Competitive Bidding in Drug Procurement: Evidence from China

Shengmao Cao, Lisa Xuejie Yi, Chuan Yu*

January 2022

Abstract

We study the impact of competitive bidding in the procurement of off-patent drugs. In 2019, China introduced competitive bidding with a quantity guarantee for thirty-one molecules in nine provinces. Using a difference-in-difference design, we show that the program reduced average drug prices by 47.4%. Generic drug firms won the majority of the bids and on average cut prices by 59.4%. We develop a model of demand and supply to quantify the trade-off between lower prices and choice distortions. Competitive bidding increases consumer welfare if policymakers consider brand preferences welfare irrelevant. The program also reduced government expenditures on insurance by 24.3%.

Keywords: competitive bidding, procurement auctions, drug prices, consumer welfare
JEL Codes: I11, I18, L51

*Contact: Cao: Stanford University, shengmao@stanford.edu; Yi: Stanford University, xuejieyi@stanford.edu; Yu: Stanford University, chuanyu@stanford.edu. We thank Luis Armona, Jose Ignacio Cuesta, Liran Einav, Matthew Gentzkow, Caroline Hoxby, Neale Mahoney, Alessandra Voena, Heidi Williams, Ali Yurukoglu, and seminar participants at the Stanford IO workshop for helpful comments. All errors are our own.

1 Introduction

Drug prices have remained high in many pharmaceutical markets after patent expiration and generic entry. Branded off-patent drugs continue to sell at high prices because consumers prefer these drugs to their generic counterparts (Caves et al. 1991; Grabowski and Vernon 1992; Bronnenberg et al. 2015).[1] Prices of many generic drugs also far exceed their production costs (Dubois and Lasio 2018; Cuddy 2020). High drug prices contribute significantly to health care expenditures in many countries, imposing a burden on patients and governments.

In an effort to reduce prices of off-patent drugs, the Chinese government launched a competitive bidding procurement program for prespecified quantities of drugs in public hospitals. While competitive bidding may reduce drug prices and benefit consumers, the quantity guarantee could distort consumer choices and reduce allocative efficiency. In this paper, we evaluate, both theoretically and empirically, the effects of the competitive bidding program on drug prices, consumer welfare, and government expenditures.

Before the competitive bidding program, hospitals in China made their own drug procurement decisions. Manufacturers set wholesale prices at the provincial level. Each hospital decided which drugs to procure and then sold them to patients at the same price.[2] Drug prices have been high for most molecules despite generic competition. Consumers' brand preferences and insurance coverage make branded drugs popular choices for hospitals despite the drugs' higher prices. In addition, sales costs and detailing costs add to the prices of all drugs.

The Chinese government implemented the competitive bidding program for thirty-one molecules in nine provinces in 2019. For each molecule, the branded drug firm and all bio-equivalent (BE) generic drug firms were eligible to participate in the bidding.[3] The program took place in two rounds. The first round involved a first-price sealed-bid auction to select the winner with the lowest

[1] A branded drug, also known as an innovator drug, is the first drug created with specific active ingredients to receive approval for use.

[2] Hospitals in China cannot charge a markup when selling drugs to consumers, as stipulated by the zero-markup policy (Fang et al. 2021). The policy has been implemented in a staggered fashion across provinces since 2009.

[3] China started certifying the bio-equivalence of generic drugs and branded drugs in 2016. Bio-equivalence certification requires rigorous laboratory evidence that the generic drug has the same clinical effectiveness and safety profile as the branded drug.

bid, with the requirement that each participating firm must bid at least 10% below its pre-bidding price. In the second round, government officials could negotiate with the winner of the first round to attempt to further reduce the price. Otherwise, an agreement at the price of the winning bid was automatically reached, and the winner could *not* renege. Once an agreement was reached, the winner committed to selling at the agreed price for the year between March 2019 and February 2020 and was guaranteed a prespecified procurement quantity.[4] Agreements were reached for twenty-five molecules, while negotiations stalled for the remaining six.

We analyze the impact of this competitive bidding program on drug prices and consumer welfare in three steps. First, to provide intuition on the potential effects of the program, we consider a stylized model of competitive bidding with one branded and one generic product. Both products have the same therapeutic value and marginal cost, but all consumers prefer the branded product. The two firms first participate in an auction. The winner commits to the winning bid and is guaranteed a sales quantity. Faced with the residual demand, the loser sets its price to maximize profits.

Our model yields four predictions. First, the generic drug firm always wins the auction. Since the generic drug firm can only sell at a price below the winning bid should it lose the auction, it will undercut any bid submitted by its rival. Second, the price change of the winner is ambiguous. By requiring the winner to commit to its bid, the program turns a Bertrand game into a Stackelberg game, which reduces competition and may more than offset the procompetitive incentive from the quantity guarantee, in particular when the guaranteed quantity is low. This concern is addressed by the requirement of a 10% price reduction in our setting; nevertheless, it highlights the importance of a reservation price in the design of procurement auctions with similar formats. Third, the price change of the losing (branded) product is also ambiguous even under a lower generic drug price, as there is a countervailing incentive to increase the price and target consumers with stronger brand preferences. Last, the effect on consumer surplus depends crucially on how we interpret consumers' brand preferences. When we assume brand preferences are welfare irrelevant,

[4]The prespecified quantity was usually between 60% and 70% of the respective procurement quantity for each of the nine provinces in the previous year. The central government split the quantity target among hospitals. A hospital that fell short of its target was subject to punishments such as reduced government funding, lower ratings, and demotion of hospital management.

competitive bidding always increases consumer surplus, conditional on a decrease in the price of the generic product. When we take consumers' brand preferences into account, the effect on consumer surplus becomes ambiguous and depends on the relative magnitudes of two countervailing forces: the price reduction and the choice distortion induced by the quantity guarantee.

Next, we turn to our empirical setting and estimate the price impact of the competitive bidding program in China. In our baseline specification, we leverage quarterly drug price and sales data at the product-province level for the twenty-five molecules for which agreements were reached. We use a difference-in-difference (DID) design in which we compare market outcomes in provinces that introduced competitive bidding with outcomes in provinces that did not.

We highlight three main empirical findings. First, generic drug firms won the bidding for twenty-two of the twenty-five molecules.[5] Second, winners on average reduced their prices by 59.4% and increased their sales by 68.4%. This large price drop provides direct evidence that markups on off-patent drugs were high before competitive bidding despite generic competition. Third, for the twenty-two molecules for which generic drug firms won the bidding, the losing branded drug firms on average reduced their prices by 7.5%. Overall, the competitive bidding program reduced average drug prices by 47.4%. Our results show that moving from decentralized procurement to competitive bidding can significantly reduce drug prices.

Our final step is to quantify the effects of the competitive bidding program on consumer welfare and government expenditures in the context of the market for hypertension drugs.[6] We focus on a key welfare trade-off: consumers benefit from lower prices, but some consumers experience a welfare loss because they are forced to substitute generic drugs for branded drugs because of the quantity guarantee.

To quantify the welfare effects of price reductions and consumer choice distortions, we need to disentangle the roles that price changes and the quantity guarantee played in reshaping consumers' drug choices in 2019. To achieve this goal, we first use data for the period prior to the competitive bidding program to estimate a model of drug demand and supply. The estimated price elasticity

[5]This result is broadly, though not completely, consistent with the first prediction of our stylized model. Branded drug firms may win the auction for reasons not captured by the model, such as capacity constraints or differences in marginal costs.

[6]The competitive bidding program covered eight main drugs that treat hypertension.

allows us to infer to what extent the sales changes in 2019 were driven by price reductions. Second, we attribute the remaining changes in sales in provinces that participated in the program relative to sales in other provinces to the quantity guarantee. Our estimates imply that price reductions and the quantity guarantee explain 21.1% and 78.9% of the sales gains by the auction winners in 2019, respectively.

Our main finding is that the net welfare effects of the competitive bidding program depend on whether policymakers consider consumers' brand preferences welfare relevant. On the one hand, if policymakers focus only on the clinical value of the drugs, the program increases consumer welfare by 39.2%, as it significantly reduces drug prices and steers consumers away from the more expensive branded drugs. On the other hand, if policymakers also care about consumers' brand preferences, the program reduces consumer welfare by 9.8%, because consumer choice distortions outweigh the benefits of the price drop. In either case, the program reduces government expenditures on insurance payments by 24.3%.

We end by highlighting one key caveat to the interpretation of our results. We focus on the short-term effects of the policy experiment and abstract from several important long-run dynamic considerations. Firms that lose the bidding may gradually exit the market, and the policy may become less effective as the market becomes more concentrated. In addition, by greatly reducing total producer surplus, the policy may limit firms' R&D incentives in the long run. Still, we consider this paper a first step toward understanding the potential of competitive bidding to reduce prices of off-patent drugs.

This paper contributes to a long line of research on the pricing of off-patent drugs. An extensive literature investigates the effects of generic entry after innovator brands go off patent in developed countries (Caves et al. 1991; Grabowski and Vernon 1992; Frank and Salkever 1997; Saha et al. 2006; Huckfeldt and Knittel 2011; Branstetter et al. 2016). Several recent studies investigate the nature of brand preferences (Colgan et al. 2015; Bronnenberg et al. 2015; Bairoliya et al. 2017) and their implications for the pricing and regulation of off-patent drugs (Atal et al. 2019). Some studies evaluate other price control policies, such as price caps (Mohapatra and Chatterjee 2016; Dean 2019) and mandatory generic substitution (Song and Barthold 2018). We show that competitive bidding has been effective in reducing off-patent drug prices in China's pharmaceutical industry though its welfare implications depend on the welfare interpretation of consumers' brand

preferences. Our results point to a policy tool to promote generic substitution and reduce prices of off-patent drugs.

This paper is closely connected to a small literature on competitive bidding in health care. Dubois et al. (2021) use cross-country variation to study centralized drug procurement in developing countries. A number of studies have investigated competitive bidding in Medicare Advantage (Song et al. 2013; Duggan et al. 2016; Cabral et al. 2018; Curto et al. 2021). The two papers closest to ours are Ji (2019) and Ding et al. (2021), which leverage the staggered roll-out of competitive bidding for the supply of durable medical equipment in the US to study its impact on equipment prices and sales. Our study uses a similar quasi-experimental design but focuses on a policy setting in an emerging economy. Our results show that competitive bidding can reduce drug prices in a setting with market frictions such as sticky brand preferences. We also contribute to the literature by quantifying the welfare effects of competitive bidding on consumers.

This paper also relates to the literature on auction theory (see Klemperer 1999 for a review). We study a novel auction format that guarantees a prespecified quantity to the winner but requires the winner to sell products at the winning bid. We show theoretically that this auction format has ambiguous price effects and empirically that it can address market frictions by inducing substitution to generic drugs and reducing prices. Our results could inform the design of similar auctions in other settings.

The remainder of the paper is organized as follows. Section 2 describes the policy background. Section 3 outlines a simple model that illustrates the potential effects of centralized competitive bidding on drug prices and consumer welfare. In Section 4, we describe the data and provide our main estimates of the effects of the program on market outcomes. In Section 5, we develop and estimate a model of demand and supply for hypertension drugs. Section 6 uses the model to quantify the welfare effects of the policy. Section 7 concludes with a discussion of our findings and their policy implications.

2 Background

As of 2015, pharmaceutical expenditures accounted for more than 40% of health care spending in China, while the OECD average was less than 20% (OECD 2018, 2019). Public hospitals

account for over 70% of pharmaceutical sales in China (Mossialos et al. 2016). In this section, we discuss major frictions in drug procurement by public hospitals in China prior to 2019 and describe the competitive bidding program.

2.1 The Scenario Prior to Competitive Bidding

Historically, public hospitals in China made their own drug procurement decisions. Pharmaceutical firms set wholesale prices for their drug products at the provincial level. Each hospital decided which products to procure and then sold the drugs to patients for inpatient use or through hospital pharmacies. The zero-markup policy, which was first implemented in 2009, mandated that hospitals sell drugs to consumers at the procurement cost (Fang et al. 2021). Most off-patent drugs have multiple generic competitors. In 2016, China started certifying the bio-equivalence (BE) between generic drugs and branded drugs. BE generic drugs have passed rigorous clinical tests demonstrating that their therapeutic efficacy and safety profile are identical to those of their branded counterparts.

Despite generic competition, prices of these off-patent drugs have been high because of several market features. First, many consumers still prefer branded drugs over BE generic drugs. As a result, branded drugs remain popular choices for hospitals despite their higher prices. Second, generous public health insurance has reduced consumers' price sensitivity and removed an important check on drug prices. By the end of 2018, around 96.8% of the population in China was covered by basic medical insurance.[7] Public insurance covers around 72% of the costs for a list of drugs that the government considers essential for public health. Finally, sales costs and detailing costs further add to drug prices. Pharmaceutical firms engage in aggressive sales activities to promote their products. Sales commissions typically amount to 20% to 30% of drug prices (Yang and Fan 2012; Zhang et al. 2014).

2.2 The Competitive Bidding Program

To mitigate market frictions and reduce drug prices, the Chinese government implemented a pilot competitive bidding program at the end of 2018. The program covers thirty-one molecules

[7] See http://www.nhc.gov.cn/mohwsbwstjxxzx/qgdlcws/lcdc.shtml. Accessed September 9, 2021.

with at least one BE generic product. It was implemented in eleven major cities, which account for around one-third of drug procurement in China and are located in nine different provinces.[8] For each molecule, the branded drug firm and all BE generic drug firms were eligible to participate in the bidding. The number of participating firms for each molecule ranged from two to six and was on average three.

The program consisted of two rounds, as outlined in Appendix Figure A.1. The first round consisted of a first-price sealed-bid auction to select the winner with the lowest bid, with the requirement that each participating firm must submit a bid at least 10% below its pre-bidding price.

In the second round, government officials negotiated with the winner from the first round to attempt to further cut the price. If officials demanded a price lower than the winner could afford, the winner could exit the negotiation without facing punishment. However, officials could always return to the winning bid at any time during the negotiation, in which case a deal would be reached automatically and the winner could not renege. As a result, the final negotiated price was always weakly below the winning bid. Once an agreement was reached, the winner committed to selling the drug at the agreed price for the year between March 2019 and February 2020 and was guaranteed a prespecified sales quantity ranging from 60% to 70% of the previous year's sales volume in public hospital procurement.

The stated goal of the policy was to bring about a large price reduction on the part of the auction winner and to induce follow-on price reductions on the part of losing firms. An agreement was reached for twenty-five molecules, while negotiations stalled on the remaining six.

To enforce the quantity guarantee, the central government imposed a quantity target on each public hospital during the specified period. For hospitals that fell short of the target, the punishments included reduced government funding, lower hospital ratings, and demotion of hospital management. Therefore, hospital management made it a priority to require physicians to persuade patients to switch to the winning products to meet the target.

[8] In China's administrative system, the first level is province and municipality, and the second level is city. Each province usually consists of around ten cities, while each municipality is built around one large city. The program covers four municipalities (Beijing, Tianjin, Shanghai, and Chongqing) and seven cities (Shenyang, Dalian, Xiamen, Guangzhou, Shenzhen, Chengdu, and Xi'an) located in five different provinces.

3 A Stylized Model of Competitive Bidding

In this section, we develop a simple model of competitive bidding and drug pricing to illustrate the potential effects of competitive bidding on drug prices and consumer welfare.

3.1 Model Setup

We consider a molecule for which one firm sells the branded product B and another firm sells a BE generic product G.[9] Both products provide an identical clinical value of 1 to a unit mass of consumers. Consumers have brand preferences for product B concerning how much more they are willing to pay for B than for G. We assume brand preferences are heterogeneous in the population and are drawn independently from a continuously differentiable distribution $\mathscr{F}$ with nonnegative support. We further assume $\mathscr{F}$ has no mass at 0. The marginal costs of both products are assumed to be 0 for simplicity. Prior to the competitive bidding, firms set prices to maximize profits simultaneously in a Bertrand game with differentiated products.

