety," Technical Report, mimeo 2024.

Anton, James J. and Dennis A. Yao, "Coordination in Split-Award Auctions," *Quarterly Journal of Economics*, 1992, *107* (2), 681–707.

Asker, John and Estelle Cantillon, "Procurement When Price and Quality Matter," *RAND Journal of Economics*, 2010, *41* (1), 1–34.

Atal, Juan Pablo, José Ignacio Cuesta, and Morten Sæthre, "Quality Regulation and Competition: Evidence from Pharmaceutical Markets," Technical Report, National Bureau of Economic Research 2022.

Athey, Susan, Dominic Coey, and Jonathan Levin, "Set-Asides and Subsidies in Auctions," *American Economic Journal: Microeconomics*, 2013, *5* (1), 1–27.

Bajari, Patrick, Robert McMillan, and Steven Tadelis, "Auctions Versus Negotiations in Procurement: An Empirical Analysis," *Journal of Law, Economics, & Organization*, 2008, *25* (2), 372–399.

Barwick, Panle Jia, Ashley T Swanson, and Tianli Xia, "A Double Dose of Reform: Insurance and Centralized Negotiation in Drug Markets," Technical Report, National Bureau of Economic Research 2025.

Brugués, Javier, "The Effects of Public Procurement on Medicine Supply," Technical Report, mimeo 2020.

Cao, Shengmao, Lisa Xuejie Yi, and Chuan Yu, "Competitive Bidding in Drug Procurement: Evidence from China," *American Economic Journal: Economic Policy*, 2024, *16* (3), 481–513.

Chaturvedi, Aadhaar, Damian R Beil, and Victor Martínez de Albéniz, "Split-Award Auctions for Supplier Retention," *Management Science*, 2014, *60* (7), 1719–1737.

Chaudhuri, Shubham, Pinelopi K. Goldberg, and Panle Jia, "Estimating the Effects of Global Patent Protection in Pharmaceuticals: A Case Study of Quinolones in India," *American Economic Review*, 2006, *96* (5), 1477–1514.

Che, Yeon-Koo, "Design Competition Through Multidimensional Auctions," *RAND Journal of Economics*, 1993, *24* (4), 668–680.

Coviello, Decio and Mario Mariniello, "Publicity Requirements in Public Procurement: Evidence from a Regression Discontinuity Design," *Journal of Public Economics*, 2014, *109*, 76–100.

_____, Andrea Guglielmo, and Giancarlo Spagnolo, "The Effect of Discretion on Procurement Performance," *Management Science*, 2018, *64* (2), 715–738.

Cripps, Martin and Norman Ireland, "The Design of Auctions and Tenders with Quality Thresholds: The Symmetric Case," *Economic Journal*, 1994, *104* (423), 316–326.

Danzon, Patricia M., Andrew W. Mulcahy, and Adrian K. Towse, "Pharmaceutical Pricing in Emerging Markets: Effects of Income, Competition, and Procurement," *Health Economics*, 2015, *24* (2), 238–252.

Dranove, David and Ginger Zhe Jin, "Quality Disclosure and Certification: Theory and Practice," *Journal of Economic Literature*, 2010, *48* (4), 935–963.

Dubois, Pierre and Laura Lasio, "Identifying Industry Margins with Price Constraints: Structural Estimation on Pharmaceuticals," *American Economic Review*, 2018, *108* (12), 3685–3724.

_____, **Yassine Lefouili, and Stéphane Straub**, "Pooled Procurement of Drugs in Low and Middle Income Countries," *European Economic Review*, 2021, *132*, 103655.

Elfenbein, Daniel W., Raymond Fisman, and Brian McManus, "Market Structure, Reputation, and the Value of Quality Certification," *American Economic Journal: Microeconomics*, 2015, *7* (4), 83–108.

Estache, Antonio and Atsushi Iimi, "Quality or Price? Evidence from ODA-Financed Public Procurement," *Public Finance Review*, 2012, *40* (4), 435–469.

Gong, Jiong, Jianpei Li, and R Preston McAfee, "Split-Award Contracts with Investment," *Journal of Public Economics*, 2012, *96* (1-2), 188–197.

Huang, Qifan, Zhentong Lu, and Castiel Chen Zhuang, "Physician Agency and the Zero-Markup Drug Policy in China: Evidence From a Structural Model," *Health Economics*, 2026, *35* (4), 653–676.

Ishitani, Catherine, "In Defense of the Middleman: Quality Failures in the Generic Pharmaceutical Market." PhD dissertation, University of Pennsylvania 2025.

Jin, Ginger Zhe and Phillip Leslie, "The Effects of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards," *Quarterly Journal of Economics*, 2003, *118* (2), 409–451.

Kang, Karam and Robert A Miller, "Winning by default: why is there so little competition in government procurement?," *The Review of Economic Studies*, 2022, *89* (3), 1495–1556.

