QUALITY REGULATION AND MARKET ACCESS: 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 Please do not cite.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. In markets with imperfectly observable quality, certification is often used to inform consumers, yet entry may remain limited. We study whether procurement design can overcome this limitation in the context of China's centralized drug procurement program. The program combines a bioequivalence (BE) requirement that certifies drug quality with an auction mechanism that guarantees immediate market access to winners, potentially leveling the playing field for entry and price competition. We document that 23% of auction winners are new entrants, alongside substantial price reductions. Demand-side evidence further shows that consumers value BE certification, confirming that it conveys meaningful information about quality while remaining an imperfect signal. Motivated by these findings, we develop and estimate a model of firm entry, procurement auctions, and post-auction competition. The model quantifies how certification and guaranteed market access jointly shape entry and competition, and shows how procurement design can transform quality regulation from a barrier to entry into a mechanism that promotes competitive entry and lowers prices.

1. Introduction

Many markets are characterized by imperfectly observable quality. When consumers cannot easily verify quality, entrants face a fundamental information problem: even high-quality firms may struggle to convince consumers of their quality, while established firms benefit from reputation and brand recognition. Pharmaceutical markets provide a prominent example, where high-quality generic entrants may struggle to attract consumers despite offering therapeutically equivalent products,

Date: June 2, 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.

allowing branded products to retain significant market power even after patent expiration. Similar challenges can arise in markets for food safety, professional services, education, and digital platforms.

While certification is a common policy response to information frictions, its effectiveness ultimately depends on whether certified firms can obtain access to customers. If market-access barriers remain, certification may improve information without generating meaningful entry or competition. For example, incumbents may possess established distribution networks, long-standing buyer relationships, or other advantages that make it difficult for entrants to reach consumers. As a result, markets may remain concentrated despite improvements in information. This observation raises a broader question: how can policy address both information frictions and market-access frictions to promote competition?

We study this question in China's pharmaceutical market. Prior to recent reforms, many off-patent drugs remained expensive despite the presence of many domestic generic manufacturers. Two frictions contributed to this outcome. First, widespread public skepticism about the safety and efficacy of domestic generics, fueled by recurring quality scandals and weak regulatory oversight, limited demand for generic products. Second, drug procurement was decentralized, with individual hospitals and regional alliances making purchasing decisions. This system favored established firms with existing reputations and hospital relationships, creating substantial barriers to market access for entrants.

In response, the Chinese government introduced two closely linked reforms. First, beginning in 2016, generic manufacturers could obtain bioequivalence (BE) certification, providing a credible signal of product quality. Second, beginning with a pilot in 2018 and expanding nationally in 2019, the government introduced centralized volume-based procurement (CVBP), which allocates large purchase contracts through competitive auctions and guarantees market access to winning firms. We use this setting to study how certification and procurement interact to shape entry and competition.

Our analysis combines several unique data sources covering the period from 2012 to 2026. We assemble a panel of prices and quantities for off-patent drugs across Chinese provinces, detailed records of BE certification applications and outcomes, and procurement-auction data containing bids, winners, and province assignments. The institutional setting is particularly useful because the reforms were introduced

sequentially. Markets first operated without certification or procurement, then transitioned to a certification regime, and finally to a regime that combines certification with centralized procurement. This variation allows us to study how firms and markets respond as information and market-access policies are introduced.

We first document several facts from the data that help distinguish the roles of information and market access. First, BE-certified generic products are able to charge higher prices, indicating that consumers value certification, but without much quantity expansion. The combination of positive price effects and limited quantity responses suggests that certification alleviates information frictions but leaves substantial market-access frictions intact. Second, procurement is associated with substantial entry by certified firms. Nearly one-quarter of procurement winners enter the market only around procurement implementation, and entrants are disproportionately likely to hold BE certification. These patterns indicate that procurement increases the return to certification by providing market access. Third, we show that procurement generates large price reductions and substantial reallocation of demand toward winning firms. Moreover, winning firms experience larger price declines even in markets where they do not receive guaranteed demand, consistent with procurement-induced scale expansion lowering costs more broadly. Taken together, the evidence suggests that certification and procurement operate through distinct but complementary channels: certification reduces information frictions, procurement relaxes market-access constraints, and procurement-induced scale expansion may further intensify competition.

To interpret these patterns, we develop a model that incorporates two frictions and a potential scale channel. The first friction is informational: consumers cannot directly observe the quality of generic drugs and therefore rely on certification when forming beliefs. The second friction concerns market access: certification may improve perceived quality without necessarily generating substantial demand for entrants. BE certification addresses the first friction by providing an informative but imperfect signal of quality, while procurement addresses the second by linking certification to guaranteed demand through centralized purchasing contracts. The model also allows procurement to generate scale-related cost reductions. Because procurement winners receive large guaranteed purchase volumes, they may lower marginal costs through economies of scale, further strengthening their ability to compete in downstream markets. The framework therefore captures three channels through which policy can affect competition: information, market access, and scale.

The model links firms' entry, certification, bidding, and pricing decisions in a unified equilibrium framework. Potential entrants differ in product quality, which affects production costs, certification outcomes, and the costs and benefits of entry. They decide whether to enter the market and whether to seek BE certification. Certification is informative but imperfect: high-quality firms are more likely to obtain certification, yet low-quality firms may also pass. Consumers observe certification status but not underlying quality and form beliefs about certified and uncertified products in equilibrium. Certified generic firms can participate in procurement auctions, where winning determines access to guaranteed demand and may affect production costs through scale expansion. Firms then compete in prices in downstream product markets. Because entry, certification, consumer beliefs, procurement participation, and market competition are jointly determined, the model highlights a central interaction: certification affects demand by changing consumer beliefs, while procurement changes the value of certification by linking it to market access and scale.