Under competitive bidding with a quantity guarantee of $\tau \in (0,1]$,[10] firms bid and set prices in a two-stage game. In the second stage, the losing firm sets a price to maximize profits, taking the winner's price and the quantity guarantee as given. If it sets a price such that its unconstrained quantity exceeds $1-\tau$, it will only sell to the $1-\tau$ consumers who value its product the most.[11]

Foreseeing the results of the second stage, both firms bid in a first-price sealed-bid auction with complete information in the first stage. The winner commits to selling at the winning bid. It is guaranteed a quantity of τ, provided that at least τ consumers prefer buying the winning product to the outside option. The reason is that while the government could mandate substitutions from the losing product to the winning product, it cannot force consumers to buy a product whose value is below its price.

[9] We consider a two-product case for simplicity of exposition. Empirically, fourteen out of the twenty-five molecules that completed the bidding had two eligible bidders, while the remaining eleven had three to six.

[10] In all equilibria that we consider in this section, all consumers purchase one product. The quantity guarantee is thus equivalent to a market-share guarantee.

[11] For example, if the branded product B loses the bidding and its price is set such that its unconstrained quantity exceeds $1-\tau$, it can sell only to the $1-\tau$ consumers with the strongest brand preferences. Some consumers, even if they prefer the branded product, would have to purchase the generic product under the quantity guarantee.

3.2 Effects of Competitive Bidding

We use the model to illustrate the potential effects of competitive bidding on drug prices and consumer surplus. We consider an example in which $\mathscr{F}$ follows a uniform distribution $\mathbb{U}[0,1]$. We discuss the intuition for four predictions from the model. Additional details on proofs and derivations of the results are provided in Appendix A.

First, generic product G always wins the auction. The intuition is that if the generic drug firm were to lose the bidding in the first stage, then in the second stage, it would have to set a price below the winning bid to sell any positive quantity, as all consumers weakly prefer the branded product. Therefore, the generic drug firm always prefers to win the bid in the first stage to secure the guaranteed quantity and will undercut any positive bid above its cost by the branded drug firm. A formal proof of this result for a general $\mathscr{F}$ is provided in Appendix A.

Second, the price change by the bid winner is ambiguous. Figure 1 compares the winning bid to the price of the generic product G under Bertrand competition. Competitive bidding reduces the price of the winning product only when τ is above a certain threshold. At low values of τ, competitive bidding turns the Bertrand game into a Stackelberg game. The bidding allows firm G to first commit to a price, thereby reducing price competition. More specifically, when τ is low, for any bid by firm G, firm B prefers targeting the remaining $1-\tau$ consumers over undercutting the bid. Thus, with complete information, firm G wins with a maximum bid of 1, which is three times the price in the original Bertrand game.

As τ increases, the branded drug firm B has a greater incentive to undercut the bid, and the winning bid decreases as a result. Competitive bidding intensifies competition by decoupling brand preferences from drug choices for τ consumers. When $\tau = 1$, the auction completely eliminates vertical product differentiation and drives prices down to marginal costs. If the policymaker's only objective is to reduce short-run drug prices, $\tau = 1$ is the optimal policy design. In practice, however, the policymaker may pick some $\tau < 1$ to preserve product varieties and avoid inducing excessive firm exits.

Third, conditional on a price reduction by the winner, the price response by the losing firm is also ambiguous. On the one hand, the losing firm has an incentive to reduce the price, as prices are strategic complements. On the other hand, it has the opposite incentive, namely, to increase prices

to target the remaining $1-\tau$ consumers who have stronger brand preferences. In the example that we consider, the second force dominates, so the price of the branded product B increases relative to that under Bertrand competition.

Last, the effect of competitive bidding on consumer surplus depends crucially on how we measure consumer surplus. We consider two different measures. The first one ignores any brand preferences and focuses only on the clinical value of the drugs. As both products deliver an identical clinical value of 1, consumers' utility is simply 1 minus the price that they pay. The second measure follows consumers' revealed preferences and treats consumers' brand preferences as welfare relevant. Therefore, forcing a consumer to switch from the branded product to the BE generic product might incur a welfare loss even if the consumer could pay a lower price.

Figure 2 plots the changes in consumer surplus from Bertrand competition to competitive bidding under both welfare measures, focusing on the range of τ where the price of the generic product decreases. There are two main takeaways. First, when we assume brand preferences are welfare irrelevant, competitive bidding always increases consumer surplus, conditional on a decrease in the price of the generic product. This reflects both the price reduction and the substitution to a cheaper product with an identical clinical value. Second, when we take consumers' brand preferences into account, the effect of competitive bidding on consumer surplus is ambiguous and depends on the value of the quantity guarantee. The direction is determined by the relative magnitudes of two countervailing forces: the price reduction and the choice distortion induced by the quantity guarantee.

3.3 Discussion

We have analyzed a highly stylized model that abstracts from several features of the actual competitive bidding program in our setting. We discuss three major simplifications and their implications for the theoretical intuitions from the model.

First, we do not consider negotiation between the bid winner and government officials. In a game with perfect information, the negotiation stage does not change firms' bidding strategies or the identity of the bid winner. Importantly, if government officials agree to accept the winning bid, the bid winner cannot renege on it, and a deal is automatically reached. This rules out a strategy

to first win the auction with a low bid and then try to negotiate a better deal afterward (or quit the negotiation and return to the status quo). Thus, the agreed price from the negotiation is always weakly lower than the winning bid. Together with the requirement of a 10% price reduction, this design helps preclude perverse outcomes such as a price increase by the bid winner in practice.

Second, we have considered a simple setting with one branded firm and one generic firm, which reflects the actual case for 14 out of the 25 molecules in our sample. The main qualitative intuition is similar in a more general setting with more than two differentiated products.[12] The firm with the lowest market share prior to competitive bidding wins the auction. The quantity guarantee intensifies market competition, but the post-bidding pricing game could be less competitive with one fewer firm. The second effect is weaker when the baseline number of firms is larger.

Last, we have assumed that the marginal costs of both products are the same and do not change after competitive bidding. In practice, heterogeneity in costs may affect the identity of the winner. In addition, the quantity guarantee could significantly cut sales and detailing costs, which could help further reduce the winner's price.

4 The Market Impact of Competitive Bidding

4.1 Data

The main data set for our empirical analysis contains quarterly product-level quantities and revenues for the thirty-one molecules covered by the competitive bidding program in each province. A product is defined as a molecule-firm combination.[13] The data were collected by China's Food and Drug Administration based on its audits of major public hospitals in twenty-four provinces in mainland China. We divide revenues by quantities to derive the average price of each product at

[12] If there are two or more homogeneous generic drug products, they price at marginal cost under both Bertrand competition and competitive bidding. In our empirical setting, we observe large variation in prices and sales among BE generic products of the same molecule, which suggests some product differentiation.

[13] The raw data are recorded at the product-dosage level, where the same product could be represented in different dosages. We convert quantity to the most common dosage of each molecule, and we aggregate quantities and revenues to the product level. For example, if we observe 1,400 units of 10 mg lisinopril tablets and 2,000 units of 20 mg lisinopril tablets, we convert the former to 700 units of 20 mg lisinopril tablets.

the province-quarter level. This is the average price that hospitals pay to procure the drugs and in principle is equal to how much they charge patients under the zero-markup policy.

The eleven cities covered by the competitive bidding program are located in nine provinces. We refer to them as *enacting provinces* and consider them directly affected by the policy. The other fifteen provinces in our data are considered *nonenacting provinces* and not affected by the policy.[14] Selection of the enacting provinces was certainly not random: enacting provinces tend to be richer, more populous and more cosmopolitan than nonenacting provinces. We refer to the twenty-five molecules for which an agreement was reached as *bidding molecules* and to the other six as *nonbidding molecules*.

We complement our main data set with two ancillary ones. First, we collected data on the set of generic products that passed the bio-equivalence test for each of the thirty-one molecules by the end of 2018 and when they passed the test. This identifies the participants eligible for competitive bidding. Second, we observe the identity of the auction winners and the final prices after the negotiation stage for the twenty-five bidding molecules. We verify that the prices inferred from our quarterly sales data at the provincial level align reasonably well with the disclosed final prices, as shown in Appendix Figure A.2.

Table 1 reports the summary statistics for all thirty-one molecules in our sample. These molecules treat a wide range of common health conditions, including diabetes, cardiovascular diseases, depression, and viral infection. We highlight three main takeaways from the table. First, compared with bidding molecules, nonbidding molecules on average bring in much smaller revenues (24.4 million versus 142.7 million RMB per quarter prior to 2019) and are sold by a significantly larger number of competing firms (45.7 versus 10.1 in 2018). Second, the average bidding molecule sells as three BE products, and more than half sell as only two. The BE requirement ensures the quality of products that participate in competitive bidding but makes the bidding less competitive. Finally, for most molecules, the branded product was significantly more expensive

[14]This definition is not perfectly precise, as each province usually has around ten cities and the policy covers only one or two cities in each province. However, we believe this is a reasonable approximation for two reasons. First, treated cities are the largest ones in each province and account for a large fraction of the population and drug prescriptions. Second, hospitals sampled by the data provider are mainly in large cities. As we show in Section 4.3, our results are robust to including only the four municipalities and excluding partially treated provinces.

than the average BE generic product prior to 2019. The median ratio of the branded to BE generic price is 1.7. As discussed in Section 2, consumers' brand preferences are a key contributor to high drug prices and a main motivation for the competitive bidding program. In the following analysis, we focus on medium and large firms with a cumulative market share of over 95% for each molecule.

In Figure 3, we plot the market shares for auction winners in enacting provinces prior to (in 2018) and after (in 2019) the competitive bidding program. Each dot is a molecule-province pair for all bidding molecules and enacting provinces. Since the quantity guarantee of the competitive bidding program went into effect in March 2019, we drop the first quarter when we compute 2019 market shares. For most molecule-province pairs, the market shares of the auction winners showed an increase in 2019, which could have been driven by both the quantity guarantee and the price reductions. Even if a winner already had a market share over 70% in 2018 and did not necessarily directly benefit from the guarantee, the price drop would still have boosted its market share. Winners' market shares were on average 74.9% in 2019. In a few cases, winners' market shares fell short of 60%. This may have occurred because of partial treatment of some provinces, inaccurate projection of the quantity guarantee, or unsuccessful enforcement by the government.

Appendix Figure A.3 shows the market shares for branded and BE generic drugs in enacting provinces before and after competitive bidding, regardless of whether they won the auction. The main takeaway is that this program induced large-scale generic substitution. Branded products, despite their significantly higher prices, commanded large market shares in 2018. Their market shares fell significantly in 2019, while the market shares of BE generic drugs increased from an average of 35.6% to 64.7%.

4.2 Empirical Strategy

In this section, we describe our empirical strategy for evaluating the effects of the competitive bidding program on various market outcomes. Let j denote molecules, m denote provinces, and t denote quarters. We estimate the effects of the policy experiment on drug prices, quantities, and market structure using a DID framework. In our baseline specification, we focus on bidding molecules and compare enacting and nonenacting provinces by using the following event-study

specification:

$$y_{jmt} = \alpha_0 + \sum_{q=4}^{-27} \beta_q Enact_m \mathbb{1}(t = q) + \lambda_{jm} + \lambda_{jt} + \varepsilon_{jmt}. \quad (1)$$

Each observation is a molecule (j)-province (m)-quarter (t) combination. Our primary outcome of interest y_{jmt} refers to the log of the average price and the log of the total quantities of molecule j in province m and quarter t. Note that these variables are aggregated across all drugs for each molecule. In addition, we examine the effects of the policy on the level of market concentration, using the Herfindahl–Hirschman index (HHI) and the number of firms selling the molecule as the outcome variables. $Enact_m$ is equal to 1 for enacting provinces and 0 for nonenacting provinces.

β_q captures how the outcome of interest moves differently in enacting provinces and nonenacting provinces q quarters before or after the policy is implemented. We normalize $\beta_{-1} = 0$ with the last quarter of 2018 as the baseline period. β_q with $q > 0$ captures the effect of the policy on the outcome variable. We include molecule-province fixed effects λ_{jm} to control for the time-invariant popularity of each molecule in each province, reflecting local preference heterogeneity. We also include molecule-quarter fixed effects λ_{jt} to control for molecule-specific time trends common to all provinces. Standard errors are clustered at the year-province level.

Three conditions are needed for the estimates to capture the causal effects of the policy on the outcome variable. First, we need $\beta_q = 0$ for $q < 0$. In other words, there should be no differential time trends for a molecule between enacting and nonenacting provinces prior to the policy experiment. This condition can be tested empirically. Second, there should be no shocks that are contemporaneous to the policy and that differently affect enacting and nonenacting provinces. A hypothetical example would be a disease outbreak in Beijing (an enacting province) in 2019. In a robustness check, we examine such potential confounders by using additional variation from the nonbidding molecules. The third condition is that bidding molecules in nonenacting provinces are not affected by the policy. This condition is the least likely to hold because the policy could have some spillover effects. For example, firms with losing bids might shift their focus to nonenacting provinces and behave differently there than they would have in a counterfactual world without the bidding. To the extent that such spillovers occur, our DID estimates pick up the difference between the policy effects in enacting and nonenacting provinces.

4.3 Effects of the Competitive Bidding Program on Market Outcomes

Figure 4 shows our main results from the DID specification in Equation (1). It compares drug prices and total quantities of bidding molecules at the provincial level between enacting and nonenacting provinces before and after 2019. Each dot represents the regression coefficient β_q for the corresponding quarter. For 2012 to 2017, we plot only the estimates for the first quarter to conserve space. Panel A shows that the average price drops by around 47.4% (58.5% after the first quarter) in enacting provinces relative to the price in nonenacting provinces. The estimates for years prior to 2019 are mostly precise zeros: although enacting provinces were not selected randomly, province-specific time trends are not a major threat to identification. Panel B shows that the policy has no appreciable effects on total quantities at the provincial level. With insurance coverage, most consumers can afford the drugs that they need, and aggregate drug demand is relatively inelastic to price changes. The competitive bidding program in large part results in substitutions between different drug products.

Appendix Figure A.4 shows the effects of the policy on measures of market concentration, where the outcomes are the HHI and the number of firms. The point estimates suggest that the policy slightly increases the HHI and reduces the number of firms by around 0.3. These results provide some tentative evidence that the policy may lead to firm exits and an increase in market concentration. We have data for only one year after the policy, and such impacts on market structure could be larger in the long run.

Our results are robust to a number of alternative specifications. First, we estimate Equation (1) with each observation weighted by the total sales quantity at the province-molecule level during 2012–2018. Appendix Figure A.5 shows that the results are almost identical to our baseline results. This confirms that our results are not driven by molecules or provinces with low sales.

Second, we define as enacting provinces only the four fully treated municipalities and drop the five partially treated provinces from the analysis. Appendix Figure A.6 shows the results of our baseline regressions on the basis of this new sample. We find a slightly larger reduction in drug prices (52.7% versus 47.4%). Including partially treated provinces may lead to an underestimate of the policy effects, but the difference is reasonably small. We consider our baseline definition of an enacting province a reasonable approximation.

Last, we estimate a triple-difference (DDD) specification where we incorporate variation from nonbidding molecules to control for province-level demand shocks that are common to all molecules (Gruber 1994). Appendix Figure A.7 shows that the results are again almost identical to our baseline results. The remaining threat to identification is contemporaneous demand shocks that are specific to bidding molecules and enacting provinces, which we cannot directly test but consider to be unlikely.

4.4 Heterogeneous Effects on Auction Winners and Losers

In Section 4.3, we present the aggregate effects of the competitive bidding program at the molecular level. In this section, we investigate the heterogeneous effects on auction winners versus losers. Auction winners must cut prices by at least 10%, as required by the policy. Auction losers, however, might adjust prices in either direction, as illustrated in Section 3. On the one hand, since prices are strategic complements, the winner's price cut might induce follow-up price reductions by the auction losers. This was part of the government's stated policy rationale of achieving price reductions "across the board." On the other hand, branded auction losers might increase their prices to target the remaining consumers with stronger brand preferences.