Krasnokutskaya, Elena and Katja Seim, "Bid Preference Programs and Participation in Highway Procurement Auctions," *American Economic Review*, 2011, *101*, 2653–2686.

Leland, Hayne E., "Quacks, Lemons, and Licensing: A Theory of Minimum Quality Standards," *Journal of Political Economy*, 1979, *87* (6), 1328–1346.

Li, Tong and Xiaoyong Zheng, "Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions," *Review of Economic Studies*, 2009, *76* (4), 1397–1429.

Liu, Chenyuan, Yi Lu, and Wanyu Yang, "Excessive Pharmaceutical Marketing

Expenditure: Policy Remedies from China," Technical Report, mimeo 2025.

Manelli, Alejandro M. and Daniel R. Vincent, "Optimal Procurement Mechanisms," *Econometrica*, 1995, *63* (3), 591–620.

Shapiro, Carl, "Investment, Moral Hazard, and Occupational Licensing," *Review of Economic Studies*, 1986, *53* (5), 843–862.

Zhang, Xuan, Nina Yin, Jiayu Zhang, Qingyuan Xue, and Bin Jiang, "Centralized Drug Procurement for Affordable and Sustainable Healthcare: Evidence from China," *Available at SSRN 5632870*, 2025.

Zhang, Yi, "Splitting Award or Winner Takes All?: Evidence from China's National Drug Procurement Auction," Technical Report, mimeo 2023.

Appendix A. Estimation Procedure

A.1. **Step 1: Price elasticity.** We estimate the price coefficient α from pre-procurement data. Aggregating the panel to the product–province–year level and restricting to pre-procurement years, we estimate

$$\log(s_{imt}/s_{0mt}) = \alpha \log(\kappa_{dmt}\, p_{imt}) + \beta\, v_{it} + \nu_{im} + \iota_{mt} + \varepsilon_{imt}, \tag{A.1}$$

where s_{imt} is the market share of product i of drug d in province m and year t, and s_{0mt} is the share of the outside option. We instrument the out-of-pocket price $\kappa_{dmt}p_{imt}$ with the drug's insurance-formulary tier, which changes within drug over time as insurance catalogs are revised, and we exclude the years in which a drug's tier changes. Product-province fixed effects ν_{im} ensure that α is identified only from within-product variation in insurance tiers over time, and province-year fixed effects ι_{mt} absorb market-level shocks.

A.2. **Step 2: Demand guarantee enforcement and product quality.** Define

$$r_{im} \equiv \log\left(\frac{s_{im}}{s_{0m}}\right) - \hat{\alpha}\log(\kappa_m p_{im}).$$

From Equation (20), r_{im} is linear in product quality and the province wedge as follows:

$$r_{im} = \widetilde{\xi}_i + \mathbf{1}\{i \neq w(m)\}z_m + u_{im} + (\alpha - \widehat{\alpha})\log(\kappa_m p_{im}).$$

We can then estimate $\{\widetilde{\xi}_i, z_m\}$ jointly by least squares:

$$\{\widehat{\widetilde{\xi}}_i, \widehat{z}_m\} = \arg\min_{\{\widetilde{\xi}_i, z_m\}} \sum_{i,m} \left[r_{im} - \widetilde{\xi}_i - \mathbf{1}\{i \neq w(m)\}z_m\right]^2. \tag{A.2}$$

We then estimate λ_m as

$$\widehat{\lambda}_m = (\widehat{s}^*_{0m} + \widehat{s}^*_{w(m)m})\,(1 - \exp(\widehat{z}_m))\,, \tag{A.3}$$

where

$$\widehat{s}^*_{0m} = \frac{1}{1 + A_m + \exp(\widehat{z}_m)B_m},\ \widehat{s}^*_{w(m)m} = \frac{A_m}{1 + A_m + \exp(\widehat{z}_m)B_m},$$
$$A_m = \exp\left(\alpha\log(\kappa_m p_{w(m)m}) + \widehat{\widetilde{\xi}}_{w(m)}\right),\ B_m = \sum_{i \neq w(m)} \exp\left(\alpha\log(\kappa_m p_{im}) + \widehat{\widetilde{\xi}}_i\right). \tag{A.4}$$

A.3. **Step 3: Marginal costs.** For a non-winner (i.e., $i \in \mathcal{L}_m$), the pricing equation can be written as

$$s^*_{im} = (1-\lambda_m)\left[1 - \frac{p_{im}}{\alpha(\tilde{c}_i + \delta_m + \epsilon_{im} - p_{im})}\right]. \tag{A.5}$$