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. We first estimate demand and recover the perceived-quality components associated with branded, certified, and uncertified products. We then recover marginal costs from pricing behavior, where costs are identified conditional on observable product status—branded, certified generic, or uncertified generic—rather than on unobserved quality type. We also use variation in procurement assignment to identify procurement-related scale parameters. Together, these steps recover the key demand and cost objects governing downstream competition.

The final step exploits firms' entry, testing, and certification decisions 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 estimate the model using an estimator that follows the identification strategy and use the estimated model to evaluate policies that address information and market-access frictions. We first ask whether quality certification alone is sufficient to induce high-quality generic entry when procurement is unavailable. We then quantify the extent to which procurement promotes competition by expanding market access and by generating scale economies. These exercises allow us to evaluate whether certification alone can promote competition or whether policies that directly expand market access are also necessary. Finally, we evaluate alternative procurement designs, including changes in demand guarantees and market assignments, to understand how procurement can encourage entry while preserving incentives for quality investment. Together, the counterfactual analyses clarify the distinct roles of information and market access in shaping entry and competition and show how procurement can transform quality regulation from a barrier to entry into a force that promotes competition.

More broadly, our findings suggest that quality regulation and competition policy should not be viewed as separate policy domains. While certification is often designed to improve information and product quality, its competitive effects depend critically on the market institutions that govern access to demand. Evaluating these policies jointly can therefore yield very different conclusions than studying either policy in isolation.

2. 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 marke. 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.

3. 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.

3.1. **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 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.

[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.

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.

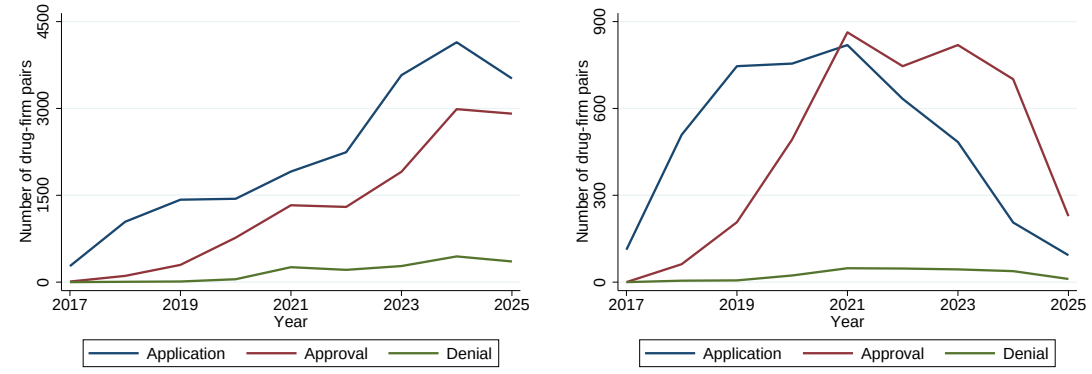
3.2. **Data Description.** We obtain quarterly drug price and sales data at the product-by-province level from MENET, a leading pharmaceutical market research company in China. The data covers 25 provinces from the first quarter of 2012 to the second quarter of 2026. Each product is a stock-keeping unit (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. Should

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; from 2019 onwards, the formulary was unified at the national level. We additionally obtain province-level average copay rates from **?**. Finally, to construct the distribution of consumer incomes, we combine household-level income data from China's 2005

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.

population census with an annual panel of province-level per-capita disposable income from the National Bureau of Statistics. We use the census microdata to estimate the within-province income distribution and scale it forward over time using the annual income panel.

3.3. **Key Data Patterns.** We next document several facts that guide the construction of our model. The evidence shows that BE certification is costly and selective, that procurement auctions provide valuable market access and generate intense competition, and that firms respond strategically to the interaction of these two institutions. These patterns motivate a model in which firms jointly decide whether to pursue certification, participate in procurement, and compete in the product market.

3.3.1. *Certification Pipeline.* 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.

TABLE 1. Auction Summary Statistics

Statistic	N	Mean	P10	P25	P50	P75	P90
Number of eligible firms	333	7.0	3	4	6	9	12
Number of bidders	333	6.7	3	4	6	8	12
Number of winners	333	4.6	2	3	4	6	8
Largest price drop w.r.t. reserve price (%)	333	69.0	28.6	56.2	76.1	88.3	93.3
Median price drop w.r.t. reserve price (%)	333	58.3	13.7	48.8	63.7	76.0	86.5
Auction has a branded bidder	333	69.1%					
Auction has a branded winner	333	7.8%					

Notes: The table reports auction-level statistics. Price-drop rows are measured relative to the reserve price.

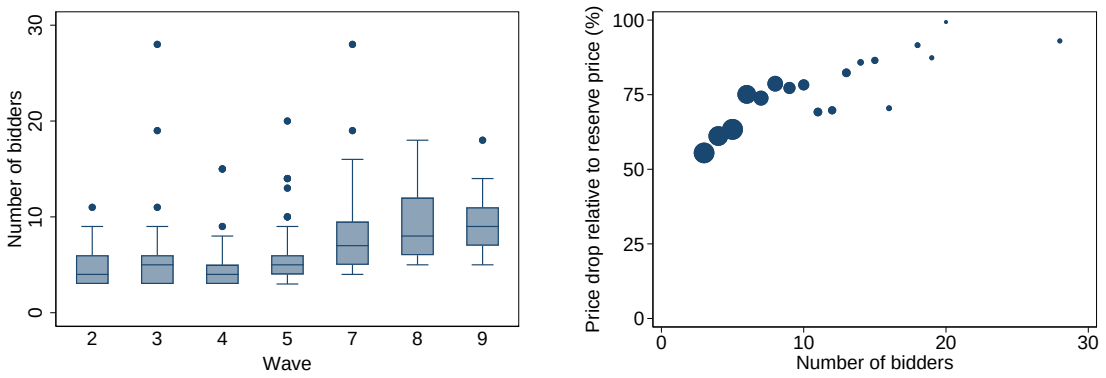
At the same time, denials are not uncommon: in the procurement sample, denied applications account for 5.1% of resolved applications over 2018–2025.