We examine price changes in enacting provinces separately for the auction winners and losers, using all firms in nonenacting provinces as the control group. Specifically, we run two versions of Equation (1), one in which we exclude all auction *losers* in enacting provinces when we construct the molecule-level prices and quantities and one in which we exclude all auction *winners*. Everything else remains the same. The estimates then correspond to the effects on auction winners and losers. Note that winners and losers are endogenously determined, so the results should be interpreted with caution.

Figure 5 shows the results. We find winners on average reduce prices by 59.4% and see an increase in sales of around 68.4%. In contrast, losing drugs show little price adjustment, and they cede market share to the auction winners. The competitive bidding program successfully achieves a steep price reduction by auction winners but falls short of inducing follow-up price reductions across the board.

To summarize, our descriptive analysis shows that the competitive bidding program signifi-

cantly reduces drug prices and induces large-scale substitutions from branded drugs to BE generic ones. While price reductions benefit consumers, (forced) generic substitutions under the quantity guarantee may reduce consumer welfare. Therefore, the impact of the program on consumer welfare is ex ante unclear. In the next section, we develop and estimate a structural model of drug demand and supply to quantify the welfare impact of the program.

5 Model and Estimation

In this section, we develop and estimate a structural model of drug demand and supply to quantify the welfare impact of the bidding program. We focus on the market for hypertension treatments since the set of bidding molecules includes seven major antihypertensive drugs.

To quantify the welfare effects of price reductions and consumer choice distortions, we need to disentangle the roles of price changes and the quantity guarantee in reshaping consumers' drug choices in 2019. To do so, we first use preprogram data to estimate a model of drug demand and supply. The estimated price elasticity allows us to infer the part of the sales change in 2019 that is driven by price reductions. Second, we attribute the remaining sales changes in enacting provinces relative to sales in nonenacting provinces to the quantity guarantee. We discuss each step in detail in this section.

5.1 Demand

We assume that each patient chooses drug product j to maximize utility under the supervision of physicians. A market is defined as province m in year t. The indirect utility of patient i in province m and year t from product j and molecule g follows a nested logit specification:

$$u_{ijmt} = \underbrace{\alpha\phi p_{jmt} + BE_{jt}\beta + \lambda_{jm} + \lambda_{mt} + \xi_{jmt}}_{\delta_{jmt}} + \zeta_{igmt} + (1-\sigma)\varepsilon_{ijmt}. \tag{2}$$

Here, ϕp_{jmt} is patient i's out-of-pocket expenditures and is equal to the listed price p_{jmt} multiplied by $\phi = 0.28$, or 1 minus the reimbursement rate. We define each molecule (or the outside option) as a nest. ζ_{igmt} is consumer i's preference shock that is common to all products of molecule g. Both ε_{ijmt} and $\zeta_{imt} + (1-\sigma)\varepsilon_{ijmt}$ follow a type-I extreme value distribution, and σ determines the degree of the within-nest correlation in preference shocks.

BE_{jt} takes the value of 1 if product j has passed the bio-equivalence test by year t. λ_{jm} and λ_{mt} are product-province and province-year fixed effects. ξ_{jmt} represents demand shocks observed by consumers but not by the econometrician. The outside option ($j = 0$) includes all other hypertension treatments not included in our sample and the option of not obtaining any treatment. The utility of the outside option is normalized to 0.

As shown in Berry (1994), our model implies the following demand equation:

$$\ln(s_{jmt}) - \ln(s_{0mt}) = \alpha\phi p_{jmt} + BE_{jt}\beta + \lambda_{jm} + \lambda_{mt} + \sigma \ln(s_{j/g,mt}) + \xi_{jmt}. \tag{3}$$

Here, $s_{j/g,mt}$ is the conditional market share of product j within molecule g. The unobserved demand shock ξ_{jmt} is likely to be correlated with both the price p_{jmt} and the conditional market share $\ln(s_{j/g,mt})$. We use instruments to consistently estimate α and σ as discussed in Section 5.3.

5.2 Supply

Firms set prices in each market to maximize profits under standard Bertrand–Nash competition. Let $\mathcal{J}_{ft}$ denote the set of products sold by firm f in year t. Let mc_{jmt} denote the marginal cost of product j in province m and year t. Then firm f's profit maximization problem in province m and year t is the following:

$$\max_{\{p_{jmt}\}_{j\in\mathcal{J}_{ft}}} \pi_{fmt}(\boldsymbol{p}) = \sum_{j\in\mathcal{J}_{ft}} s_{jmt}(p_{jmt} - mc_{jmt}).$$

Let J_{mt} be the total number of products in province m and year t, and let Δ_{mt} be a J_{mt}-by-J_{mt} matrix whose (k,r)-th term is $-\dfrac{\partial s_{rmt}}{\partial p_{kmt}}$ if products k and r are produced by the same firm and 0 otherwise. Let $\vec{s}_{mt}$, $\vec{p}_{mt}$, and $\vec{mc}_{mt}$ represent J_{mt}-by-1 vectors of market shares, prices, and marginal costs for all products in market mt. The first-order conditions in the firms' profit maximization problem can be written in vector notation as follows:

$$\vec{s}_{mt} - \Delta_{mt}(\vec{p}_{mt} - \vec{mc}_{mt}) = 0.$$

This implies

$$\vec{mc}_{mt} = \vec{p}_{mt} - \Delta_{mt}^{-1}\vec{s}_{mt}. \tag{4}$$

We derive elements of Δ_{mt} from the demand system and use Equation (4) to compute the marginal cost of each product.

5.3 Identification

We estimate Equation (3) with data for years prior to 2019 at the product-province-year level. To convert sales to market shares, we assume the market size to be 1.2 times the total quantity sold in each market. Since we include province-year fixed effects, our demand estimates stay the same under alternative assumptions about market size.

With product-province fixed effects, the identification of demand-side parameters relies on the time-series variation in drug prices and sales. Year-on-year changes in prices are likely to be correlated with unobserved demand shocks that firms observe before setting prices. The within-nest conditional market share is also endogenous, as any demand shock that increases a product's market share also increases its within-nest market share. Thus, OLS regressions overestimate both parameters. We use instruments to consistently estimate α and σ.

The first instrument is the average price of the focal product in all other provinces in the same year, which is the classic Hausman instrument (Hausman 1996). The identifying assumption is that conditional on the fixed effects, demand shocks are uncorrelated across markets, while common cost shocks pass through to prices in all markets. The second instrument, also known as the BLP instrument (Berry et al. 1995), is the total number of drugs for each molecule in each market. The identifying assumption is that product entry and exit are not correlated with contemporaneous unobserved demand shocks. This assumption is likely to be valid since approval of new drugs takes time. Our preferred specification is the nested logit model estimated with the two instruments. We also estimate a simple logit model without a nest structure.

It is important to note that in estimating demand, we do not leverage the policy variation that we document in our reduced-form analysis in Section 4. This is because while we could measure the effects of the competitive bidding program on both auction winners and losers, we cannot separate the effects of the price changes from those of the quantity guarantee. As a result, we take a two-step approach whereby we first use data for the period prior to the policy to estimate consumers' price elasticity and then recover the effects of the quantity guarantee from residual sales changes not explained by the price effects. We discuss how we quantify the effects of the quantity guarantee in more details in Section 5.5.

5.4 Estimation Results

Column (1) in Table 2 presents the results from an OLS estimation of a simple logit model. The estimate implies an average own-price elasticity of around 0.61. Column (2) presents the results of an instrumental variables (IV) regression with the Hausman instrument. The estimated average own-price elasticity is 1.46. Note that in the logit model, we do not allow drugs within the same molecule to substitute more strongly with each other than drugs across different molecules do, which could bias our estimates.

Therefore, we also employ the nested logit model shown in Columns (3) and (4) in Table 2, estimated using OLS and IV regressions, respectively. The estimated own-price elasticity is similar to the estimate under the simple logit model for both the OLS and IV regressions. In Column (4), the coefficient of 0.41 on conditional market share suggests that consumers are indeed more likely to substitute between different products of the same molecule than between products of different molecules. In addition, our estimates suggest that passing the BE test does not have an appreciable impact on drug sales, which reflects limited awareness or understanding of BE certification among consumers.

Appendix Table A.1 shows the results of the first-stage regressions. The first column summarizes the first stage of the logit model, and the last two columns summarize the nested logit model. Comparing the three columns, we can see that the Hausman instrument is highly predictive of the average price while the BLP instrument mainly helps identify the coefficient of the conditional market share. The F statistics are reasonably large.

The drug-province fixed effects represent province-specific drug preferences. We plot the distributions of the drug-province fixed effects separately for branded drugs, BE generic drugs, and other generic drugs in Figure 6. The branded drugs on average have higher values of fixed effects than both types of generic drugs. This is direct evidence of the existence of brand preferences: that is, consumers derive more utility from consuming branded drugs. On the other hand, the difference between generic drugs with and without a bio-equivalence certificate is not significant. Combined with the small and nonsignificant effect of passing the bio-equivalence test in Table 2, this implies that patients may have lacked the information needed to choose higher-quality generic drugs in the pre-auction period.

With the demand estimates at hand, we derive the price elasticity and recover the marginal costs of each product by means of Equation (4). Figure 7 presents the distribution of markups for branded and generic drugs separately. Compared with generic drugs, branded drugs sell with higher markups. In addition, we find that the distribution of markups for BE generic drugs is almost the same as that for non-BE generic ones.

5.5 Effects of the Quantity Guarantee

So far, we have estimated the price coefficient and the within-nest preference correlation based on data for the period prior to 2019. Under the assumption that these parameters remain the same in 2019, we can recover the effects of price changes after competitive bidding on drug sales in 2019. We estimate demand shifters at the product-province level that explain changes in sales in 2019 that are not driven by the price change.

Define consumer i's utility from product j in market m in 2019 as follows:

$$\tilde{u}_{ijm,2019} = \tilde{\delta}_{jm,2019} + \zeta_{igm,2019} + (1-\sigma)\varepsilon_{ijm,2019}. \tag{5}$$

Here, $\tilde{\delta}_{jm,2019} = \delta_{jm,2019} + \gamma_{jm,2019}$, and $\gamma_{jm,2019}$ is the demand shifter. We calculate $\delta_{jm,2019}$ based on the estimated values in 2018, adjusting for price changes in 2019. We then estimate $\gamma_{jm,2019}$ for each product-province pair by matching the observed market shares and the predicted market shares in 2019.[15]

Part of the demand shifter may capture changes in consumer preferences across time instead of impacts of the quantity guarantee. We use the average demand shifter in nonenacting provinces as an estimate for the time trend from 2018 to 2019 in consumers' preferences for each drug product:

$$\hat{\tau}_j = \frac{1}{|NE|}\sum_{k\in NE}\hat{\gamma}_{jk,2019}.$$

[15] One minor issue is that the set of products is not identical in 2018 and 2019 because products entered or exited some provinces. We exclude products that were new to some provinces in 2019, as we do not have marginal cost estimates for these products. We calculate the demand shifter for the set of product-province pairs in 2018. For a few losing drugs sold only in 2018 and not in 2019 in some provinces, we impute $\hat{\gamma}_{jm,2019}$ by means of the estimated demand shifters of the same product or products in the same molecule in other provinces. Appendix B provides more details on the imputation.

Here, NE is the set of nonenacting provinces. We then subtract it from the estimated $\hat{\gamma}_{jm,2019}$ in enacting provinces to recover the net policy demand shifter:

$$\tilde{\gamma}_{jm,2019} = \hat{\gamma}_{jm,2019} - \hat{\tau}_j.$$

We assume that absent the competitive bidding program, consumers' preferences for each drug product would have trended similarly in enacting and nonenacting provinces.

These policy demand shifters provide a reduced-form description of the effects of the quantity guarantee on consumer demand. Appendix Figure A.8 shows the distribution of the policy demand shifters. We see that the quantity guarantee significantly shifts consumer demand from losing products to auction winners. Such choice distortions may potentially reduce consumer surplus and offset the welfare gains from price cuts. In the next section, we quantify the welfare effects of the competitive bidding program.

6 The Welfare Effects of Competitive Bidding

In this section, we use the model described in Section 5 to quantify the welfare effects of the competitive bidding program in the market for hypertension treatments. The standard measure of consumer welfare, which is based on the revealed preference paradigm, does not capture the possibility that consumers' preferences for specific drug products may be driven by misinformation rather than by true utility. We first introduce an alternative welfare measure in the same spirit as the welfare analysis in Section 3 and then quantify the welfare effects of the competitive bidding program under the standard and alternative welfare measures.

6.1 Welfare Measures

We follow Train (2015) and distinguish between consumers' decision utility and actual utility. The former determines consumer choices, while the latter determines consumer welfare. In our setting, decision utility includes both the policy demand shifter (that is, the effects of the quantity guarantee) and consumers' brand preferences. Actual utility does not include the policy demand shifter and may or may not include brand preferences, depending on whether one deems brand preferences welfare relevant.

Formally, let $\delta'_{jm,2019}$ denote the intrinsic therapeutic value of product j in market m in 2019. For a non-BE generic product, we define $\delta'_{jm,2019}$ to be equal to the perceived value $\delta_{jm,2019}$. For a branded or BE generic product, we define $\delta'_{jm,2019}$ to be the average perceived value of all branded and BE generic products of the same molecule g in market m (excluding the price component $\alpha\phi p_{jmt}$). This definition relies on the fact that all BE generic products are certified to have the same therapeutic value as the branded product. Specifically,

$$\delta'_{jm,2019} = \alpha\phi p_{jm,2019} + \frac{1}{|\{k \in g, k \in BE_{m,2019}\}|} \sum_{k \in g, k \in BE_{m,2019}} (\delta_{km,2019} - \alpha\phi p_{km,2019}),$$

where $BE_{m,2019}$ is the set of branded and BE generic drugs available in province m in 2019.

Consumer i's decision utility in 2019 is as follows:

$$\tilde{u}^{D}_{ijm,2019} = \delta_{jm,2019} + \hat{\tau}_j + \tilde{\gamma}_{jm,2019} + \zeta_{igmt} + (1-\sigma)\varepsilon_{ijmt}. \tag{6}$$

Her actual utility is as follows:

$$\begin{cases} \tilde{u}^{A1}_{ijm,2019} = \delta_{jm,2019} + \hat{\tau}_j + \zeta_{igmt} + (1-\sigma)\varepsilon_{ijmt} & \text{if brand preference is relevant;} \\ \tilde{u}^{A0}_{ijm,2019} = \delta'_{jm,2019} + \hat{\tau}_j + \zeta_{igmt} + (1-\sigma)\varepsilon_{ijmt} & \text{if brand preference is irrelevant.} \end{cases} \tag{7}$$

Welfare loss occurs when consumers' choice utility is misaligned with their actual utility. We discuss additional details on the calculation of consumer surplus in Appendix C.

6.2 Welfare Analysis

We measure the effects of the competitive bidding program on consumer surplus, producer surplus, and government expenditures. To shed light on the trade-off between price reductions and potential choice distortions, we first examine two separate scenarios in which we consider effects from just the price changes or just the quantity guarantee. We then consider the full policy scenario, which combines the two effects. The policy scenarios are described as follows:

- **Baseline (status quo):** No policy intervention. Consumers' decision utility follows Equation (6) but without the $\tilde{\gamma}_{jm,2019}$ (quantity guarantee) term. All firms set prices freely in Bertrand–Nash competition.

- **Case 1 (price constraint only):** Consumers' decision utility follows equation (6) but without the $\tilde{\gamma}_{jm,2019}$ term. Auction winners set the observed negotiated prices. Other firms adjust prices strategically.