The non-winner's observed share is $(1-\lambda_m)$ times its share in the unrestricted market segment. Define the unrestricted-segment share as

$$s^f_{im}\left(\alpha, \{\tilde{\xi}_j\}_{j\in\mathcal{J}_m}\right) = \frac{s^*_{im}}{1-\lambda_m} = \frac{\exp(\alpha\log(\kappa_m p_{im}) + \tilde{\xi}_i)}{1+\sum_k \exp(\alpha\log(\kappa_m p_{km}) + \tilde{\xi}_k)},$$

where $\mathcal{J}_m$ is the set of active products in province m. Substituting this expression into the non-winner pricing equation, Eq. (A.5), gives the marginal cost inversion:

$$p_{im}\left(1 + \frac{1}{\alpha\left[1 - s^f_{im}\left(\alpha, \{\tilde{\xi}_k\}_k\right)\right]}\right) = \tilde{c}_i + \delta_m + \epsilon_{im}.$$

Plugging in the demand estimates from previous steps in the LHS and denoting it as $\hat{c}_{im}$, the above equation can be written as

$$\hat{c}_{im} = \tilde{c}_i + \delta_m + \epsilon_{im} + g_{im},$$

where g_{im} is the generated estimation error:

$$g_{im} \equiv p_{im}\left(\frac{1}{\hat{\alpha}\left[1 - s^f_{im}\left(\hat{\alpha}, \{\hat{\tilde{\xi}}_k\}_k\right)\right]} - \frac{1}{\alpha\left[1 - s^f_{im}\left(\alpha, \{\tilde{\xi}_k\}_k\right)\right]}\right).$$

Under consistency of the demand estimates in the previous steps, g_{im} converges to zero asymptotically. In finite samples, however, $\hat{c}_{im}$ is a nonlinear function of the estimated demand parameters, so the generated estimation error need not be centered at zero. We write

$$g_{im} = b_{im} + \tilde{g}_{im},$$

where b_{im} captures the finite-sample plug-in bias induced by the nonlinear share and markup inversion, and $\tilde{g}_{im}$ is the remaining centered estimation error. We estimate b_{im} by bootstrap and denote it by $\hat{b}_{im}$.

Finally, we recover product-level costs $\tilde{c}_i$ and province effects δ_m by projecting the bias-corrected pseudo-costs on product and province fixed effects:

$$\{\widehat{\tilde{c}}_i, \widehat{\delta}_m\} = \arg\min_{\{\tilde{c}_i\}_i, \{\delta_m\}_m} \sum_m \sum_{i\in\mathcal{L}_m} \left(\widehat{c}_{im} - \hat{b}_{im} - \tilde{c}_i - \delta_m\right)^2, \tag{A.6}$$

subject to the normalization

$$\sum_m \delta_m = 0.$$

We restrict attention to observations for which the implied non-winner pricing FOC is economically meaningful. In particular, since $\alpha < 0$, positive marginal cost requires that $\hat{c}_{im} > 0$, or equivalently,

$$-\alpha(1 - s^f_{im}) > 1.$$

In implementation, we check this condition using the estimated demand parameters and exclude observations that are close to the boundary. Identification of $\{\tilde{c}_i\}_i$ and $\{\delta_m\}_m$ requires the product-province network among non-winner observations to be connected. Products or provinces with only a small number of non-winner observations may be weakly identified.

We then estimate θ_γ by minimizing the difference between the observed bids and the model-predicted bids. We can use an NLS approach to minimize

$$\hat{\theta}_\gamma = \arg\min_{\theta_\gamma} \sum_i \left[b_i - \mathbb{E}\left\{ \beta(\hat{c}_i(\theta_\gamma), n^C, n^U) | \theta_\gamma, \hat{\alpha}, \{\hat{\lambda}_m\}_m, \{\hat{\delta}_m\}_m, \hat{G}_{\tilde{\xi}|v}, \hat{G}_{\tilde{c}|v}) \right\} \right]^2,$$

where $\hat{c}_i(\theta_\gamma)$ is given by

$$\hat{c}_i(\theta_\gamma) = \widehat{\tilde{c}_i} \left(1 + \theta_\gamma \sum_m \text{size}_m \mathbf{1}(w_m = i) \right).$$

A.4. **Step 4: Testing, entry, and type distributions.** Let

$$\zeta_1 \equiv (\phi, \psi)', \qquad \zeta_2 \equiv (\theta'_e, \theta'_\eta)', \qquad \zeta \equiv (\zeta'_1, \zeta'_2)',$$

where both entry cost parameters θ_e and test cost parameters θ_η are two-dimensional. Thus, ζ contains six parameters. For a fixed candidate ζ_1, we first recover the type-specific entry and testing probabilities, together with $(\omega, \bar{\xi})$, from the observed entry, testing, certification, quality, and cost-distribution moments. We then compute profits and choose ζ so that the recovered probabilities satisfy the model's entry and testing optimality conditions.