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.

3.3.2. *Auction Access and Bidder Competition.* Table 1 summarizes the participation and bidding patterns in the centralized procurement auctions. We focus on rounds 2 to 9, for which we observe the complete set of bids, except for round 6, which covers only insulin products for diabetes treatment and follows a different set of rules. Among the auctions in our analysis sample, the median auction has six eligible firms, six bidders, and four winners. Essentially, all auctions have at least three bidders and two winners. Thus, the program typically combines a certification screen with a split-award auction, rather than using a single-winner procurement rule. Most eligible firms participate, likely because guaranteed market access is substantial and the residual market is limited.

Panel A of Figure 2 shows the distribution of bidder counts across procurement waves. 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.

FIGURE 2. Auction Participation and Outcomes



(A) Bidders by Wave (B) Price Drop Relative to Reserve Price

Notes: Panel A plots box plots of the number of bidders by procurement round. Panel B groups auctions by bidder count and plots the average percentage reduction of the lowest bid relative to the reserve price.

Auctions generate large price reductions, with meaningful dispersion across auctions. The median largest price drop is 76.1% relative to the reserve price, while the 25th and 75th percentiles are 56.2% and 88.3%. Almost all auctions have at least one bidder bidding below the reserve price. Branded firms participate in many auctions but rarely win: 69.1% of auctions have at least one branded bidder, while only 7.8% have a branded winner.

Panel B of Figure 2 groups auctions by the number of bidders and plots the largest price drop relative to the reserve price. Auctions with more bidders tend to have larger price drops. This relationship is descriptive rather than causal, since auctions with more bidders may differ in underlying market conditions. Nevertheless, the figure shows that realized procurement prices are far below reserve prices across much of the bidder-count distribution.

4. 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 by certified products around procurement implementation, consistent with guaranteed market access increasing the returns to

certification and entry. Finally, we estimate the effects of procurement on prices and quantities. 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.

4.1. **Certification Has Demand Value but Limited Standalone Expansion.** Our first result examines the market value of BE certification. If certification alone conveys meaningful information and generates a large quantity response, then the procurement-access channel would be less important, since certification would already help generic entrants overcome market-access barriers. If certification raises perceived quality but does not substantially increase quantities, then procurement can play an important role by creating demand for certified generic entrants.

To study the market value of BE certification, we estimate how BE approval affects market performance 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.

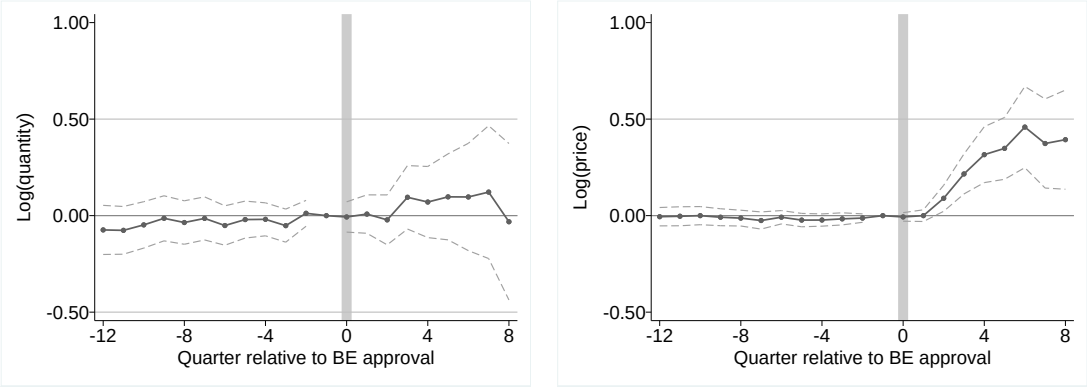
We first estimate a stacked event-study specification. Let c index BE approval cohorts, defined by the quarter in which a generic product obtains BE approval. For each cohort, we stack drug-firm pairs first approved in quarter c with the control products described above, and define event time as $t - c$. 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. The specification includes drug-firm-cohort and drug-quarter-cohort fixed effects, so the comparison is within product and within drug-quarter market. Standard errors are clustered at the drug level.

The key identifying assumption is that, absent BE approval, treated products would have followed the same path as control products within the same drug-quarter market, after removing permanent product differences and common drug-quarter shocks. This assumption rules out unobserved shocks that arrive at the time of BE

FIGURE 3. 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.

approval and differentially affect treated products relative to never-approved products in the same market.

Figure 3 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 close to zero around the approval quarter and do not display a persistent upward shift after approval. This suggests that certification by itself does not immediately solve the market-access problem faced by generic firms.

Panel B shows a different pattern for prices. The price response is small in the approval quarter, but rises gradually over the following quarters and reaches economically meaningful magnitudes. This dynamic pattern is consistent with BE certification changing perceived product quality or consumers' willingness to pay.

Table 2 reports pooled difference-in-differences estimates of the effect of BE approval on log quantity, log price, and log revenue. The pooled estimates reinforce the event-study evidence. BE approval has no detectable effect on log quantity, while prices rise modestly and sales increase through the price margin. On average, generic products raise prices by 9.6% after BE approval, but their prices remain well below those of the corresponding branded products. Before procurement, the median BE generic price is about 54.2% of the corresponding branded product's price.

TABLE 2. BE Approval and Product Outcomes

	Log(Quantity)	Log(Price)	Log(Revenue)
BE Approval	0.019	0.096	0.115
	(0.038)	(0.025)	(0.045)
N	70,493	70,493	70,493
R^2	0.877	0.984	0.919
Drug-Firm FE	Yes	Yes	Yes
Drug-Quarter FE	Yes	Yes	Yes

Notes: Observations are drug-firm-quarter pairs from the full procured-drug sample. All specifications include drug-firm and drug-quarter fixed effects.