- **Case 2 (quantity guarantee only):** Consumers' decision utility follows Equation (6). All firms set prices freely in Bertrand–Nash competition. In particular, auction winners' prices are not constrained by the observed negotiated prices.

- **Case 3 (full policy):** Consumers' decision utility follows Equation (6). Auction winners' prices are set at the observed negotiated prices, and other producers can respond strategically. This scenario essentially replicates the actual policy setting.

Figure 8 summarizes the results on consumer surplus. As expected, price reductions increase consumer surplus.[16] The effects of the quantity guarantee depend on the nature of consumers' brand preferences. When consumers' brand preferences are welfare relevant, the quantity guarantee distorts consumer choices and reduces consumer welfare. When consumers' brand preferences are welfare irrelevant, the quantity guarantee improves consumer welfare by offsetting choice distortions caused by brand preferences. Under the full policy, when brand preferences are welfare relevant, the choice distortion effect looms large, and consumer welfare decreases by 9.8%. When brand preferences are welfare irrelevant, however, price reductions and generic substitutions both improve consumer welfare and add up to a 39.2% increase in consumer surplus.

Turning to firms and government spending, we find that the competitive bidding program reduces producer surplus by 24.4% and reduces government expenditures on insurance reimbursement by 24.3%, as shown in Figure 9. When these effects are taken together, total social surplus decreases by 14.5% under revealed preference but increases by 7.8% if policymakers consider brand preferences welfare irrelevant.

Overall, our analysis shows that the competitive bidding program was effective in reducing drug prices. The quantity guarantee, which is usually distortionary in ordinary market settings, may also improve consumer welfare by correcting choice distortions due to misjudged brand preferences. Competitive bidding is an appealing tool to reduce the prices of off-patent drugs and to

[16] Appendix Figure A.9 shows that price reductions by the auction winner also lead to a modest reduction on prices of other drug products, which contributes to the welfare gains.

guide consumers to switch toward generic drugs if policymakers deem consumers' brand preferences welfare irrelevant.

7 Conclusion

In this paper, we evaluated a novel policy experiment run by the Chinese government to limit the prices of off-patent drugs: competitive bidding with a quantity guarantee. We first built a simple model to predict the moving forces behind this competitive bidding initiative. Our model predicts that the generic drug firm always wins the bidding; however, the price changes are ambiguous ex ante.

Our empirical analysis yielded three main findings consistent with the model predictions. First, competitive bidding was effective in reducing drug prices: auction winners on average cut prices by 59.4%, and the average drug price fell by 47.4%. However, the losing firms did not cut prices significantly in response, and the policy fell short of achieving price reductions across the board. Second, the bidding program led to large-scale generic substitution. We built a simple model to quantify the welfare impacts for hypertensive drugs that were included in the bidding. Under the assumption that consumers' choices of branded drugs over cheaper BE generic ones are driven by misinformation or misperceptions, we found that this policy increased consumer surplus in China's pharmaceutical market by 39.2%. In addition, government expenditures decreased by 24.3%. Finally, the policy resulted in heavy losses for drug firms and has led firms that lost auctions to exit. Over time, such effects may increase market concentration and make competitive bidding less effective in limiting drug prices.

Overall, our results show that competitive bidding with a quantity guarantee is a useful addition to the toolkit of policy instruments for limiting drug prices. A unique advantage of competitive bidding over common policies such as cost-plus price regulations and reference pricing is that it leverages market forces to identify a price level that is still profitable for the winner. This ensures that the product remains available to consumers after the price restriction is imposed, and it therefore allows all patients to benefit from cheaper drugs. This could be a useful approach for policymakers to consider, especially in countries in which public hospitals and pharmacies constitute a large part of the pharmaceutical market. That said, policymakers should be mindful of its

negative impact on firms' surplus and its long-term effects on the market structure.

References

Atal, J., Cuesta, J. I., and Sæthre, M. (2019). Quality Regulation and Competition: Evidence from Pharmaceutical Markets. Working Paper, PIER.

Bairoliya, N., Karaca-Mandic, P., McCullough, J. S., and Petrin, A. (2017). Consumer Learning and the Entry of Generic Pharmaceuticals. Working Paper 23662, National Bureau of Economic Research.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841–890.

Berry, S. T. (1994). Estimating Discrete-Choice Models of Product Differentiation. *RAND Journal of Economics*, 25(2):242–262.

Branstetter, L., Chatterjee, C., and Higgins, M. J. (2016). Regulation and Welfare: Evidence from Paragraph IV Generic Entry in the Pharmaceutical Industry. *RAND Journal of Economics*, 47(4):857–890.

Bronnenberg, B. J., Dubé, J.-P., Gentzkow, M., and Shapiro, J. M. (2015). Do Pharmacists Buy Bayer? Informed Shoppers and the Brand Premium. *Quarterly Journal of Economics*, 130(4):1669–1726.

Cabral, M., Geruso, M., and Mahoney, N. (2018). Do Larger Health Insurance Subsidies Benefit Patients or Producers? Evidence from Medicare Advantage. *American Economic Review*, 108(8):2048–87.

Caves, R. E., Whinston, M. D., Hurwitz, M. A., Pakes, A., and Temin, P. (1991). Patent Expiration, Entry, and Competition in the US Pharmaceutical Industry. *Brookings Papers on Economic Activity. Microeconomics*, 1991:1–66.

Colgan, S., Faasse, K., Martin, L. R., Stephens, M. H., Grey, A., and Petrie, K. J. (2015). Perceptions of Generic Medication in the General Population, Doctors and Pharmacists: A Systematic Review. *BMJ Open*, 5(12).

Cuddy, E. (2020). Competition and Collusion in the Generic Drug Market. Working Paper.

Curto, V., Einav, L., Levin, J., and Bhattacharya, J. (2021). Can Health Insurance Competition Work? Evidence from Medicare Advantage. *Journal of Political Economy*, 129(2):570–606.

Dean, E. B. (2019). Who benefits from pharmaceutical price controls?: Evidence from india. Technical report, Center for Global Development.

Ding, H., Duggan, M., and Starc, A. (2021). Getting the Price Right? The Impact of Competitive Bidding in the Medicare Program. Working Paper 28457, National Bureau of Economic Research.

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

Dubois, P., Lefouili, Y., and Straub, S. (2021). Pooled Procurement of Drugs in Low and Middle Income Countries. *European Economic Review*, 132:103655.

Duggan, M., Starc, A., and Vabson, B. (2016). Who Benefits When the Government Pays More? Pass-through in the Medicare Advantage Program. *Journal of Public Economics*, 141:50–67.

Fang, H., Lei, X., Shi, J., and Yi, X. (2021). Physician-Induced Demand: Evidence from China's Drug Price Zero-Markup Policy. Working Paper 28998, National Bureau of Economic Research.

Frank, R. G. and Salkever, D. S. (1997). Generic Entry and the Pricing of Pharmaceuticals. *Journal of Economics & Management Strategy*, 6(1):75–90.

Grabowski, H. G. and Vernon, J. M. (1992). Brand Loyalty, Entry, and Price Competition in Pharmaceuticals After the 1984 Drug Act. *Journal of Law and Economics*, 35(2):331–350.

Gruber, J. (1994). The incidence of mandated maternity benefits. *The American economic review*, pages 622–641.

Hausman, J. A. (1996). Valuation of New Goods Under Perfect and Imperfect Competition. In *The Economics of New Goods, Studies in Income and Wealth*, volume 58. University of Chicago Press.

Huckfeldt, P. J. and Knittel, C. R. (2011). Pharmaceutical Use Following Generic Entry: Paying Less and Buying Less. Working Paper 17046, National Bureau of Economic Research.

Ji, Y. (2019). The Impact of Competitive Bidding in Health Care: The Case of Medicare Durable Medical Equipment. Working Paper.

Klemperer, P. (1999). Auction Theory: A Guide to the Literature. *Journal of Economic Surveys*, 13(3):227–286.

Mohapatra, D. and Chatterjee, C. (2016). Price Control and Access to Drugs: The Case of India's Malaria Market. Working paper.

Mossialos, E., Ge, Y., Hu, J., and Wang, L. (2016). *Pharmaceutical Policy in China: Challenges and Opportunities for Reform*. World Health Organization.

OECD (2018). Pharmaceutical Expenditure. *Health at a Glance 2018: OECD Indicators*.

OECD (2019). Pharmaceutical Expenditure. *Health at a Glance 2019: OECD Indicators*.

Saha, A., Grabowski, H., Birnbaum, H., Greenberg, P., and Bizan, O. (2006). Generic Competition in the US Pharmaceutical Industry. *International Journal of the Economics of Business*, 13(1):15–38.

Song, Y. and Barthold, D. (2018). The Effects of State-Level Pharmacist Regulations on Generic Substitution of Prescription Drugs. *Health Economics*, 27(11):1717–1737.

Song, Z., Landrum, M. B., and Chernew, M. E. (2013). Competitive Bidding in Medicare Advantage: Effect of Benchmark Changes on Plan Bids. *Journal of Health Economics*, 32(6):1301–1312.

Train, K. (2015). Welfare Calculations in Discrete Choice Models When Anticipated and Experienced Attributes Differ: A Guide with Examples. *Journal of Choice Modelling*, 16:15–22.

Yang, Z. and Fan, D. (2012). How to Solve the Crisis Behind Bribegate for Chinese Doctors. *The Lancet*, 379(9812):e13–e15.

Zhang, Y., Yu, Y.-S., Tang, Z.-H., Chen, X.-H., and Zang, G.-Q. (2014). Crack Down on Medical Corruption: An Urgent Matter in China. *European Journal of Internal Medicine*, 25(1):e2–e3.

Table 1: List of Molecules and Summary Statistics

	Disease Category	Revenues (Million)	Price (RMB)	Sales (Million)	Brand–BE Price Ratio	# Firms	# BE Firms
Bidding Molecules							
Clopidogrel bisulfate	Arteriosclerosis	561.2	4.6	124.9	1.6	3	3
Montelukast sodium	Asthma	108.2	9.2	12.0	2.4	5	2
Gefitinib	Cancer	102.0	391.0	0.3	2.2	3	2
Imatinib mesylate	Cancer	152.0	122.2	1.7	9.5	4	2
Pemetrexed disodium	Cancer	320.5	2275.4	0.1	3.6	16	2
Atorvastatin calcium	Cholesterol	462.3	8.3	56.5	1.3	8	3
Rosuvastatin calcium	Cholesterol	218.1	7.1	31.7	1.4	7	6
Diosmectite	Diarrhea	10.3	1.6	6.6	1.8	20	5
Levetiracetam	Epilepsy	58.2	8.3	7.1	1.5	6	2
Tenofovir	HIV	30.6	35.1	1.6	1.7	7	4
Entecavir	Hepatitis B	412.4	22.2	20.4	1.8	10	5
Amlodipine besylate	Hypertension	184.3	3.4	53.5	3.4	47	6
Enalapril maleate	Hypertension	4.1	1.0	4.2	1.2	16	2
Fosinopril sodium	Hypertension	16.0	2.8	5.7	1.4	3	2
Irbesartan	Hypertension	98.0	3.6	27.2	2.6	17	4
Irbesartan–hydrochlorothiazide	Hypertension	62.1	3.8	16.6	1.6	7	3
Lisinopril	Hypertension	0.9	1.4	0.6	1.4	14	2
Losartan potassium	Hypertension	89.5	4.1	22.1	1.6	11	2
Cefuroxime axetil	Infection	31.9	2.5	12.5	1.4	11	4
Escitalopram	Mental Illness	74.4	11.9	6.5	1.7	6	4
Olanzapine	Mental Illness	147.8	10.6	14.3	2.8	7	2
Paroxetine HCL	Mental Illness	60.7	6.4	9.6	2.2	6	2
Risperidone	Mental Illness	30.3	1.9	16.1	2.3	10	2
Dexmedetomidine HCL	Painkiller	150.2	158.7	1.0	1.3	6	2
Flurbiprofen axetil	Painkiller	181.9	68.1	2.7	1.1	3	2
Nonbidding Molecules							
Captopril	Hypertension	1.9	0.5	3.9	19.5	49	3
Amoxicillin	Infection	8.9	0.4	22.9	2.9	59	4
Azithromycin, injection	Infection	61.0	35.4	1.7	1.2	44	2
Azithromycin, tablet	Infection	23.6	6.4	3.7	5.0	85	2
Alfacalcidol	Osteoporosis	41.0	1.7	24.5	1.8	13	2
Tramadol HCL	Painkiller	10.1	3.8	2.7	1.8	24	2

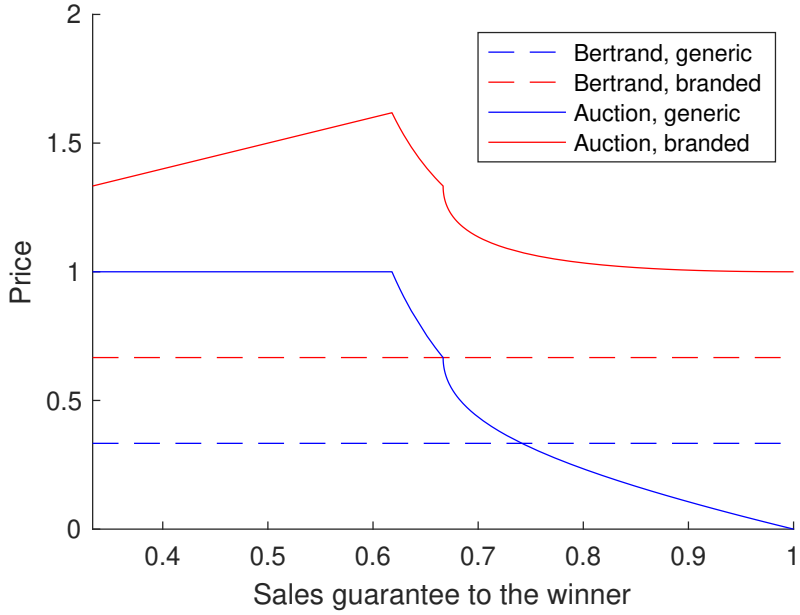
Notes: This table shows the 25 bidding and 6 nonbidding molecules in our sample. Quantities for each molecule are measured in the unit of the same dosage for the molecule (e.g., 10 mg lisinopril tablets). The brand–BE price ratio measures the ratio between the price of the branded drug and the average price of generic drugs from the same molecule. The number of firms lists the number of active firms selling the molecule in 2018. The number of BE firms includes both generic drugs that have passed the bio-equivalence test by 2018 and the branded drug.