A.4.1. *Recovering equilibrium probabilities conditional on ζ_1.* For given ζ_1, we recover (i) the equilibrium type-specific entry probabilities under each regime, q^0_τ and q_τ, (ii) the equilibrium type-specific testing probabilities, ρ_τ, and (iii) two parameters, $(\omega, \bar{\xi})$, using the following eight moment conditions. The first set of moments are associated with entry probabilities for each regime, the testing probability, and the certification

probability.

$$\Pr(\text{Entry}^0) = \phi q_H^0 + (1-\phi) q_L^0, \tag{A.7}$$

$$\Pr(\text{Entry}) = \phi q_H + (1-\phi) q_L, \tag{A.8}$$

$$\Pr(\text{Tested} \mid \text{Entry}) = \frac{\phi \rho_H q_H + (1-\phi)\rho_L q_L}{\phi q_H + (1-\phi) q_L}, \tag{A.9}$$

$$\Pr(\text{Certified} \mid \text{Tested}) = \frac{\phi \rho_H q_H + \psi(1-\phi)\rho_L q_L}{\phi \rho_H q_H + (1-\phi)\rho_L q_L}. \tag{A.10}$$

Note that we observe the sample analogue of the left-hand side of the above equations: the entry probabilities under each regime, $\Pr(\text{Entry}^0)$ and $\Pr(\text{Entry})$, the testing probability conditional on entry, $\Pr(\text{Tested} \mid \text{Entry})$ and the certification probability conditional on testing, $\Pr(\text{Certified} \mid \text{Tested})$.

The next three moments are associated with the average unobserved quality in demand.

$$\mathbb{E}[\widetilde{\xi}_i^0] = \frac{\omega \phi q_H^0}{\phi q_H^0 + (1-\phi) q_L^0} + \bar{\xi}, \tag{A.11}$$

$$\mathbb{E}[\widetilde{\xi}_i \mid C] = \frac{\omega \rho_H \phi q_H}{\rho_H \phi q_H + \psi \rho_L (1-\phi) q_L} + \bar{\xi}, \tag{A.12}$$

$$\mathbb{E}[\widetilde{\xi}_i \mid U] = \frac{\omega (1-\rho_H) \phi q_H}{(1-\rho_H)\phi q_H + (1-\psi\rho_L)(1-\phi) q_L} + \bar{\xi}, \tag{A.13}$$

Another moment concerns the relationship between the cost signal distributions conditional on certification or not.

$$\frac{G_{c|C}(c;\widehat{\theta}_\gamma) - F_{c|L}(c;\widehat{\theta}_\gamma)}{G_{c|U}(c;\widehat{\theta}_\gamma) - F_{c|L}(c;\widehat{\theta}_\gamma)} = \frac{\mu_C}{\mu_U}, \tag{A.14}$$

where

$$\mu_C = \frac{\rho_H \phi q_H}{\rho_H \phi q_H + \psi \rho_L (1-\phi) q_L}, \quad \mu_U = \frac{(1-\rho_H)\phi q_H}{(1-\rho_H)\phi q_H + (1-\psi\rho_L)(1-\phi) q_L}.$$

Note that Equation (A.14) holds for any c values such that the denominator is nonzero. Rather than imposing Equation (A.14) at a single CDF evaluation point, we recover the common ratio, μ_C/μ_U using an integrated CDF-distance criterion.

$$\frac{\mu_C}{\mu_U} = \arg\min_r \int_c \left[\left(G_{c|C}(c;\widehat{\theta}_\gamma) - F_{c|L}(c;\widehat{\theta}_\gamma) \right) - r \left(G_{c|U}(c;\widehat{\theta}_\gamma) - F_{c|L}(c;\widehat{\theta}_\gamma) \right) \right]^2 dG_c(c).$$

It can be shown that, conditional on ζ_1, the observed moments recover

$$\mathcal{P}(\zeta_1) \equiv \left(q_H^0(\zeta_1), q_L^0(\zeta_1), q_H(\zeta_1), q_L(\zeta_1), \rho_H(\zeta_1), \rho_L(\zeta_1) \right),$$

as well as $\omega(\zeta_1)$ and $\bar{\xi}(\zeta_1)$. Then we can recover the cost signal distribution for the high-type firm using Equation (22).

This step requires the corresponding aggregate moments to be interior, so that certified firms, uncertified firms, and pre-policy entrants are all observed. We restrict the parameter space to candidates satisfying

$$\mu_C, \mu_U, \phi \in (0,1), \quad \omega > 0, \quad 0 \leq \psi < \Pr(\text{Certified} \mid \text{Tested}),$$

and for each $\tau \in \{L, H\}$,

$$q_\tau^0, q_\tau, \rho_\tau \in [0,1].$$

The restriction $\omega > 0$ fixes the type labeling so that the high type has higher mean demand quality. We also require $\mu_C \neq \mu_U$, and that Equation (22) is a proper CDF: it must be weakly increasing, lie between zero and one, and satisfy the appropriate boundary conditions. Candidates of ζ_1 violating any of these restrictions are assigned an infinite objective value.