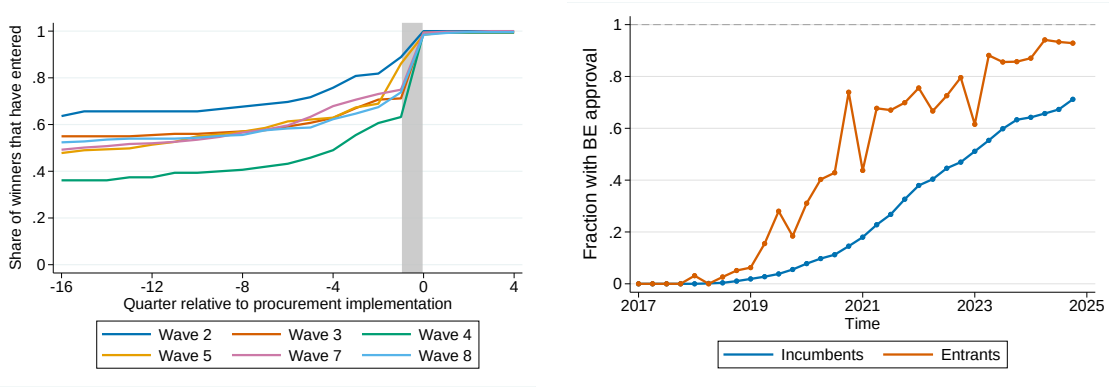
Together, Figure 3 and Table 2 indicate that BE certification has demand value, but that this value is expressed mainly through higher prices rather than through a large standalone increase in quantities in decentralized hospital markets.

4.2. **Procurement Is Associated with New Entry.** Our second result is that guaranteed market access created by procurement is associated with substantial entry by certified firms. Panel A of Figure 4 shows the cumulative share of winners that have entered by quarter relative to procurement implementation, separately by procurement wave. The share of winners rises sharply at implementation: among drug-firm pairs in waves 2 through 9, 24.3% of winners first appear at or after procurement implementation and do not sell the product beforehand. This pattern holds across waves, indicating that procurement access is an important driver for entry.

Panel B shows that entrants are much more likely than incumbents to have BE approval, especially after 2020, when procurement become central to the market. Thus, the entrants appearing around procurement are not simply additional low-certification suppliers; many enter in a BE-ready state. Taken together, the patterns suggest that certification and procurement access are complements: BE approval helps firms satisfy eligibility and quality requirements, while the procurement award provides the market access that makes entry attractive.

4.3. **Auctions Lower Prices and Reallocate Quantities.** Our third result examines the effects of procurement on market outcomes. We estimate these effects 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

FIGURE 4. Entry Timing and BE Approval Around Procurement



(A) Winner Entry (B) BE Approval: Entrants vs. Incumbents

Notes: Panel A plots the cumulative share of winners that have entered by quarter relative to procurement implementation and procurement wave. Panel B plots the fraction of non-branded drug-firm products with BE approval over time. Incumbents are products observed before 2017Q1; entrants are products in their first observed quarter.

studies use waves 2 through 8, so that each cohort has sufficient post-procurement observations. We estimate:

$$Y_{dimtc} = \sum_{k=-16}^{12} \beta_k \text{Treat}_{djmc} \mathbf{1}\{t - t_c = k\} + \lambda_{dimc} + \lambda_{tc} + \epsilon_{dimtc}, \tag{2}$$

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 5 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 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.

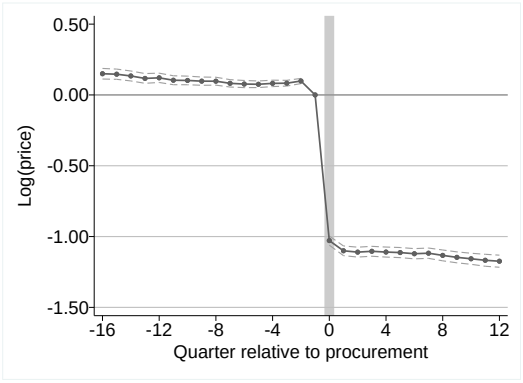
Panels E and F show non-winners, including losing bidders and non-participants. Their prices decline more moderately, by roughly 0.15–0.30 log points after procurement, while quantities fall by more than one log point and remain depressed. The residual market therefore contracts for firms that do not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

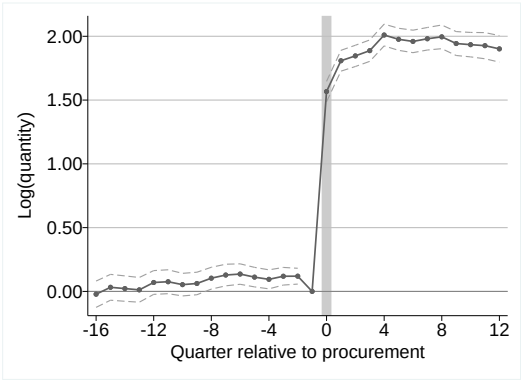
5. A Model of Quality Certification, Entry, and Procurement

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 firms and may restrict consumer choice through demand guarantees. Third, procurement outcomes affect downstream product-market competition because winning firms must supply

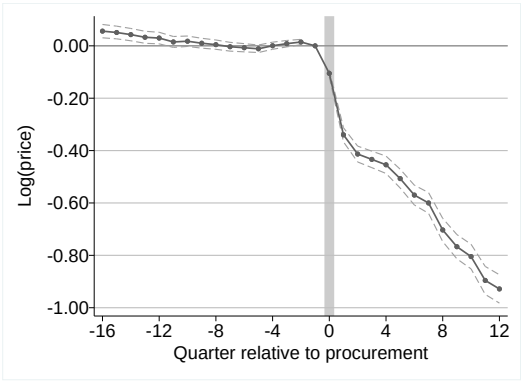
FIGURE 5. Procurement Effects on Prices and Quantities