Table 2: Demand Estimation

	(1) Logit OLS	(2) Logit IV	(3) NLogit OLS	(4) NLogit IV
Price	-0.587	-1.414	-0.681	-1.362
	(0.230)	(0.486)	(0.0856)	(0.388)
Passed BE	0.175	0.141	-0.0210	0.0601
	(0.0733)	(0.0756)	(0.0219)	(0.0618)
Cond. Mkt Share			0.939	0.412
			(0.0176)	(0.188)
N	3205	3191	3205	3191
Province-Drug FE	Yes	Yes	Yes	Yes
Province-Year FE	Yes	Yes	Yes	Yes

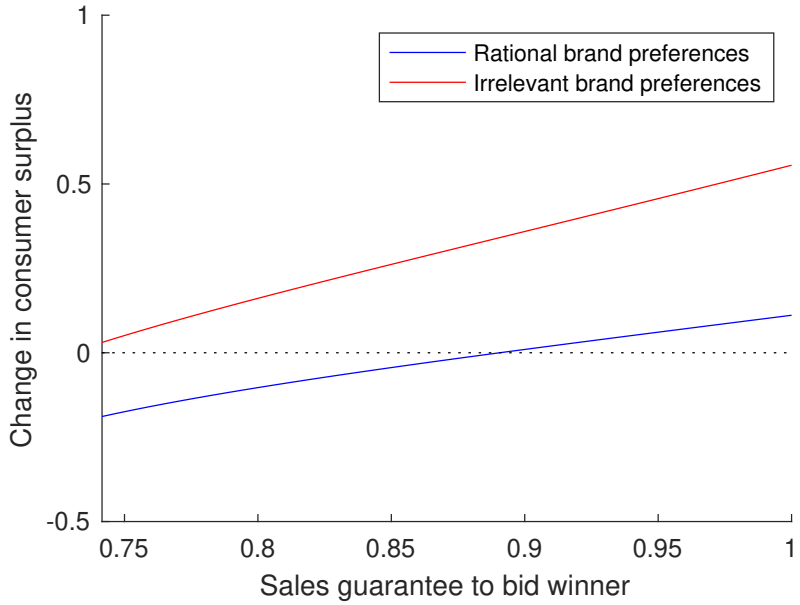
Notes: This table shows the estimation results of the demand system in Equation (3). Price is the out-of-pocket price paid by patients. Columns (1)–(2) use the logit specification, and Columns (3)–(4) use the nested logit specification. We use data for 2012–2018, the period prior to the competitive bidding program. We include province-drug and province-year fixed effects. Columns (1) and (3) show the OLS estimates with no instruments used. Columns (2) and (4) show the IV estimates based on the instruments described in Section 5.3. Appendix Table A.1 shows the first-stage regressions. Standard errors are clustered at the province-year level and are reported in parentheses.

Figure 1: Equilibrium Prices under Bertrand Competition and Competitive Bidding



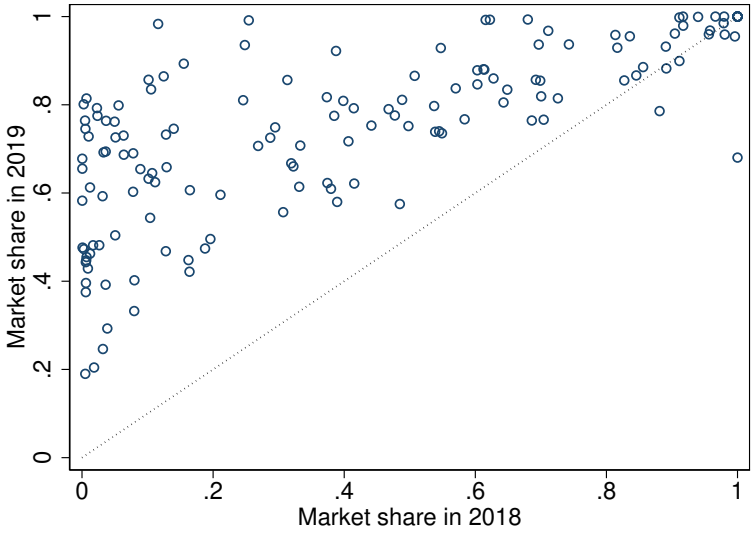
Notes: This figure shows equilibrium prices under both Bertrand competition and competitive bidding as in the model from Section 3. See Appendix A for the derivations.

Figure 2: Changes in Consumer Surplus from Bertrand Competition to Competitive Bidding



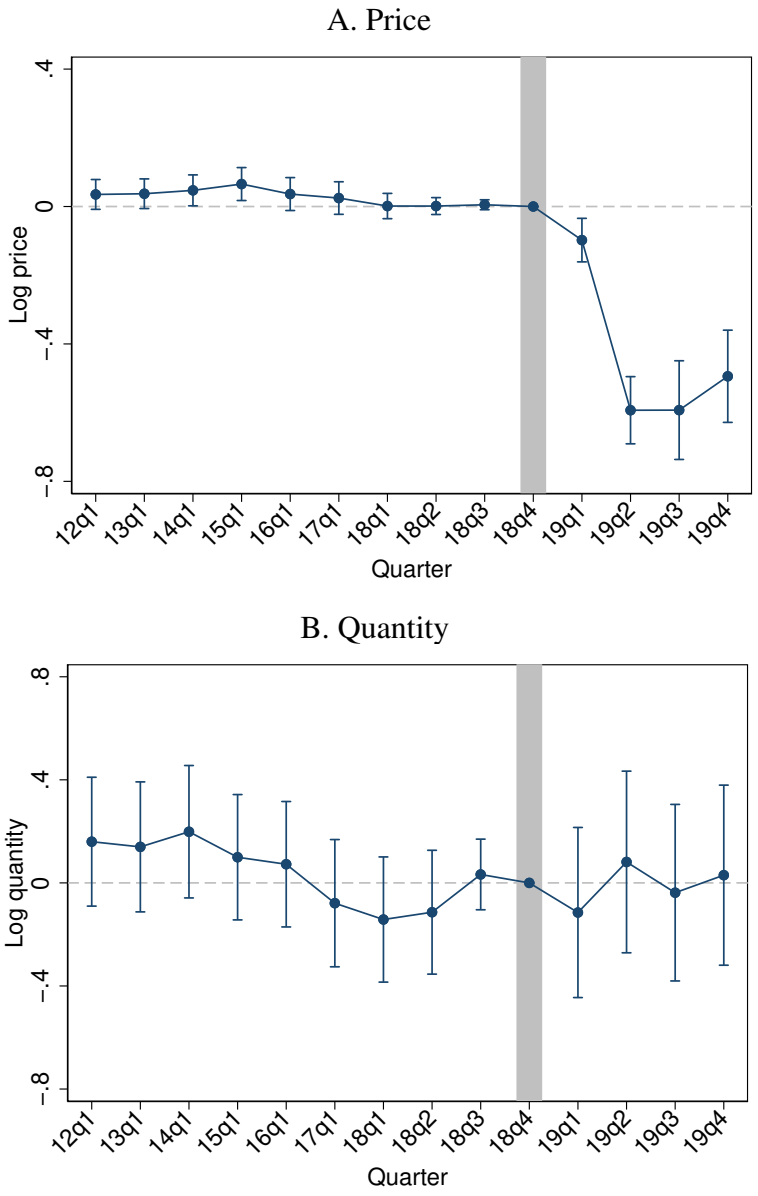
Notes: This figure shows the changes in consumer surplus with the move from Bertrand competition to competitive bidding as in the model from Section 3. See Appendix A for the derivations.

Figure 3: Market Shares of Auction Winners in Enacting Provinces



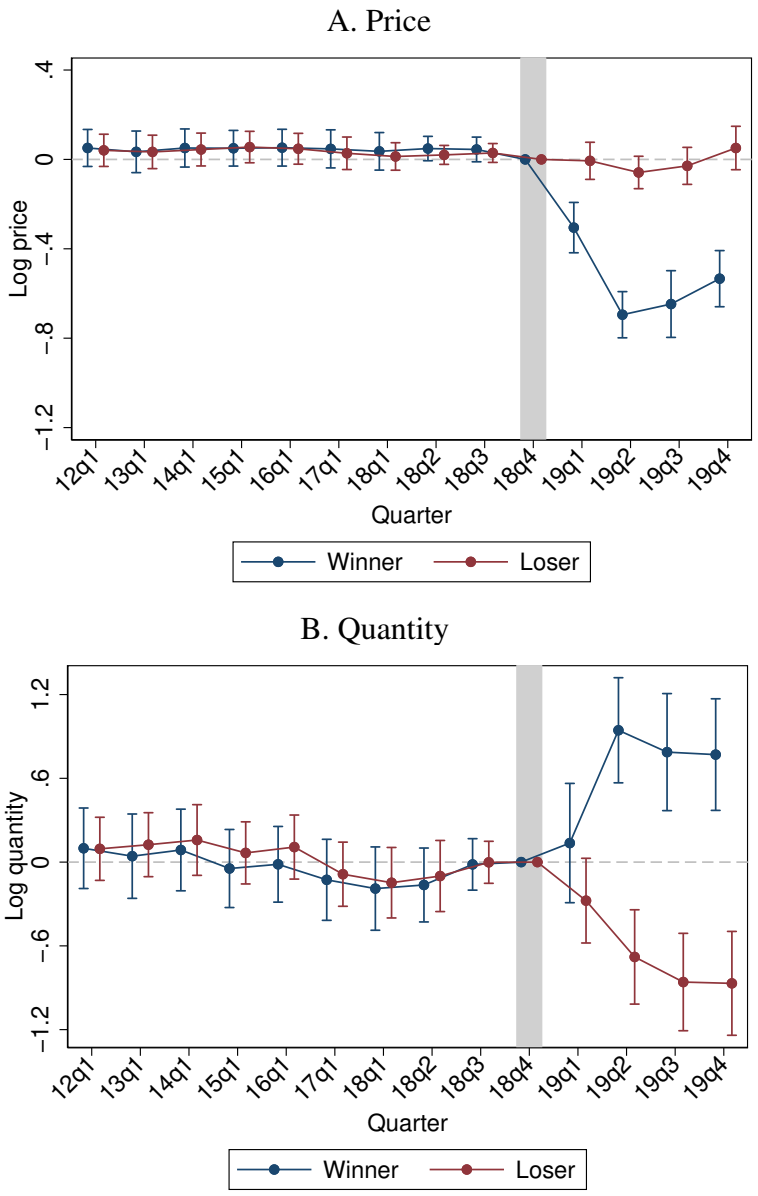
Notes: This figure plots the market shares for auction winners in enacting provinces for the periods prior to (2018) and after (2019) the competitive bidding program. Since the quantity guarantee of the competitive bidding went into effect in March 2019, we drop the first quarter when computing the 2019 market shares. Each dot is a molecule-province pair for all bidding molecules and enacting provinces.

Figure 4: Policy Effects on Drug Prices and Sales



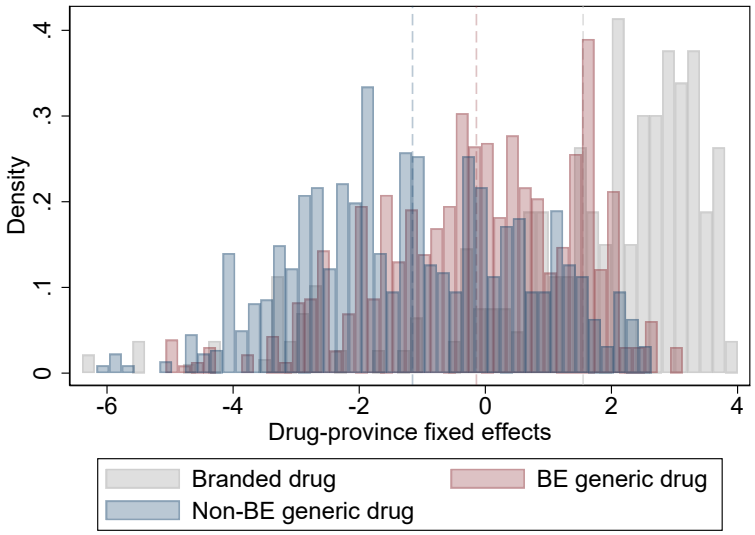
Notes: This figure shows the effects of competitive bidding on drug prices and sales, estimated using the specification shown in Equation (1). Each dot represents the regression coefficient for the corresponding quarter, and each line segment represents the 95% confidence interval with standard errors clustered at the province-year level. The coefficient of the fourth quarter in 2018 is normalized to zero. For 2012 to 2017, we show only the coefficients for the first quarter of each year to conserve space.

Figure 5: Heterogeneous Effects on Auction Winners and Losers



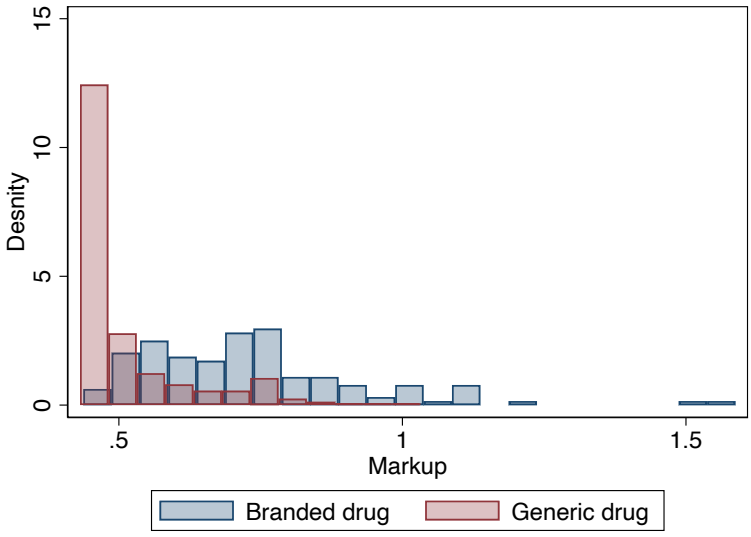
Notes: Similarly to Figure 4, this figure pools the results from two regressions where we estimate the effects of competitive bidding on drug prices and sales based on the specification shown in Equation (1) separately for firms that won (in blue) and lost (in red) the auction. Each dot represents the regression coefficient for the corresponding quarter, and each line segment represents the 95% confidence interval with standard errors clustered at the province-year level. The coefficient of the fourth quarter in 2018 is normalized to zero. From 2012 to 2017, we show only the coefficients for the first quarter of each year to conserve space.

Figure 6: Illustration of Brand Preferences



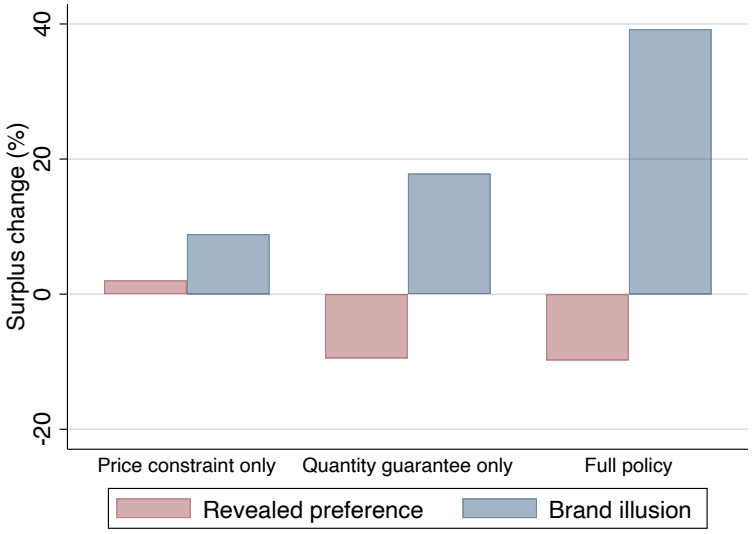
Notes: This figure illustrates brand preferences. We treat the estimated drug-province fixed effects in Equation (1) as province-specific drug preferences and plot the distribution of these fixed effects for branded drugs, BE generic drugs, and non-BE generic drugs separately. The dash lines represent the average drug-province fixed effects for the three types of products.

Figure 7: Markup for Branded and Generic Drugs



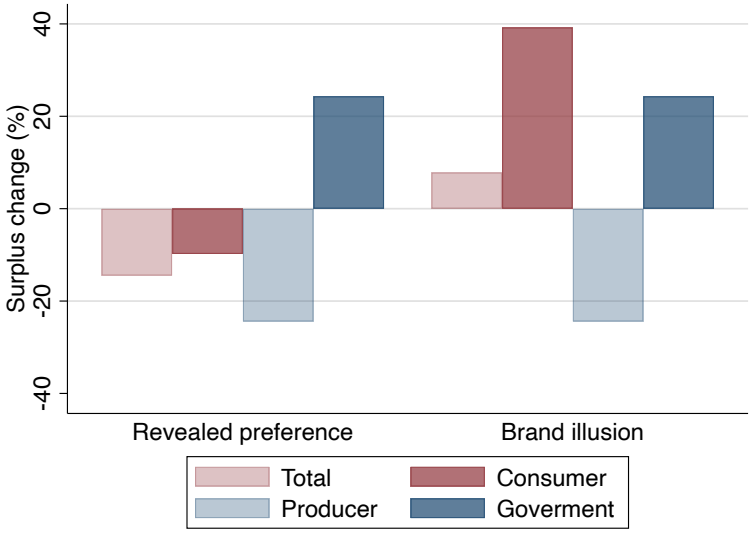
Notes: This figure plots the distribution of markups in 2018 for branded and generic drugs separately. The markups are calculated from Equation (4) as the differences between prices and marginal costs.