A.4.2. *Expected profits.* For each candidate ζ_1, we use the demand and cost primitives from Steps 1–4, together with $F_{c|H}(\cdot;\zeta_1)$, to compute profits conditional on each realized market configuration. Let

$$\Pi_{C,\tau}(n^C, n^U; \zeta_1) \qquad \text{and} \qquad \Pi_{U,\tau}(n^C, n^U; \zeta_1)$$

denote the expected operating profits of a type-τ firm when it is certified or uncertified, respectively, and when the market contains n^C certified and n^U uncertified firms in total. These conditional profit objects incorporate the equilibrium pricing and bidding calculations for a fixed market configuration. They integrate over the focal firm's type-specific cost signal, the status-specific demand-quality distribution, and the corresponding cost and quality distributions of rival firms. The pricing and bidding simulations use the estimated distributions $G_{\widetilde{\xi}|v}$ directly, so ω and $\bar{\xi}$ enter the moment inversion above but do not enter the continuation-profit calculation separately.

Let J denote the number of potential firms for the drug, and let $n^{C,-i}$ and $n^{U,-i}$ denote the numbers of certified and uncertified rivals of a focal firm. The candidate probabilities imply that a potential rival is certified, uncertified, or does not enter

with probabilities

$$\begin{aligned}
p_C(\zeta_1) &= \phi q_H(\zeta_1)\rho_H(\zeta_1) + \psi(1-\phi)q_L(\zeta_1)\rho_L(\zeta_1),\\
p_U(\zeta_1) &= \phi q_H(\zeta_1)[1-\rho_H(\zeta_1)] + (1-\phi)q_L(\zeta_1)[1-\psi\rho_L(\zeta_1)],\\
p_0(\zeta_1) &= 1 - p_C(\zeta_1) - p_U(\zeta_1).
\end{aligned}$$

Under the model's independent type and cost draws across firms, $(n^{C,-i}, n^{U,-i}, J-1-n^{C,-i}-n^{U,-i})$ follows the multinomial distribution implied by (p_C, p_U, p_0). Its probability mass function is

$$\begin{aligned}
P_{-i}(n^C, n^U; \zeta_1) = \frac{(J-1)!}{n^C! n^U! (J-1-n^C-n^U)!} [p_C(\zeta_1)]^{n^C} [p_U(\zeta_1)]^{n^U} \\
\times [p_0(\zeta_1)]^{J-1-n^C-n^U}.
\end{aligned} \tag{A.15}$$

The expected continuation profits are then

$$\pi_{C,\tau}(\zeta_1) = \sum_{n^C=0}^{J-1} \sum_{n^U=0}^{J-1-n^C} \Pi_{C,\tau}(n^C+1, n^U; \zeta_1) P_{-i}(n^C, n^U; \zeta_1), \tag{A.16}$$

$$\pi_{U,\tau}(\zeta_1) = \sum_{n^C=0}^{J-1} \sum_{n^U=0}^{J-1-n^C} \Pi_{U,\tau}(n^C, n^U+1; \zeta_1) P_{-i}(n^C, n^U; \zeta_1). \tag{A.17}$$

The arguments (n^C+1, n^U) and (n^C, n^U+1) make explicit that the conditional profit functions count the focal firm, whereas (n^C, n^U) in P_{-i} counts only rivals.

We compute the expected operating profit $\pi^0_\tau(\zeta_1)$ in the pre-auction regime analogously, integrating the corresponding conditional profit over the candidate distribution of the number and types of rival entrants.

Define the gross gain from testing by

$$\Delta_H(\zeta_1) = \pi_{C,H}(\zeta_1) - \pi_{U,H}(\zeta_1), \qquad \Delta_L(\zeta_1) = \psi\left[\pi_{C,L}(\zeta_1) - \pi_{U,L}(\zeta_1)\right]. \tag{A.18}$$

The factor ψ in Δ_L accounts for the fact that a low-type firm receives the certified continuation payoff only when it passes the test. Given that the testing cost is realized after entry, the value of entry in the post-policy regime is

$$V_\tau(\zeta) = \pi_{U,\tau}(\zeta_1) + \int \max\{\Delta_\tau(\zeta_1) - \eta, 0\}\, dF_\eta(\eta; \theta_\eta), \qquad \tau \in \{L, H\}. \tag{A.19}$$

Thus, the entry value includes the option value of testing.