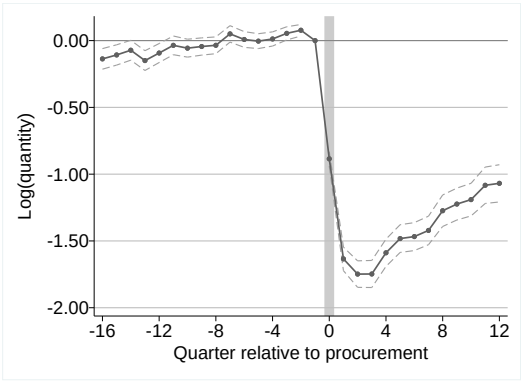
(A) Winners, Selected Provinces



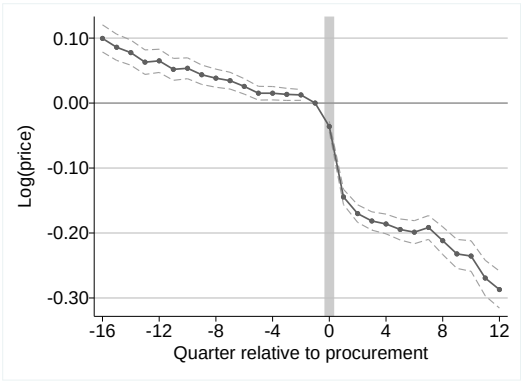
(B) Winners, Selected Provinces



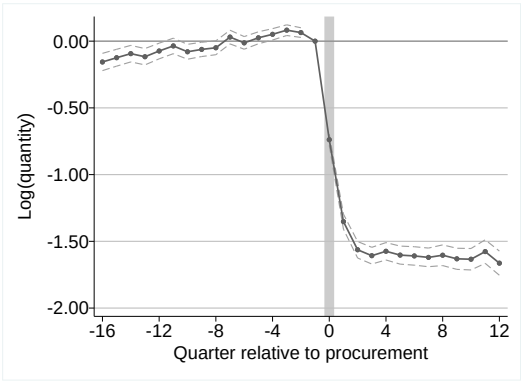
(C) Winners, Other Provinces



(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.

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.

5.1. **Setup.** There is one branded incumbent firm and $\bar{N}$ potential entrants producing generic drugs. Each potential entrant has a predetermined quality type $\theta \in \{H, L\}$, where a fraction ϕ 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. The model unfolds over four stages.

In the first stage, each potential entrant of type θ draws an entry cost e from distribution $F_{e|\theta}$ and decides whether to enter the market. Firms that enter produce a drug of their predetermined quality type θ.

In the second stage, 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 τ_i drawn from distribution F_τ. 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.[2]

We assume that the BE test is informative but imperfect:

$$\begin{aligned} \Pr(v = C \mid H) &= 1, \\ \Pr(v = C \mid L) &= \psi \in [0, 1]. \end{aligned} \tag{3}$$

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.

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

In the third stage, certified generic firms in the centralized procurement auction.[3] For simplicity, we assume that participation is costless. Auction participants privately observe a cost signal c drawn from the type-specific distribution $F_{c|\theta}$ and submit bids in the procurement auction. We assume that the density $f_{c|\theta}$ is log-concave.

In the final stage, 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.

5.1.1. *Consumer Demand.* We model drug choice as a joint decision between patients and physicians. Conditional on selecting a particular molecule or treatment category, a patient chooses among available products under physician supervision. Substitution across molecules or unrelated treatments is absorbed into the outside option.

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

$$u_{ijm} = \alpha\kappa_m p_{im} + \beta\mu(v_i) + \xi_i + \nu_{ijm}, \tag{4}$$

where $\kappa_m p_{im}$ denotes the patient's out-of-pocket payment, equal to the list price multiplied by copay rate κ_m.[4] The term $\mu(v_i)$ represents consumers' equilibrium belief that product i is high quality:

$$\mu(v_i) = \Pr(\theta_i = H \mid v_i).$$

For the branded drug, quality is known and consumers believe it is high quality with certainty, so $\mu(B) = 1$. For generic drugs, the posterior beliefs $\mu(C)$ and $\mu(U)$ are determined endogenously in equilibrium through firms' entry and certification decisions.

The term ξ_i captures a product-level demand shock that is orthogonal to quality certification and brandedness. We assume that ξ_i is independently drawn from distribution F_ξ after entry and procurement outcomes but before downstream price

[3]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]For drugs covered by the National Health Insurance system, κ_m is approximately 28%, although copay rates and coverage vary somewhat across provinces and drugs.

competition. The realization of ξ_i is observed by all firms prior to pricing. This timing preserves symmetry among auction participants while allowing for ex-post vertical differentiation across products.

Finally, ν_{ijm} follows a Type I extreme value distribution. The outside option includes alternative products outside our sample as well as the option of foregoing treatment, and its mean utility is normalized to zero.

We model the government 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 may choose freely among all available products. This formulation reflects the implementation of procurement quotas in practice, where physicians are strongly incentivized to prescribe the winning drug.

Let $Q^R_{im}(p^w_m, \xi_i)$ denote demand for product i, as the procurement winner, from consumers whose choices are restricted by the procurement program, where p^w_m is the winning price. Let $Q^U_{im}(p_{mi}, v_i, \xi_i, \mathbf{p}_{m,-i}, \mathbf{v}_{-i}, \boldsymbol{\xi}_{-i})$ denote demand for product i from consumers with unrestricted choice sets, where i might or might not be the procurement winner. This demand depends on the characteristics of product i, as well as the prices, certification statuses, and brandedness of competing products in market m.[5]

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

$$mc_{im} = c_i - \sum_{m'} \gamma_{m'} \mathbf{1}\{w_{m'} = i\} + \delta_m + \epsilon_{im}, \tag{5}$$

where c_i is the firm-specific cost signal observed prior to bidding, $w_{m'}$ denotes the winner of market m', and ϵ_{im} is an i.i.d. cost shock across firms and markets, drawn from F_ϵ.