Figure 8: Changes in Consumer Surplus



Notes: This figure shows the change in consumer surplus under three counterfactual scenarios, as in Figure A.9. We measure consumer surplus in two ways. In red bars, we measure consumer surplus by using consumers' revealed preferences. In blue bars, we assume that brand preference is misjudged and welfare irrelevant, as discussed in Section 6.1.

Figure 9: Changes in Total Surplus



Notes: This figure shows the change in total surplus from full competitive bidding and decomposes the change into changes in consumer, producer, and government surplus. The left four bars show the results if we measure consumer surplus based on consumers' revealed preferences. The right four bars show the results if we assume that brand preference is misjudged and welfare irrelevant. See Section 6.2 for more details.

Online Appendix, Not for Publication

A Proofs and Derivations

In this section, we provide details of the proofs and derivations for the stylized model outlined in Section 3.

Consider a molecule with one firm selling a branded product B and another firm selling a BE generic product G. Both products offer an identical clinical value of 1 to a unit mass of consumers. Consumers have heterogeneous brand preferences for product B that are drawn independently from a continuously differentiable distribution $\mathscr{F}$ with non-negative support. We further assume $\mathscr{F}$ has no mass at 0. Marginal costs of both products are the same at c. Prior to the competitive bidding, firms set prices to maximize profits simultaneously in a Bertrand game with differentiated products. We use p_B^* and p_G^* to denote equilibrium prices under the Bertrand competition.

The government introduces a competitive bidding with sales guarantee $\tau \in (0,1]$. The format is first-price sealed-bid auction, and we assume firms have complete information on $\mathscr{F}$ and marginal costs. The winner commits to selling at the winning bid, and is guaranteed a quantity of τ, or all consumers who value the winning product no less than the winning bid if there are fewer than τ such consumers. Then the losing firm sets a price to maximize profits taking the winner's price and quantity guarantee as given. If it sets a price such that its unconstrained quantity exceeds $1-\tau$, the losing firm will only sell to $1-\tau$ of consumers who value its product (relatively) the most. We use p_B' and p_G' to denote equilibrium prices under the competitive bidding.

A.1 The Generic Drug Always Wins

We prove this result for a general $\mathscr{F}$ with non-negative support by contradiction. Let F denote its cumulative distribution function. Suppose there exists an equilibrium where the branded drug submits a bid $p_1 < 1$ and the generic drug submits a bid $p_1' \geq p_1$, and the branded drug wins the auction. The generic drug chooses a price p_2 to maximize profits after losing the auction. We must have $p_2 < p_1$, otherwise, the generic drug will have a market share of 0.

The generic drug's profit is $(p_2 - c)\min\{1-\tau, F(p_1 - p_2)\}$. Note that the generic drug always chooses p_2 such that $F(p_1 - p_2) \leq 1-\tau$ because its market share is at most $1-\tau$. Its profit thus simplifies to $(p_2 - c)F(p_1 - p_2)$, subject to $F(p_1 - p_2) \leq 1-\tau$.

Now consider a unilateral deviation by the generic drug: the generic drug cuts its bid to $p_1 - \varepsilon >$

0 and wins the auction. Let $p^B(\varepsilon)$ represent the best response by the branded drug. We must have $p^B(\varepsilon) > p_1 - \varepsilon$. Otherwise, the branded drug will have an unconstrained market share of 1. That contradicts to the fact that the loser can only receive a maximum quantity of $1-\alpha$ due to the quantity guarantee. Therefore, the branded drug could be better off by increasing the price, thus it should always choose $p^B(\varepsilon)$ such that $1 - F(p^B(\varepsilon) - p_1 + \varepsilon) \leq 1 - \tau$. The generic drug's profit under this deviation is $(p_1 - \varepsilon - c)\max\{\tau, F(p^B(\varepsilon) - p_1 + \varepsilon)\}$, which can be simplified to $(p_1 - \varepsilon - c)F(p^B(\varepsilon) - p_1 + \varepsilon)$ subject to $F(p^B(\varepsilon) - p_1 + \varepsilon) \geq \tau$, or $p^B(\varepsilon) \geq p_1 + F^{-1}(\tau) - \varepsilon$.

For this deviation to be unprofitable, we need:

$$\underbrace{(p_2 - c)F(p_1 - p_2)}_{\text{Generic losing the auction}} \geq \underbrace{(p_1 - \varepsilon - c)F(p^B(\varepsilon) - p_1 + \varepsilon)}_{\text{Generic winning the auction}}. \tag{A.1}$$

By continuity, this inequality holds for $\varepsilon \to 0^+$, which is

$$(p_2 - c)F(p_1 - p_2) \geq (p_1 - c)F(p_3 - p_1), \tag{A.2}$$

where $p_3 = p^B(0) \geq p_1 + F^{-1}(\tau) > p_1$. The fact that $F^{-1}(\tau) > 0$ for any $\tau > 0$ comes from the assumption that $\mathscr{F}$ has no mass at 0.

Next, consider a unilateral deviation by the branded drug: the branded drug increases its bid above p_1' and loses the auction. Let p_3' represent its optimal price after losing the auction. Similarly, for this deviation to be unprofitable, we need:

$$\begin{aligned}
\underbrace{(p_1 - c)(1 - F(p_1 - p_2))}_{\text{Branded winning the auction}} &\geq \underbrace{(p_3' - c)(1 - F(p_3' - p_1'))}_{\text{Branded losing the auction}} \\
&\geq (p_3 - c)\min\{1 - \tau, 1 - F(p_3 - p_1')\} \\
&\geq (p_3 - c)\min\{1 - \tau, 1 - F(p_3 - p_1)\} \\
&= (p_3 - c)(1 - F(p_3 - p_1)).
\end{aligned} \tag{A.3}$$

The second line comes from the fact that p_3' is a best response to p_1' and weakly more profitable than p_3. The third line comes from $p_1' \geq p_1$. The last line comes from $F(p_3 - p_1) \geq \alpha$.

As $p_2 < p_1$, inequality A.2 implies $F(p_1 - p_2) \geq F(p_3 - p_1)$. As $p_3 > p_1$, inequality A.3 implies $F(p_3 - p_1) \geq F(p_1 - p_2)$. It follows that $F(p_3 - p_1) = F(p_1 - p_2) > 0$. Then equation A.2 holds only if $p_2 \geq p_1$, which contradicts to the fact that $p_2 < p_1$. So there does not exist an equilibrium where the branded drug wins the auction.

A.2 Derivations of Equilibrium Prices

To analytically solve the equilibrium prices, we assume $\mathscr{F}$ follows a uniform distribution $\mathbb{U}(0,1)$. For simplicity, we also assume the marginal costs of both firms are 0. Then firms' profits under the Bertrand competition are given by:

$$\begin{aligned} \pi_B &= p_B(1-p_B+p_G) \\ \pi_G &= p_G(p_G-p_B) \end{aligned}$$

The best response functions are $p_B = \dfrac{1+p_G}{2}$, and $p_G = \dfrac{p_B}{2}$. Equilibrium prices are $p_B^* = \dfrac{2}{3}$, and $p_G^* = \dfrac{1}{3}$. Equilibrium market shares are $s_B^* = \dfrac{2}{3}$, and $s_G^* = \dfrac{1}{3}$.

Now consider a competitive bidding with a quantity guarantee τ between $\dfrac{1}{3}$ and 1. From Appendix A.1, the generic product always wins the auction. We solve the equilibrium prices.

Let $p'_G \leq 1$ denote the winning bid of the generic drug G. When $p'_G < 1$, in equilibrium, this bid must make the firm B exactly indifferent between undercutting and conceding to the bid. Otherwise, if the bid is higher, firm B has the incentive to undercut; if the bid is lower, firm G has the incentive to raise the bid and still wins the auction. When $p'_G = 1$, this bid must make firm B weakly prefer conceding to the bid than undercutting, because firm G cannot further increase the bid, otherwise, it will receive zero market share despite the guarantee, as no consumers value the product more than 1.

We hope to write p'_G as a function of τ and proceed case by case. When $p'_G < 1$, the indifference condition is

$$p'_G \times \max\{1 - \frac{p'_G}{2}, \tau\} = (1-\tau)(p'_G + \tau) \tag{A.4}$$

The left-hand side of Equation (A.4) is the profit of firm B from undercutting the bid by $\varepsilon \to 0^+$. When it undercuts the bid, firm B will win the auction. In response, firm G will either set a price of $1-\tau$ so that it has a market share of exactly $1-\tau$, or a price of $\dfrac{p'_G}{2}$ and has a market share of $\dfrac{p'_G}{2}$. It will choose the former if $1-\tau \geq \dfrac{p'_G}{2}$ and the latter otherwise.

The right-hand side of Equation (A.4) is the profit of firm B from conceding to the bid. It's easy to see that firm G's best response is to target the remaining consumers, set a price of $p'_G + \tau$, and captures all the remaining $1-\tau$ consumers.

Case I: $1-\frac{p'_G}{2}<\tau$. Equation (A.4) reduces to $p'_G\tau=(1-\tau)(p'_G+\tau)$, which gives us

$$p'_G=\frac{\tau(1-\tau)}{2\tau-1}.$$

For this to hold, we need $2(1-\tau)<p'_G<1$, or $\tau\in(\frac{\sqrt{5}-1}{2},\frac{2}{3})$.

Case II: $1-\frac{p'_G}{2}\geq\tau$. Equation (A.4) reduces to $p'_G(1-\frac{p'_G}{2})=(1-\tau)(p'_G+\tau)$, which gives us

$$p'_G=\tau-\sqrt{3\tau^2-2\tau}.$$

For this to hold, we need $3\tau^2-2\tau\geq 0$ and $p'_G\leq 2(1-\tau)$, or $\tau\in[\frac{2}{3},1]$.

Case III: $p'_G=1$. Firm B must weakly prefer conceding to the bid than undercutting. Similarly, we need

$$1\times\max\{0.5,\tau\}\leq(1-\tau)(1+\tau).$$

It follows that $\tau\leq\frac{\sqrt{5}-1}{2}$. We also need to impose the incentive compatibility constraint, i.e., firm G must be weakly better off participating in the auction than the Bertrand competition. Its profit under the Bertrand competition is $\frac{1}{3}$, so this implies $\tau\geq\frac{1}{3}$. No equilibrium would exist if $\tau<\frac{1}{3}$.

Collecting the cases, we have:

$$p'_G=\begin{cases}1, & \frac{1}{3}\leq\tau<\frac{\sqrt{5}-1}{2}\\ \frac{\tau(1-\tau)}{2\tau-1}, & \frac{\sqrt{5}-1}{2}\leq\tau<\frac{2}{3}\\ \tau-\sqrt{3\tau^2-2\tau}, & \frac{2}{3}\leq\tau\leq 1\end{cases}.$$

The branded product always sets $p'_B=p'_G+\tau$ to target the remaining $1-\tau$ consumers:

$$p'_B=\begin{cases}1+\tau, & \frac{1}{3}\leq\tau<\frac{\sqrt{5}-1}{2}\\ \frac{\tau^2}{2\tau-1}, & \frac{\sqrt{5}-1}{2}\leq\tau<\frac{2}{3}\\ 2\tau-\sqrt{3\tau^2-2\tau}, & \frac{2}{3}\leq\tau\leq 1\end{cases}.$$

The results are plotted in Figure 1.

B Imputation of Policy Demand Shifters

As in Section 5.5, the estimated demand shifter for drug j in province m in 2019 is denoted as $\hat{\gamma}_{jm,2019}$ in Equation (5). Since drugs might enter or exit the market across years, our sample is not a balanced panel of product-year-province combinations. As discussed in Section 5.5, we exclude products that were new to some provinces in 2019, as we do not have marginal cost estimates for these products. For a few losing drugs that were only sold in 2018 but not in 2019 in some provinces, we impute $\hat{\gamma}_{jm,2019}$ using the estimated demand shifters of the same product or products in the same molecule in other provinces. There are a total of 38 such observations.

More specifically, in the first step, if the product is missing in an enacting province in 2019, we impute its policy demand shifter using the average estimated policy effect of the same product across all other enacting provinces. If it is missing in an non-enacting province, we use the average estimated policy effect of the same product across all other non-enacting provinces.

$$\hat{\gamma}_{jm\in M,2019} = \frac{\sum_{k\in M, j\in \mathscr{J}_{k,2019}} \hat{\gamma}_{jk,2019}}{|\{k : k \in M, j \in \mathscr{J}_{k,2019}|\}}, \text{ where } M \in \{E, NE\},$$

where $\mathscr{J}_{m,2019}$ is the set of available drugs in province m in 2019, E is the set of enacting provinces, and NE is the set of non-enacting provinces. In this way, we can impute the policy demand shifters for 32 out of the 38 product-province observations.

We impute the rest 6 drug-province pairs using the average estimated policy demand shifters across all products of the same molecule, excluding the auction winner.[17] Specifically,

$$\hat{\gamma}_{j\in g, m\in M,2019} = \frac{\sum_{k\in M} \sum_{l\in g, l\in \mathscr{J}^{loser}_{k,2019}} \hat{\gamma}_{lk,2019}}{\sum_{k\in M} \{|l : l \in g, l \in \mathscr{J}^{loser}_{k,2019}\}|}, \text{ where } M \in \{E, NE\},$$

where g is the molecule that drug j belongs to and $\mathscr{J}^{loser}_{m,2019}$ is the set of available drugs in province m in 2019, excluding the auction winner if the province is an enacting province.

[17] All auction winners are not missing in 2019. So these products are all losing products or products in the non-enacting provinces.

C Welfare when Decision Utility and Actual Utility Differ

As discussed in Sections 5.5 and 6.1, in our setting, a consumer's decision utility can differ from her actual utility in two ways. First, under the quantity guarantee, consumers' decision utility for drug j is shifted by $\tilde{\gamma}_{jm,2019}$. Second, consumers' brand preferences could be a misperception and welfare irrelevant, which might only contribute to decision utility but not actual utility. In this section, we describe the calculation of consumer surplus when the decision utility differs from the actual utility.

Our method builds on Train (2015)'s example where the consumers' anticipated and experienced attributes differ. Suppose the decision utility is given by u_{ijmt} and the actual utility is given by $\tilde{u}_{ijmt} = u_{ijmt} + \Delta u_{jmt}$. The true consumer surplus is given by

$$CS = \frac{1}{\alpha} E(\tilde{u}_{ijmt}) = \frac{1}{\alpha} E(u_{ijmt}) + \frac{1}{\alpha} \sum s_{jmt} \times \Delta u_{jmt}, \tag{C.5}$$

where s_{jmt} is the market share of drug j chosen under the decision utility u_{ijmt}.

Now we apply this method to different utility functions in Section 6.1. First, we consider the case without the policy demand shifters, but we assume consumers' brand preferences are welfare irrelevant, thus does not affect the actual utility. Consumer i's decision utility is given by:

$$\tilde{u}^{A1}_{ijm,2019} = \delta_{jm,2019} + \hat{\tau}_j + \zeta_{igmt} + (1-\sigma)\varepsilon_{ijmt},$$

while her actual utility is

$$\tilde{u}^{A0}_{ijm,2019} = \delta'_{jm,2019} + \hat{\tau}_j + \zeta_{igmt} + (1-\sigma)\varepsilon_{ijmt}$$

The consumer surplus is given by

$$CS = \frac{1}{\alpha} E(\tilde{u}^{A1}_{ijm,2019}) + \frac{1}{\alpha} \sum s_{jmt} \times (\delta'_{jm,2019} - \delta_{jm,2019}),$$

where s_{jmt} is the market share of drug j chosen under the choice utility $\tilde{u}^{A1}_{ijm,2019}$.