A.4.3. *Equilibrium restrictions and estimation.* The model-predicted entry and testing probabilities are

$$q_\tau^{0,M}(\zeta) = F_e\big(\pi_\tau^0(\zeta_1);\theta_e\big)\,, \tag{A.20}$$

$$q_\tau^{M}(\zeta) = F_e(V_\tau(\zeta);\theta_e)\,, \tag{A.21}$$

$$\rho_H^{M}(\zeta) = F_\eta(\Delta_H(\zeta_1);\theta_\eta)\,, \tag{A.22}$$

$$\rho_L^{M}(\zeta) = F_\eta(\Delta_L(\zeta_1);\theta_\eta)\,. \tag{A.23}$$

We compare these probabilities with the probabilities recovered from the observed moments. Define the six-dimensional equilibrium residual

$$h(\zeta) \equiv \begin{pmatrix} q_H^0(\zeta_1) - q_H^{0,M}(\zeta) \\ q_L^0(\zeta_1) - q_L^{0,M}(\zeta) \\ q_H(\zeta_1) - q_H^{M}(\zeta) \\ q_L(\zeta_1) - q_L^{M}(\zeta) \\ \rho_H(\zeta_1) - \rho_H^{M}(\zeta) \\ \rho_L(\zeta_1) - \rho_L^{M}(\zeta) \end{pmatrix}. \tag{A.24}$$

The final-stage estimator is

$$\widehat{\zeta} = \arg\min_{\zeta\in\mathcal{Z}} h(\zeta)' W h(\zeta), \tag{A.25}$$

where W is positive definite and $\mathcal{Z}$ incorporates the feasibility restrictions above, including positivity of the scale parameters in θ_e and θ_η. Because $h(\zeta)$ contains six residuals and ζ contains six parameters, the equilibrium system is exactly identified dimensionally. If a feasible exact root exists, Equation (A.25) is equivalent to solving

$$h(\zeta) = 0,$$

and the choice of W does not affect the solution. When simulation error or sampling variation prevents an exact root, we use the minimum-distance formulation and hold the simulation draws fixed across candidate parameter values.

Equality between the number of equilibrium restrictions and parameters does not by itself establish identification. We therefore examine the singular values of the numerical Jacobian, profile the objective with respect to ϕ and ψ, and use multiple starting values to assess weak identification and multiplicity of feasible roots. Because this step uses generated distributions and simulated profits, we conduct inference by bootstrapping the full drug-level estimation procedure, including Steps 1–3.

Appendix B. Procurement Effects on Prices and Quantities

This appendix estimates the effects of procurement on prices and quantities using a stacked event-study specification at the drug (d)-firm (i)-province (m)-quarter (t) level. Let c index procurement cohorts, defined by the quarter of procurement implementation t_c. For each cohort, we stack treated drug-firm-province cells with a set of control drugs that are never procured during the sample period. The plotted event studies use waves 2 through 8, so that each cohort has sufficient post-procurement observations. We estimate:

$$Y_{dimtc} = \sum_{k=-16}^{12} \beta_k \text{Treat}_{djmc} \mathbf{1}\{t - t_c = k\} + \lambda_{dimc} + \lambda_{tc} + \epsilon_{dimtc}, \tag{A.26}$$

where we normalize β_{-1} to zero and include endpoint bins for event times below -16 and above 12. Product-market-cohort fixed effects, λ_{dimc}, absorb permanent differences across drug-firm-province cells within each stack. Cohort-by-quarter fixed effects, λ_{tc}, absorb quarterly shocks common to all observations in a stack. The outcome Y_{dimtc} is either the log price or log quantity.

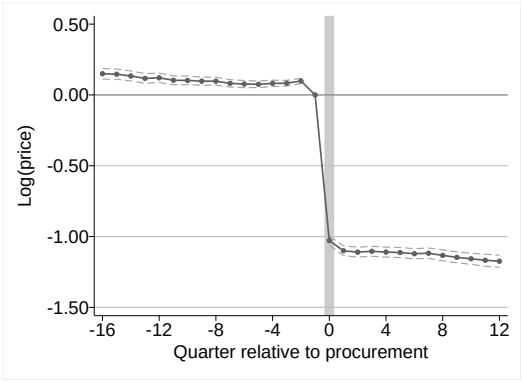
We estimate separate event studies for the three product-market positions created by procurement: winners in their selected provinces, where they are guaranteed market access; winners in other provinces, where another winner is guaranteed market access; and non-winners, which include both losing bidders and non-participants.

The key identifying assumption is that, absent procurement, treated drug-firm-province cells would have followed the same path as never-treated control cells within the same cohort stack, after accounting for permanent product-market differences and common cohort-quarter shocks. This rules out unobserved shocks that arrive exactly at a cohort's procurement date and differentially affect procured drugs relative to comparable never-treated drugs. It also requires that never-treated drugs are not substantially affected by spillovers from procurement.