The parameters $\gamma_m \geq 0$ and δ_m are publicly observed and common across firms. The term

$$\sum_{m'} \gamma_{m'} \mathbf{1}\{w_{m'} = i\}$$

[5]Because only certified generic firms are eligible to participate in the procurement auction, restricted consumers choose only among procurement-winning products. As a result, demand from restricted consumers depends only on the winner's price and not separately on certification status or brandedness.

captures the effect of procurement outcomes on production efficiency.[6] Winning additional markets may reduce marginal costs through economies of scale, production complementarities, or distribution efficiencies. The parameter δ_m captures market-specific determinants of production costs.

After the procurement auction and before the pricing stage, realized marginal costs become publicly observed.[7]

5.1.3. *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 consumers 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 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.

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

5.2.1. *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

[6]We allow γ_m to vary with market characteristics, including market size and the demand-guarantee rate λ_m.

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

[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.

each market, including the winner identity w_m and the winning price p_m^w. To simplify notation, define firm i's post-auction state in market m as

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

Conditional on rivals' prices $\mathbf{p}_{m,-i}$, firm i's post-auction profit in market m is

$$\pi_m(\mathbf{p}_{m,-i}, \Omega_{im}) = \begin{cases} (p_m^w - mc_{im})\Big[\lambda_m Q_{im}^R(p_m^w, \xi_i) \\ \qquad + (1-\lambda_m) Q_{im}^U(p_m^w, v_i, \xi_i, \mathbf{p}_{m,-i}, \mathbf{v}_{-i}, \boldsymbol{\xi}_{-i})\Big], & \text{if } i = w_m, \\ \max\limits_{p_i} (p_i - mc_{im})(1-\lambda_m) \\ \qquad \times Q_m^U(p_i, v_i, \xi_i, \mathbf{p}_{m,-i}, \mathbf{v}_{-i}, \boldsymbol{\xi}_{-i}), & \text{if } i \neq w_m. \end{cases} \tag{6}$$

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.[9]

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

$$\beta(c \mid n),$$

where c denotes the firm's private cost signal and n denotes the number of auction participants.[10] Since auction participation is costless, the number of participants equals the number of certified generic firms. Certification outcomes are publicly observed prior to bidding, so firms observe n 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, because procurement outcomes affect firm-wide production efficiency through (5), bidding incentives are interdependent across markets.

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

[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 θ affects the distribution of cost signals through $F_{c|\theta}$, but conditional on the realized signal c, it does not directly affect downstream marginal costs. Hence bidding strategies depend on the realized signal rather than directly on θ.

of rival costs. Let

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

denote firm i's expected continuation payoff in market m conditional on winning with bid b. This continuation value is derived from the post-auction profit function in (6), 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, continuation profits depend directly on the submitted bid. Similarly, let

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

denote firm i's expected continuation payoff conditional on not being assigned market m. Given rivals' strategies, firm i's expected payoff from bid b is

$$\begin{aligned}\pi^a(b, c_i, n; \beta) = \sum_{m \in M} \Bigg\{ & \Pr\left(\beta^{-1}(b) \in \left[c^{(r_m)}, c^{(r_m-1)}\right]\right) \\ & \times \mathbb{E}_{\mathbf{c}_{-i}}\Big[\pi_m^w\big(b, c_i, \beta(\mathbf{c}_{-i}), n\big) - \pi_m^\ell\big(c_i, \beta(\mathbf{c}_{-i}), n\big) \\ & \qquad \Big| \, \beta^{-1}(b) \in \left[c^{(r_m)}, c^{(r_m-1)}\right]\Big] \\ & + \mathbb{E}_{\mathbf{c}_{-i}}\Big[\pi_m^\ell\big(c_i, \beta(\mathbf{c}_{-i}), n\big) \Big| \, \beta^{-1}(b) \notin \left[c^{(r_m)}, c^{(r_m-1)}\right]\Big] \Bigg\}, \end{aligned} \tag{7}$$

where $\beta^{-1}(b) \in \left[c^{(r_m)}, c^{(r_m-1)}\right]$ corresponds to firm i submitting the r_mth lowest bid and thus getting assigned market m. We adopt the convention $c^{(0)} = -\infty$. In equilibrium,

$$\beta(c_i \mid n) \in \arg\max_b \, \pi^a(b, c_i, n; \beta). \tag{8}$$

The equilibrium bidding condition highlights the tradeoff created by centralized procurement. Lower bids increase the probability of securing desirable markets but also reduce downstream prices conditional on winning. When markets are assigned only to the lowest bidder ($r_m = 1$ for all m), firms bid more aggressively than under decentralized price competition because raising the bid directly reduces both the probability of winning and the cost of providing drugs. More generally, markets allocated to highly ranked bidders (small r_m) generate stronger competitive pressure and lower equilibrium bids.[11] In the Appendix, we provide sufficient conditions under which the symmetric monotonic equilibrium bidding strategy is unique.

[11] Under log-concavity of the underlying density, order statistics are unimodal.