Second, we consider the case with the policy demand shifters, but assume consumers' brand preferences are welfare relevant. Consumer i's decision utility is given by:

$$\tilde{u}^{D}_{ijm,2019} = \delta_{jm,2019} + \hat{\tau}_j + \tilde{\gamma}_{jm,2019} + \zeta_{igmt} + (1-\sigma)\varepsilon_{ijmt},$$

while her actual utility is given by:

$$\tilde{u}^{A1}_{ijm,2019} = \delta_{jm,2019} + \hat{\tau}_j + \zeta_{igmt} + (1-\sigma)\varepsilon_{ijmt}.$$

The consumer surplus is given by

$$CS = \frac{1}{\alpha} E(\tilde{u}^{D}_{ijm,2019}) - \frac{1}{\alpha} \sum s_{jmt} \times \tilde{\gamma}_{jm,2019},$$

where s_{jmt} is the market share of drug j chosen under the choice utility $\tilde{u}^{D}_{ijm,2019}$.

Last, we consider the case with the policy demand shifters, and assume consumers' brand preferences are welfare irrelevant. Consumer i's decision utility is given by:

$$\tilde{u}^{D}_{ijm,2019} = \delta_{jm,2019} + \hat{\tau}_j + \tilde{\gamma}_{jm,2019} + \zeta_{igmt} + (1-\sigma)\varepsilon_{ijmt},$$

while her actual utility is given by:

$$\tilde{u}^{A0}_{ijm,2019} = \delta'_{jm,2019} + \hat{\tau}_j + \zeta_{igmt} + (1-\sigma)\varepsilon_{ijmt}$$

The consumer surplus is given by

$$CS = \frac{1}{\alpha} E(\tilde{u}^{D}_{ijm,2019}) + \frac{1}{\alpha} \sum s_{jmt} \times (\delta'_{jm,2019} - \delta_{jm,2019} - \tilde{\gamma}_{jm,2019}),$$

where s_{jmt} is the market share of drug j chosen under the choice utility $\tilde{u}^{D}_{ijm,2019}$.

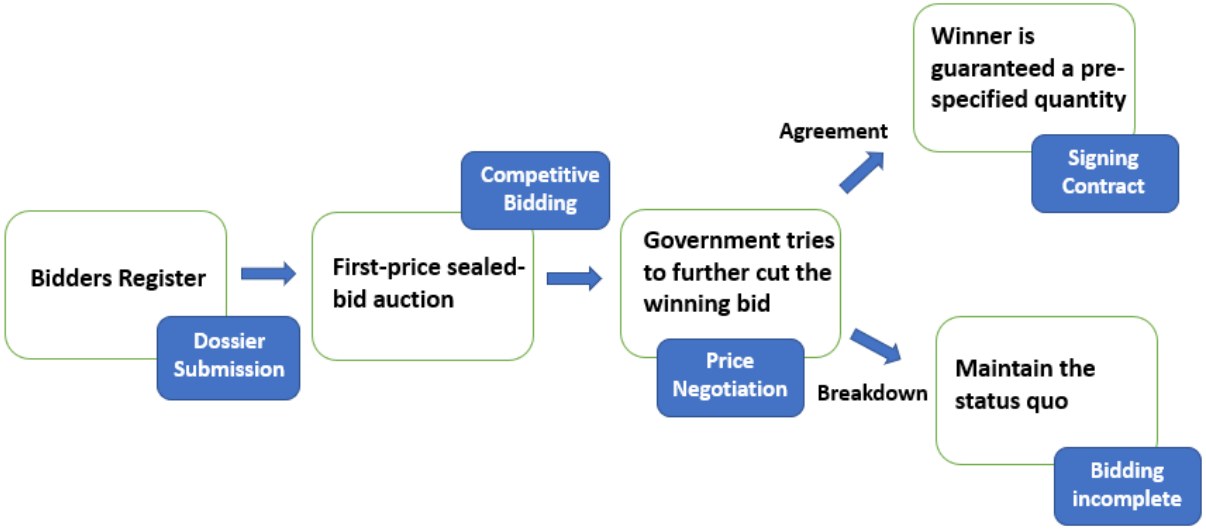
D Supplemental Tables and Figures

Table A.1: First Stage of Demand Estimation

	Logit IV	Nested Logit IV	
	(1)	(2)	(3)
	Price	Price	Cond. Mkt Share
Hausman IV	0.126	0.126	-0.0106
	(0.0113)	(0.0112)	(0.0514)
BLP IV		-0.00722	-0.173
		(0.00347)	(0.0288)
Passed BE	-0.0238	-0.0239	0.197
	(0.00848)	(0.00852)	(0.0679)
N	3191	3191	3191
F	67.20	48.33	14.55
Province-Drug FE	Yes	Yes	Yes
Province-Year FE	Yes	Yes	Yes

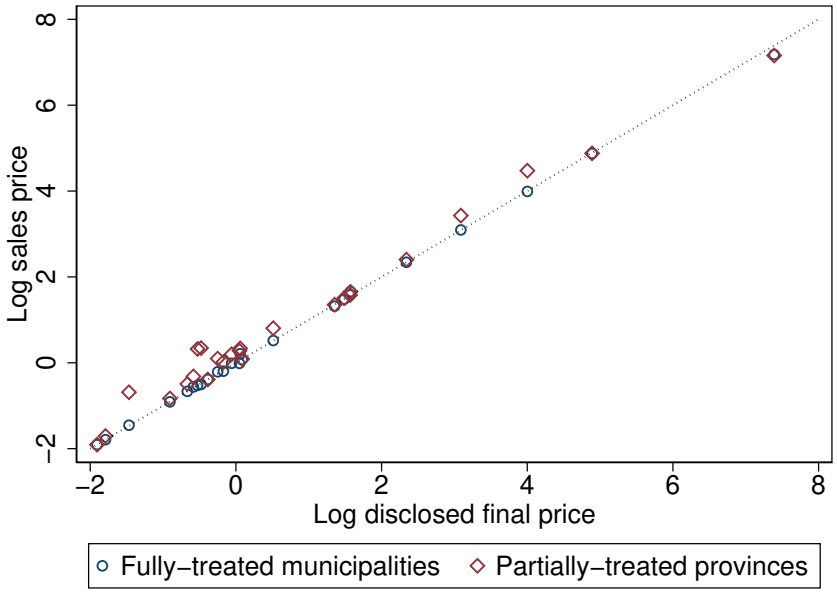
Notes: This table shows the first-stage regressions of the demand estimation using instrumental variables, corresponding to columns (2) and (4) in Table 2. Hausman price IV is defined as the average price of the same product sold in other provinces in the same year. BLP IV is defined as the total number of drugs in each molecule in the market. Standard errors are clustered at the province-year level and are reported in the parentheses.

Figure A.1: Bidding and Negotiation Process



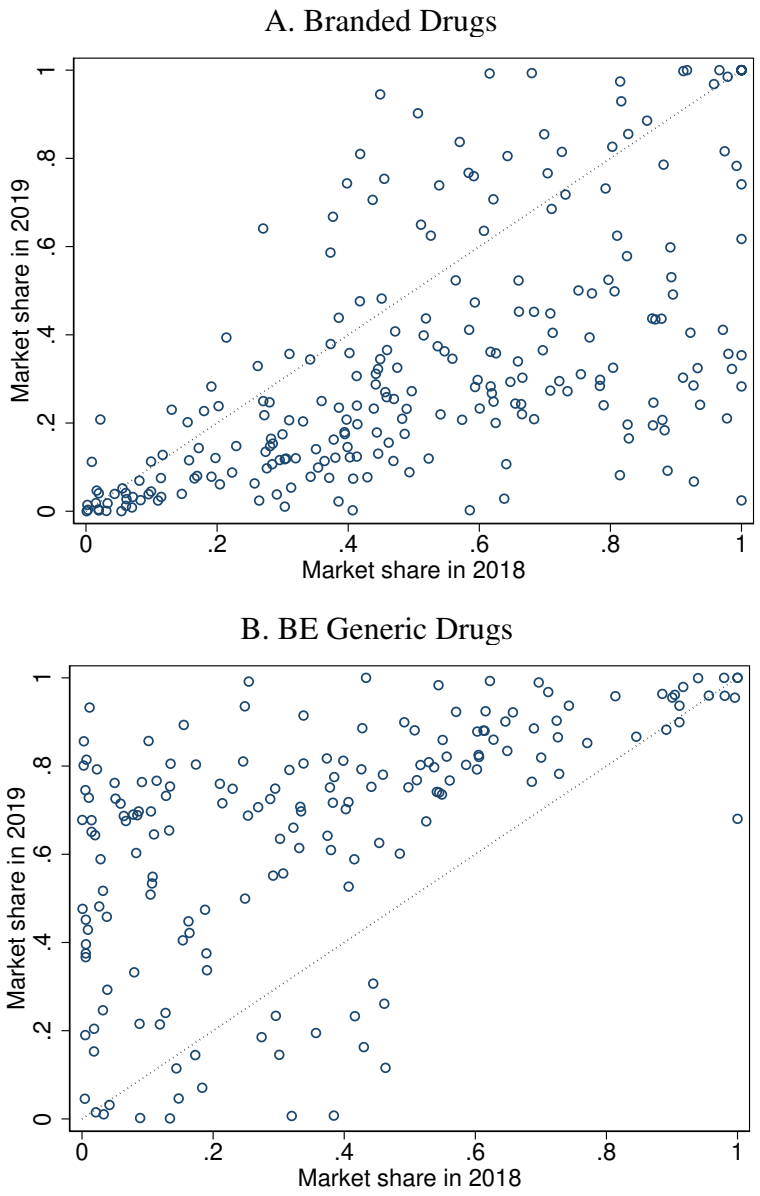
Notes: This figure shows the detailed bidding and negotiation process. The drug producers first register for the competitive bidding and submit their dossiers. Then they participate in a first-price sealed-bid auction. The bidder with the lowest bidding price is selected as the auction winner. In some scenarios, the government would initiate a negotiation process with the auction winner in an attempt to further cut down the price. If a deal is not reached at this stage, the status quo would be maintained. Otherwise, the auction winner would be guaranteed a pre-specified quantity and sell their drugs at the final negotiated price.

Figure A.2: Compare Sales Price and Disclosed Final Price



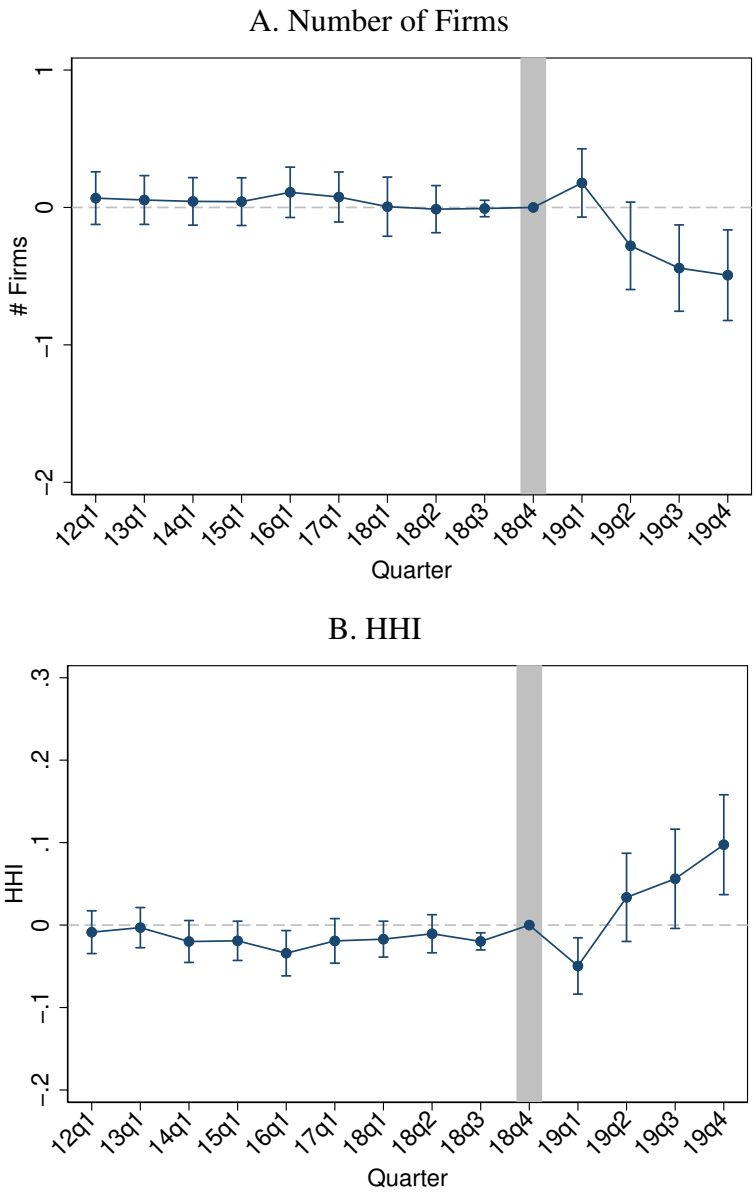
Notes: This figure compares the final negotiated prices disclosed by the government and the sales prices observed in the data. The x-axis represents the log of final prices disclosed by the government, and y-axis represents the log of sale prices of the winning product in enacting provinces observed in our data. We categorize provinces by whether they are fully-treated (i.e., four municipalities including Beijing, Shanghai, Chongqing and Tianjin) or partially-treated with only one or two large cities covered by the competitive bidding. The gray dotted line represents the 45-degree line.

Figure A.3: Market Shares of Branded and Generic Drugs in Enacting Provinces



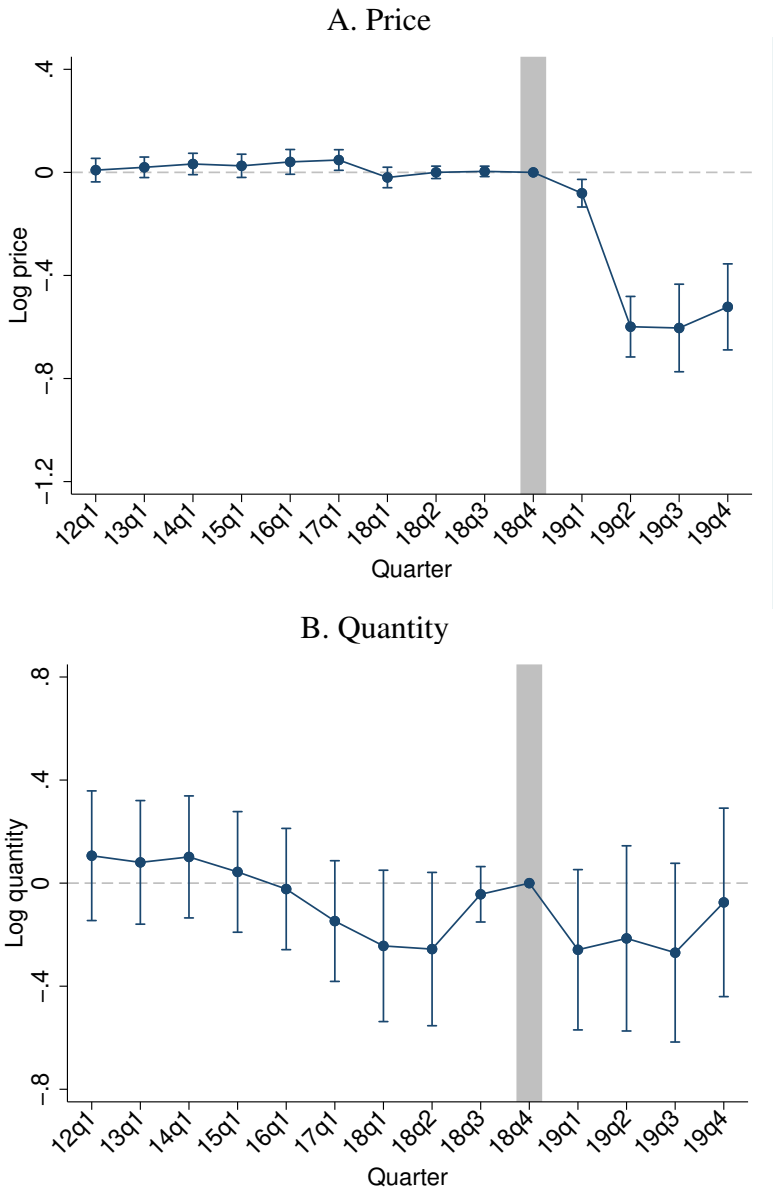
Notes: This figure plots the market shares for branded drugs (Panel A) and BE generic drugs (Panel B) in enacting provinces prior to the competitive bidding (2018) and after (2019). Since the market guarantee of the competitive bidding went into effect in March 2019, we drop the first quarter when computing market shares in 2019. Each dot is a molecule-province pair for all bidding molecules and enacting provinces.