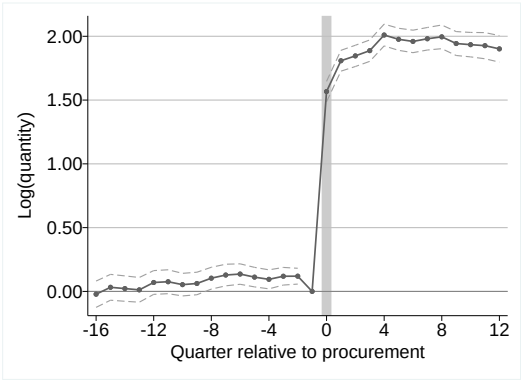
Figure A1 reports the estimates. Panels A and B show winners in their selected provinces. Prices fall by about one log point immediately after procurement and remain low. Quantities rise by roughly 1.5 log points at implementation and stay close to 2 log points above the pre-procurement level. These patterns capture the direct procurement effect: the auction lowers transaction prices, while province selection gives winners immediate access to demand.

Panels C and D show winners in provinces they did not select. Quantities fall sharply after procurement. This decline is partly mechanical: these provinces are

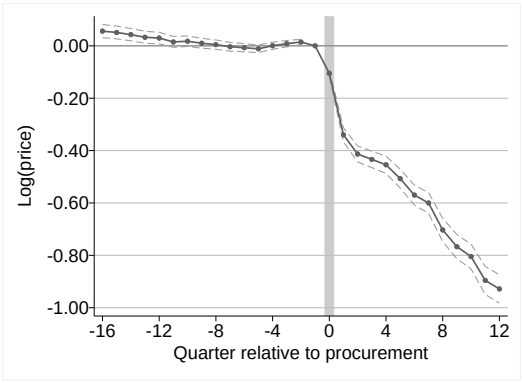
FIGURE A1. Procurement Effects on Prices and Quantities