5.2.3. *Test-Taking Strategies.* Let $\pi_{v,\theta}$ denote the expected payoff of a type-θ firm with certification outcome v. For a firm that passes the BE test ($v_i = C$), expected profit is

$$\pi_{C,\theta} = \mathbb{E}_{(c_i,n)}\left[\pi^a(\beta(c_i), c_i, n; \beta) \mid \theta\right]. \tag{9}$$

For firms that either fail or do not take the BE test ($v_i = U$),

$$\pi_{U,\theta} = \sum_{m\in M} \mathbb{E}\left[\pi_m(\mathbf{p}_{-i}, U, mc_{im}, \xi_i, \mathbf{v}_{-i}, \mathbf{mc}_{m,-i}, \boldsymbol{\xi}_{-i}, w_m, p_m^w) \mid \theta\right], \tag{10}$$

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

A type-θ firm takes the BE test if and only if expected gains exceed testing cost:

$$\pi_{C,\theta}\Pr(v = C \mid \theta) + \pi_{U,\theta}\Pr(v = U \mid \theta) - \tau_i \geq \pi_{U,\theta}.$$

Using (3), the equilibrium probability that a type-θ firm takes the test is

$$\rho_\theta = \begin{cases} F_\tau(\pi_{C,H} - \pi_{U,H}) & \text{if } \theta = H, \\ F_\tau\big(\psi(\pi_{C,L} - \pi_{U,L})\big) & \text{if } \theta = L. \end{cases} \tag{11}$$

Low-quality firms take the test with positive probability only when $\psi > 0$, since otherwise passing is impossible.

5.2.4. *Entry and Quality Choice.* Expected profit from entering as a high-quality firm is

$$\pi_H = \mathbb{E}_\tau\big(\pi_{C,H} - \tau \mid \tau \leq \pi_{C,H}\big)F_\tau(\pi_{C,H} - \pi_{U,H}) + \pi_{U,H}\big(1 - F_\tau(\pi_{C,H} - \pi_{U,H})\big).$$

Expected profit from entering as a low-quality firm is

$$\pi_L = \mathbb{E}_\tau\Big[\pi_{C,L} - \tau \mid \tau \leq \psi(\pi_{C,L} - \pi_{U,L})\Big]F_\tau\Big(\psi(\pi_{C,L} - \pi_{U,L})\Big) + \pi_{U,L}\Big[1 - F_\tau\big(\psi(\pi_{C,L} - \pi_{U,L})\big)\Big].$$

The probabilities of entering as high or low quality are

$$q_H = F_{e|H}(\pi_H) \text{ and } q_L = F_{e|L}(\pi_L).$$

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

$$\mu(C) = \frac{\rho_H \phi q_H}{\rho_H \phi q_H + \rho_L \psi(1-\phi) q_L}, \tag{12}$$

$$\mu(U) = \frac{(1-\rho_H)\phi q_H}{(1-\rho_H)\phi q_H + \big[1 - \rho_L + \rho_L(1-\psi)\big](1-\phi)q_L}, \tag{13}$$

which enter demand through (4).

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.

5.3. **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

$$u^0_{ijm} = \alpha\kappa_m p_{im} + \beta\mu^0_i + \xi_i + \nu_{ijm},$$

where $\mu_i^0 = 1$ for the branded drug and $\mu_i^0 = \mu^0$ 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_θ^0 denote the expected operating profit of a type-θ entrant in this regime. Entry satisfies

$$q_\theta^0 = F_{e|\theta}(\pi_\theta^0).$$

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

$$\mu^0 = \frac{\phi F_{e|H}(\pi_H^0)}{\phi F_{e|H}(\pi_H^0) + (1-\phi) F_{e|L}(\pi_L^0)}.$$

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, raises demand for certified products, and changes entry and testing incentives.

Centralized procurement adds a second channel. In addition to affecting consumer beliefs, certification can determine access to procurement auctions and subsequent market access. Winning procurement may generate guaranteed demand and scale-related cost advantages, which affect firms' incentives to enter, test, bid, and compete in downstream markets. Thus, in the model, certification and procurement interact: certification provides information about quality, while procurement changes the payoff from becoming certified by linking certification to market access and scale.

6. 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.[12]

6.1. **Model Primitives.** The primitives of the model consist of: (i) the entry environment, including the type-specific entry-cost distributions $F_{e|\theta}$ and the fraction of high-quality firms among potential entrants, ϕ; (ii) the certification technology, including the testing-cost distribution F_τ and the probability that a low-quality firm passes the BE test, ψ; (iii) demand-side parameters, including price sensitivity α, quality sensitivity β, and the distribution of product-level demand shocks F_ξ; and (iv) supply-side parameters, including the type-specific cost-signal distributions $F_{c|\theta}$, market-specific scale parameters γ_m, market-specific cost shifters δ_m, and the distribution of idiosyncratic cost shocks F_ϵ. We take as given the market-specific copay

[12]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.

rates κ_m, the government-guarantee parameters λ_m, the procurement assignment rule $\{r_m\}_{m\in M}$, and the potential number of entrants $\bar{N}$.

6.2. **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 $\mu(v)$ 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.

6.2.1. *Step 1.* Using observed prices, quantities, certification statuses, and product availability, we identify the price coefficient α, product-level demand shocks ξ_i, and the perceived-quality indices

$$(\beta\mu(C),\ \beta\mu(U),\ \beta\mu^0).$$

Demand alone identifies the product of the quality-sensitivity parameter β and consumers' posterior beliefs. The separation of β from posterior beliefs comes from the equilibrium restrictions on entry, testing, certification, and cost-signal distributions described below.

6.2.2. *Step 2.* Next, we recover marginal costs from the downstream pricing first-order conditions.[13] We then use the cost equation to decompose marginal costs into firm cost signals c_i, market-specific cost shifters δ_m, scale effects from procurement γ_m, and firm-market-specific cost shocks ϵ_{im}. Variation in procurement assignment around the relevant bid-rank thresholds identifies the scale parameters γ_m, and the panel structure of the data identifies the firm cost signals c_i.