Figure A.4: Policy Effects on Market Concentration



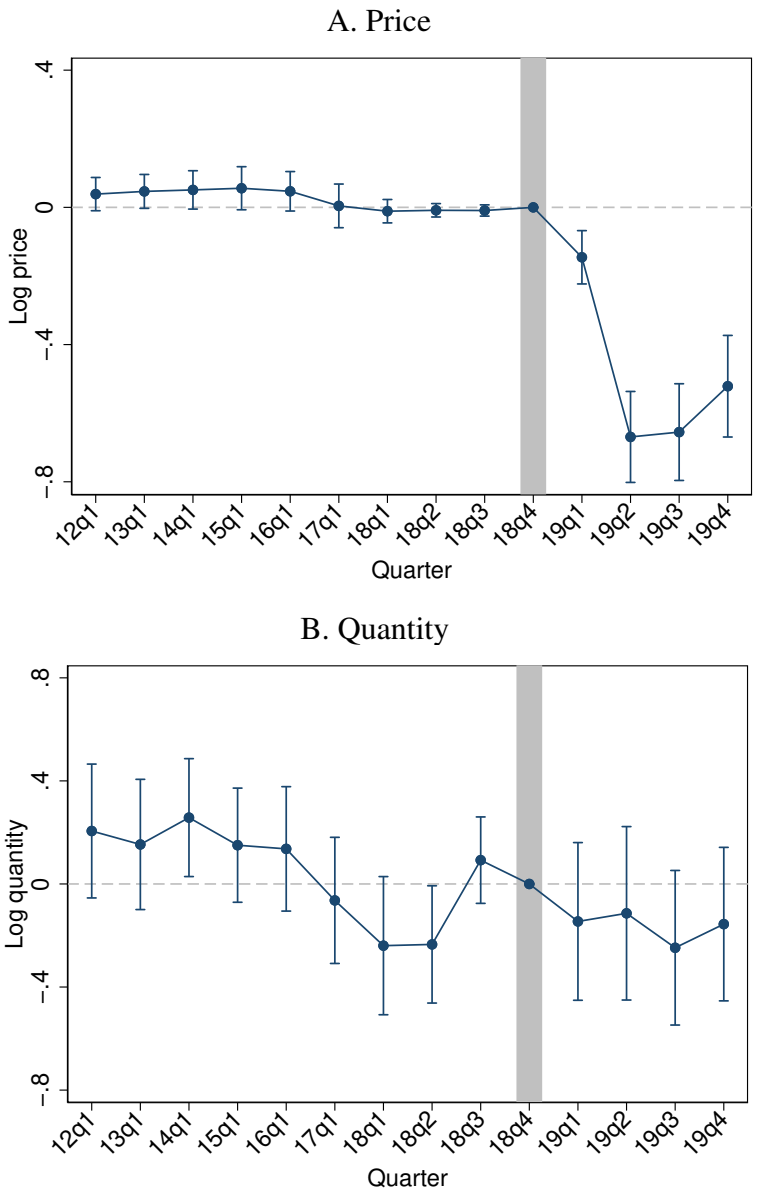
Notes: This figure shows the effects of the competitive bidding on the number of active firms and Herfindahl-Hirschman Index (HHI), estimated using the specification shown in Equation (1). Each dot represents the regression coefficient for the corresponding quarter, and each line segment represents the 95% confidence interval using standard errors clustered at the province-year level. The coefficient of the fourth quarter in 2018 is normalized to zero. From 2012 to 2017, we only show coefficients for the first quarter of each year to conserve space.

Figure A.5: Policy Effects on Drug Prices and Sales, Sales Weighted



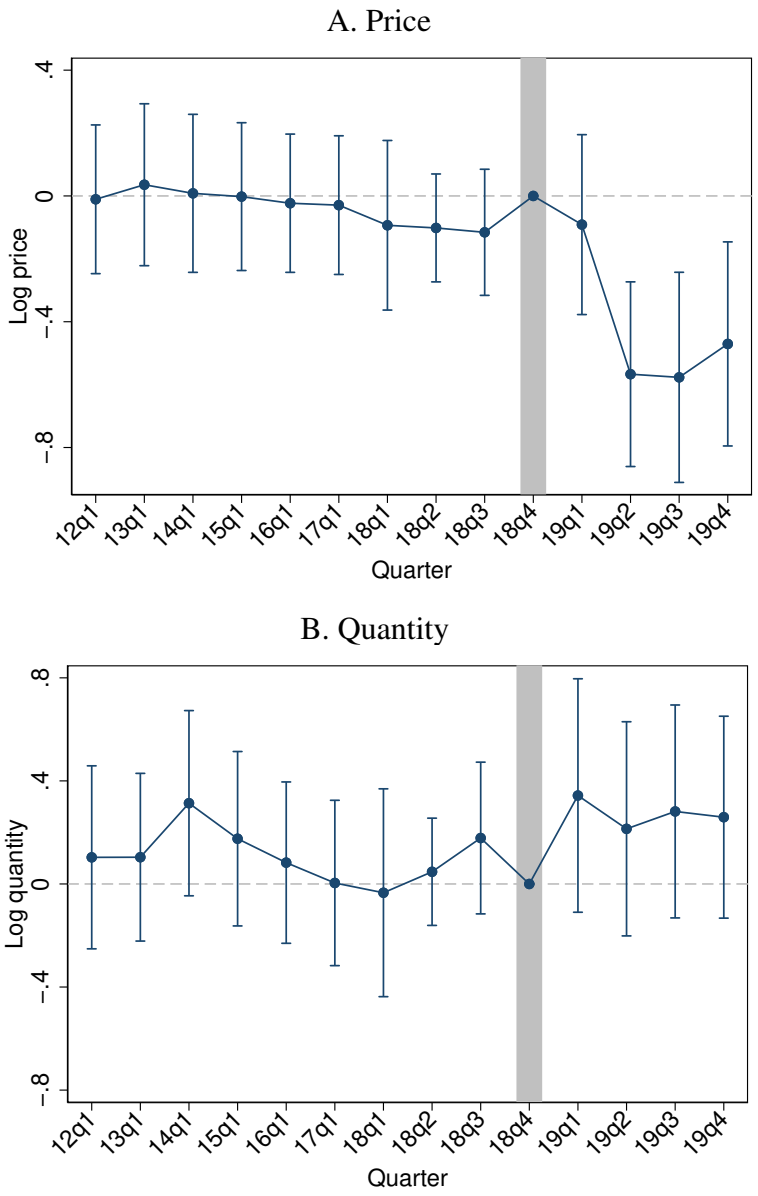
Notes: This figure shows the effects of the competitive bidding on drug prices and sales, similar to Figure 4. The only difference is that each observation is weighted by the total sales quantity at the province-molecule level during 2012-2018 in this figure.

Figure A.6: Policy Effects on Drug Prices and Sales, Excluding Partially Treated Provinces



Notes: This figure shows the effects of the competitive bidding on drug prices and sales, similar to Figure 4. The only difference is that only four municipalities, Beijing, Tianjin, Shanghai, and Chongqing, are included as enacting provinces in this figure. We drop provinces where seven other enacting cities are located, as these provinces are only partially treated.

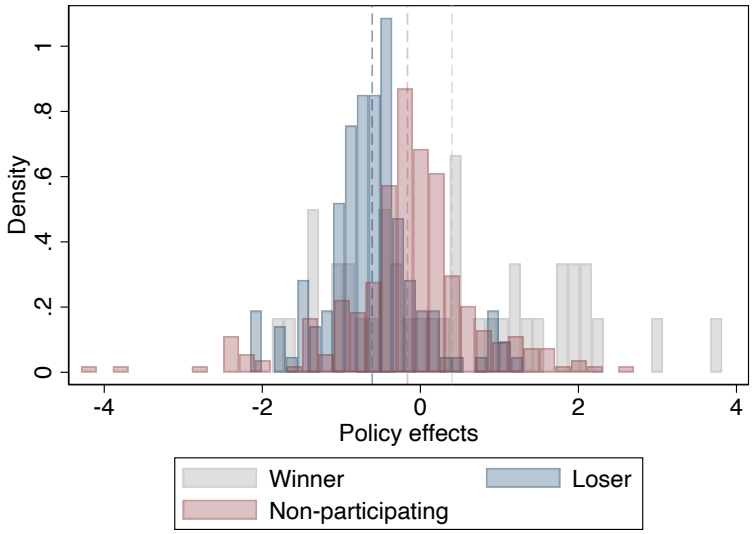
Figure A.7: Policy Effects on Drug Prices and Sales, DDD Specification



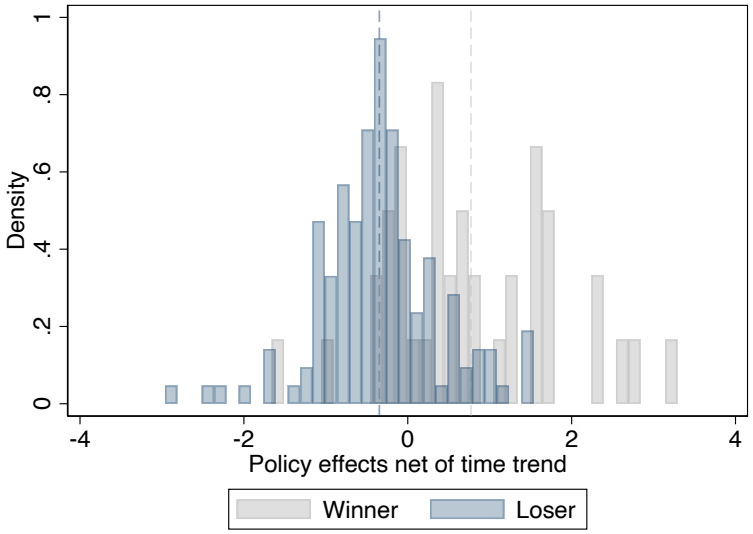
Notes: This figure shows the effects of the competitive bidding on drug prices and sales, estimated using the DDD specification described in Section 4.3. Each dot represents the regression coefficient for the corresponding quarter, and each line segment represents the 95% confidence interval with standard errors clustered at the province-year level. The coefficient of the fourth quarter in 2018 is normalized to zero.

Figure A.8: Distribution of Policy Demand Shifters

A. Policy Shifters

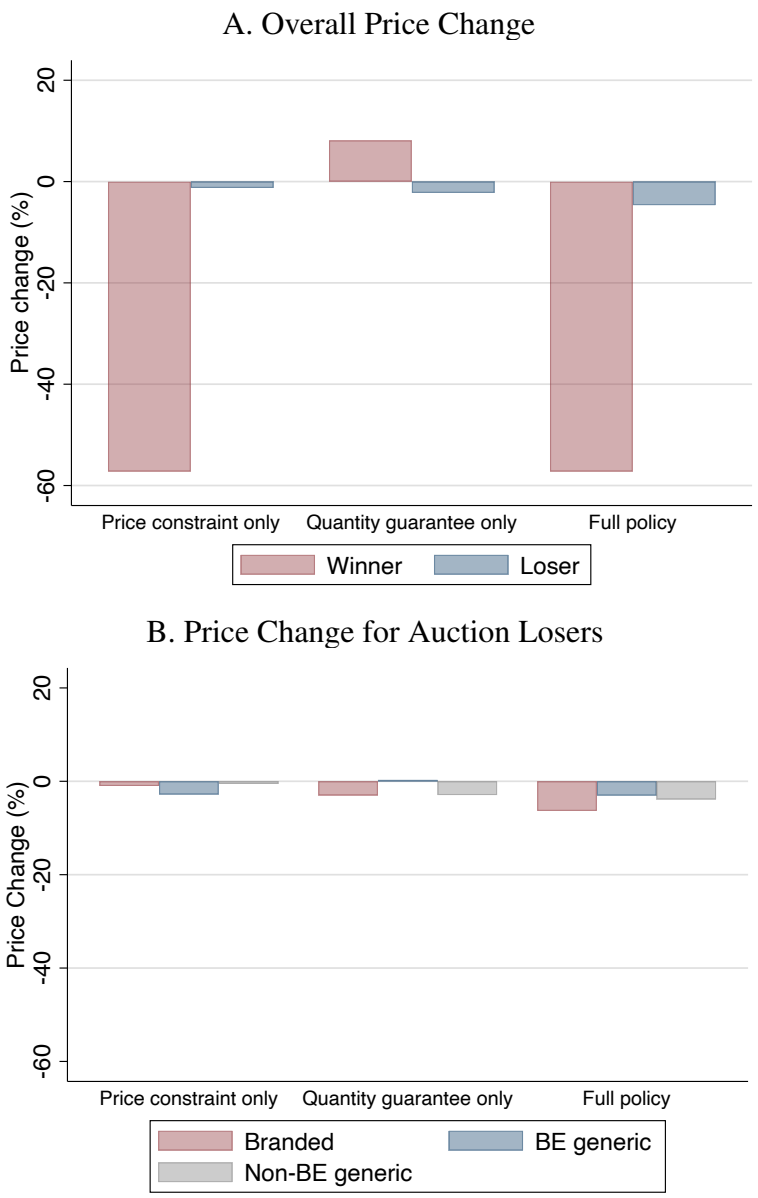


B. Policy Shifters Net of Time Trend



Notes: This figure shows the effects of the policy on drug demand. We assume there exists a policy demand shifter for each available product in each market during the post-auction period, and consumers' price sensitivity and the within-nest preference correlation remain unchanged after the auction. We estimate policy demand shifters at the province-product level such that the observed and predicted market shares during the post-auction period are perfectly matched. Panel A plots the distribution of the policy demand shifters for auction winners and losers in enacting provinces and products in non-enacting provinces separately. The policy effects for products in non-enacting provinces can reflect the time trend, so we normalize all policy demand shifters such that the mean for these products is 0. In Panel B, we plot the net policy shifters for auction winners and losers in enacting provinces by subtracting the corresponding average policy shifters in non-enacting provinces for the same product. The dash lines represent the average value of each group separately.

Figure A.9: Price Changes under Counterfactual Scenarios



Notes: This figure shows price changes under three counterfactual scenarios. The status quo is the pre-auction period in 2018. In the first scenario, consumer demand remains fixed, prices of auction winners are fixed at the post-auction prices, and other firms can freely adjust their prices. In the second scenario, consumer demand is distorted by the quantity guarantee and all firms can freely adjust their prices. In the third scenario, consumer demand is distorted by the quantity guarantee, prices of auction winners are fixed at the post-auction prices, and other firms can freely adjust their prices. The third scenario corresponds to the actual competitive bidding that took place. We calculate price changes for auction winners and losers separately. Panel A shows the price changes for winners and losers separately. Panel B shows the price changes for branded and generic losers separately. See Section 6 for more details.