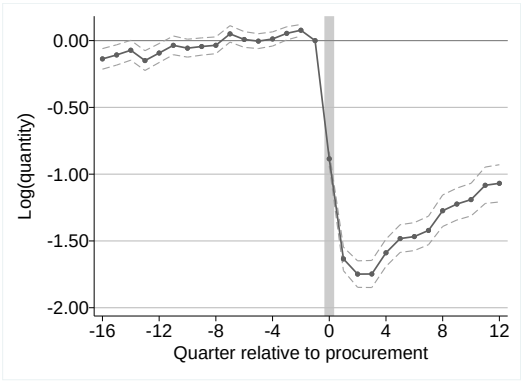
(A) Winners, Selected Provinces



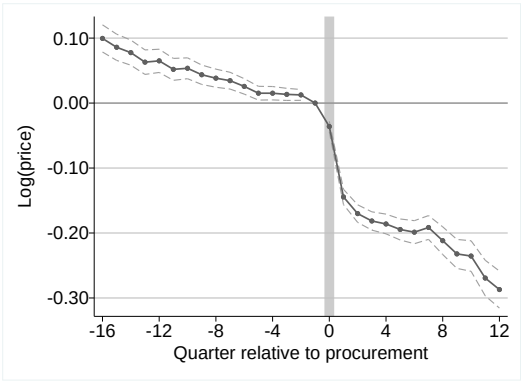
(B) Winners, Selected Provinces



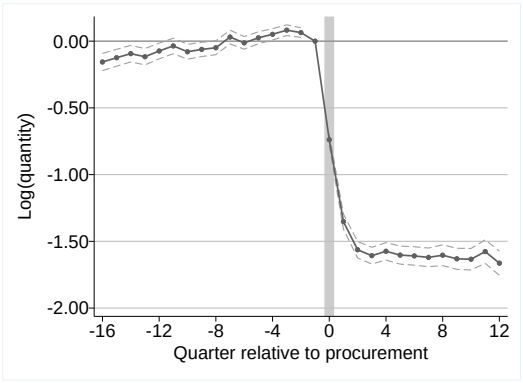
(C) Winners, Other Provinces



(D) Winners, Other Provinces



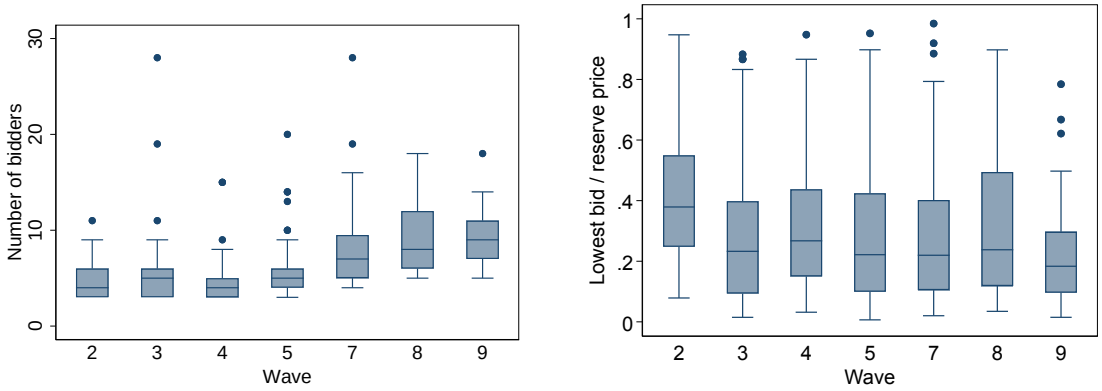
(E) Losers & Non Bidders



(F) Losers & Non Bidders

Notes: Event time is measured in quarters relative to procurement implementation. The coefficient for event quarter −1 is normalized to zero. Controls are never-treated products from non-procured drugs with generic competition.

FIGURE A2. Auction Participation and Outcomes By Wave



(A) Number of Bidders (B) Lowest Bid Relative to Reserve Price

Notes: Panel A shows the number of bidders by procurement round. Panel B shows the lowest bid relative to the reserve price by procurement round.

selected by other auction winners, redirecting contracted demand away from a given winner in markets it did not choose. Prices fall more gradually, reaching nearly one log point below the pre-procurement level by the end of the event window. This gradual decline is consistent with more intensive competition following the sharp price reductions by selected winners. It is also consistent with cost reductions or scale economies from winning procurement.

Panels E and F show non-winners, including losing bidders and non-participants. Their prices decline more moderately, by roughly 0.15–0.30 log points after procurement, while quantities fall by more than one log point and remain depressed. The residual market therefore contracts for firms that do not obtain procurement access.

Overall, the estimates show that centralized procurement lowers prices, reallocates demand toward selected winners, and generates spillovers in residual markets. The comparison between Panels C and E suggests that, even in provinces they did not select, winners experience larger price declines than non-winners. This pattern may reflect scale economies: by gaining procurement access in selected provinces, winners expand output, improve efficiency, and lower marginal costs across markets.

APPENDIX C. ADDITIONAL FIGURES AND TABLES

TABLE A1. CVBP Procurement Rounds: Scope and Timeline

Round	Molecules	Announcement	Bidding	Implementation
Round 1	25	Sep 2019	Sep 2019	Dec 2019
Round 2	33	Dec 2019	Jan 2020	Apr 2020
Round 3	56	Jul 2020	Aug 2020	Nov 2020
Round 4	45	Jan 2021	Feb 2021	May 2021
Round 5	62	Jun 2021	Jun 2021	Nov 2021
Round 6	6	Sep 2021	Nov 2021	Dec 2021
Round 7	61	Jun 2022	Jul 2022	Nov 2022
Round 8	40	Mar 2023	Mar 2023	Jul 2023
Round 9	42	Aug 2023	Nov 2023	Mar 2024
Round 10	62	Nov 2024	Dec 2024	Apr 2025
Round 11	55	Jul 2025	Oct 2025	Feb 2026

Notes: All rounds are national and cover all provinces. The announcement date is the date of the official procurement guideline document released by the National Healthcare Security Administration (NHSA). The bidding date is the date of the sealed-bid auction. The implementation date is the date from which hospitals were required to fulfill procurement quotas. Molecule counts are defined at the molecule–dosage-form level and are taken from official bid documents for Rounds 1–9 and from NHSA public announcements for Rounds 10–11; bid data are not available in our dataset for Rounds 10 and 11. Round 1 is the national expansion of a prior pilot conducted in 11 major cities (the "4+7 pilot"), which had bidding in December 2018 and implementation in March 2019, covering the same 25 molecules. Round 6 is a special insulin round covering 6 insulin product categories.

TABLE A2. Demand-Guarantee Rates: Logistic Projection on Market Observables

	(1)	(2)	(3)	(4)
Official guarantee rate	2.437	2.729	2.843	2.049
	(0.272)	(0.287)	(0.279)	(0.382)
Contract tenure (years)	0.122	0.002	0.044	0.045
	(0.027)	(0.027)	(0.026)	(0.026)
Two winners	0.798	0.752	0.789	0.569
	(0.082)	(0.095)	(0.092)	(0.111)
Three winners	0.695	0.435	0.459	0.393
	(0.067)	(0.070)	(0.068)	(0.073)
Auction-round dummies		Yes	Yes	Yes
Province dummies			Yes	Yes
Therapeutic-class dummies				Yes
Dosage-form dummies				Yes
N (identified market cells)	12,494	12,494	12,494	12,494
R^2	0.833	0.844	0.855	0.857

Notes: Nonlinear least-squares estimates of the logistic specification for λ_m, fit to the market wedges recovered from Equation (21). Columns add covariate blocks sequentially; column (4) is the full specification used to construct the fitted values. Markets are drug-province-year cells served by a single procurement winner. Winner-count categories are relative to auctions with four or more winners.