This step recovers the distributions of cost signals conditional on certification and testing outcomes, as well as the cost shock distributions F_ϵ. Although consumers observe only whether a product is certified, the researcher observes whether an uncertified product failed the BE test or was not tested. We therefore distinguish between two uncertified states: U^F, denoting products that failed the BE test, and

[13]In the baseline estimation, we do not use the bidding first-order conditions when calibrating the government-guarantee parameters λ_m. Alternatively, the bidding first-order conditions can help identify λ_m, since λ_m determines the value of winning each market.

U^N, denoting products that were not tested. Let the recovered conditional cost-signal distributions be

$$G_{c|C},\, G_{c|U^F},\, G_{c|U^N}.$$

The BE testing technology implies that high-type products pass the test with probability one, while low-type products pass with false-positive probability ψ. Therefore, a product that fails the BE test must be low type. It follows that

$$G_{c|U^F}(c) = F_{c|L}(c).$$

By contrast, certified products and non-tested products are mixtures of high and low types. A key timing assumption is that the cost signal c_i is realized and observed by the firm only after entry, testing, and certification decisions have been made. Therefore, these decisions do not select firms on the realized cost signal. Certification and testing outcomes change the mixture of latent types, but not the within-type distribution of c_i. With that, for any realization c,

$$G_{c|C}(c) = \mu(C)F_{c|H}(c) + (1-\mu(C))\,F_{c|L}(c), \tag{14}$$

$$G_{c|U^N}(c) = \mu(U^N)F_{c|H}(c) + \left(1-\mu(U^N)\right)F_{c|L}(c). \tag{15}$$

Here, $\mu(U^N)$ denotes the probability that a non-tested entrant is high type:

$$\mu(U^N) = \frac{(1-\rho_H)\phi q_H}{(1-\rho_H)\phi q_H + (1-\rho_L)(1-\phi)q_L}.$$

This object differs from $\mu(U)$, the posterior belief for uncertified products in demand, because consumers pool products that failed the BE test and products that were not tested.

6.2.3. *Step 3.* We then use a system of equilibrium restrictions to recover

$$\left(\beta, \phi, \psi, \rho_H, \rho_L, q_H, q_L, q_H^0, q_L^0\right). \tag{16}$$

While (β, ϕ, ψ) are model parameters, the rest are equilibrium objects: type-specific testing probabilities ρ_θ and type-specific entry probabilities before and after procurement, q_θ^0 and q_θ.

The first set of restrictions comes from observed entry, testing, and certification outcomes: the testing probability, the certification probability conditional on testing, and entry probabilities in the pre-certification and post-procurement regimes. The second set comes from the perceived-quality indices identified in Step 1, $(\beta\mu(C), \beta\mu(U), \beta\mu^0)$. The third set comes from the cost-signal mixture restrictions in

Step 2, which provide additional discipline because certified products and non-tested products are different mixtures of high- and low-quality types. Specifically, using equations (14) and (15), together with $G_{c|U^F}(c) = F_{c|L}(c)$, we obtain the restriction

$$\mu(U^N)\{G_{c|C}(c) - G_{c|U^F}(c)\} - \mu(C)\{G_{c|U^N}(c) - G_{c|U^F}(c)\} = 0. \tag{17}$$

This restriction holds for any c in the common support of the relevant cost-signal distributions. In principle, the cost-signal distributions therefore provide many over-identifying restrictions. Here we aggregate (17) into one moment by integrating it against a prespecified weighting function. Together, the observed entry, testing, certification moments, the perceived-quality indices, and the cost-mixture moment provide restrictions for the objects in (16).

6.2.4. *Step 4.* Finally, given the recovered beliefs, the false-positive probability, and the demand and cost parameters from the previous steps, we solve the model to compute the type-specific profits relevant to testing ($\pi_{C,H},\ \pi_{U,H},\ \pi_{C,L},\ \pi_{U,L}$) and entry ($\pi_H,\ \pi_L,\ \pi_H^0,\ \pi_L^0$). These model-implied profits convert the reduced-form cutoff probabilities recovered in Step 3 into restrictions on the primitive distributions. Specifically, the testing moments identify

$$F_\tau(\pi_{C,H} - \pi_{U,H}) \quad \text{and} \quad F_\tau\left(\psi(\pi_{C,L} - \pi_{U,L})\right),$$

while the entry moments identify, for $\theta \in \{L, H\}$

$$F_{e|\theta}(\pi_\theta), \quad F_{e|\theta}(\pi_\theta^0).$$

Under parametric restrictions on F_τ and $F_{e|\theta}$, these restrictions identify the parameters of the testing-cost and entry-cost distributions.

7. Results and Conclusion

We are currently completing the estimation and will report the corresponding counterfactual results in the next version. We plan to use the estimated model to study how procurement changes entry and competition in the presence of information frictions. The counterfactuals are organized around three questions. First, is certification alone sufficient to induce high-quality generic entry? Second, how much of procurement's effect operates through guaranteed market access versus scale-related cost reductions? Third, how can procurement design be improved?

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 ($\gamma_m = 0$), 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 studies procurement design. Procurement design changes the reward to certification by altering both the allocation and the size of guaranteed demand. We therefore consider alternative assignment rules and demand-guarantee rates, varying r_m and λ_m. In particular, we examine a more concentrated assignment rule that allocates larger markets to the lowest bidder, thereby increasing the scale benefits of winning but reducing the number of firms that receive procurement access. This design may lower procurement prices further by intensifying competition for the top rank and allowing the winner to exploit larger scale economies. However, when certification is imperfect, increasing the prize from winning procurement also raises the value of certification for low-quality firms that may pass the test. As a result, a more concentrated design can increase low-quality entry or certification attempts, and may reduce consumer welfare if a larger share of restricted consumers are shifted toward lower-quality winners.

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.

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.

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, 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.

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 Chinaâs National Drug Procurement Auction," Technical Report, mimeo 2023.

Appendix A. 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.