IDB WORKING PAPER SERIES Nº IDB-WP-1441

More Information, Lower Price? Access to Market-Based Reference Prices and Gains in Public Procurement Efficiency

José Martinez-Carrrasco
Otavio Conceição
Ana Lucia Dezolt

Inter-American Development Bank
Institutions for Development Sector, Fiscal Management Division
Office of Strategic Planning and Development Effectiveness

March 2023

More Information, Lower Price? Access to Market-Based Reference Prices and Gains in Public Procurement Efficiency

José Martinez-Carrrasco (Inter-American Development Bank)

Otavio Conceição (Sao Paulo School of Economics)

Ana Lucia Dezolt (Inter-American Development Bank)

March 2023

Cataloging-in-Publication data provided by the
Inter-American Development Bank
Felipe Herrera Library

Martínez Carrasco, José.
More information, lower price?: access to market-based reference prices and gains in public procurement efficiency /
José Martinez-Carrrasco, Otavio Conceição, Ana Lucia Dezolt.
p. cm. — (IDB Working Paper Series ; 1441)
Includes bibliographic references.
1. Government purchasing-Brazil. 2. Electronic invoices-Brazil. 3. Information storage and retrieval systems-Brazil.
4. Prices-Brazil. I. Canozzi, Otavio. II. Dezolt, Ana Lucía. III. Inter-American Development Bank. Fiscal Management Division.
IV. Inter-American Development Bank. Office of Strategic Planning and Development Effectiveness. V. Title. VI. Series.
IDB-WP-1441

http://www.iadb.org

More Information, Lower Price? Access to Market-Based Reference Prices and Gains in Public Procurement Efficiency *

José Martinez-Carrrasco[†] Otavio Conceição[‡] Ana Lucia Dezolt[§]

March 19, 2023

Abstract

The paper examines the impact of providing market-based reference prices on public procurement efficiency in Brazil. Specifically, the study focuses on the State Secretariat of Health (SES) in Rio Grande do Sul and the algorithm developed by the local tax administration to calculate representative reference prices for pharmaceutical products. Unlike previous studies, reference prices are calculated based on the universe of local business-to-business transactions. The study finds that SES procurement officers' access to this information caused a significant reduction in purchase unit prices, particularly for products characterized by a higher ex-ante unit price, a smaller number of suppliers, and purchased by a smaller number of public institutions. The gains in efficiency are attributed to the use of up-to-date market information, which is particularly useful for products where information asymmetry is more likely to exist between procurement officers and private providers.

Keywords: Public procurement, reference prices, electronic invoicing, Brazil.
JEL Codes: D40, H51, H57.

---

*We thank the technical staff of the State Secretariat of Finance of Rio Grande do Sul (SEFAZ/RS), especially Evelise Dalmoro, Guilherme Petry, Marcos Bosio, Eduardo Akira, and Israel Fama. We remain responsible for any errors.

[†]Inter-American Development Bank (IDB). E-mail: jalejandrom@iadb.org

[‡]Sao Paulo School of Economics (EESP FGV). E-mail: otavio.canozzi@gmail.com

[§]Inter-American Development Bank (IDB). E-mail: anapa@iadb.org

1 Introduction

In the last decades, government purchases have become an important space for the adoption of new technologies, especially those related to big data and artificial intelligence. The use of technology holds the promise of helping bureaucrats make better decisions, reducing costs, and improving procurement outcomes. This includes the utilization of machine learning algorithms to predict and prevent malfeasance using e-procurement systems (Huber and Imhof, 2019; Gallego et al., 2021; De Michele and Vieyra, 2022), the real-time monitoring of public expenditures (Lauletta et al., 2019), and the establishment of competitive parameters for auctions (Decarolis, 2014; Fazekas and Tóth, 2017; Ferraz et al., 2015). The potential impacts are huge, as nearly USD 11 trillion are spent on procurement globally every year, and therefore even savings of 1 percent can amount to USD 110 billion yearly (Fazekas and Blum, 2021). Given its economic relevance, improving the efficiency of public procurement is of paramount importance for many governments worldwide.

This paper studies the impacts of an intervention that provided frontline procurement officers of the State Secretariat of Health (SES) of Rio Grande do Sul, Brazil, with market-based representative prices for the purchase of medicines. In most tendering processes, public procuring bodies such as SES are required to define reference prices in auctions for the acquisition of off-the-shelf goods.[1] This is a critical phase of the tender process and is often subject to some discretionary action on the part of procurement officers.[2] On the one hand, discretion may allow competent procuring agencies to exploit their local expertise to ensure equal dimensions that are difficult to cover by explicit contractual terms (Baltrunaite et al., 2021). On the other hand, it

[1] Off-the-shelf goods are products for which the price is often the unique relevant dimension to consider in the procurement process. Such products vary from pharmaceuticals to printing paper. The quality, for example, is typically not a matter of concern, unlike services like roads and school construction, which do not qualify as off-the-shelf (Ferraz et al., 2015).

[2] The importance of the reference prices is that they are often correlated with the final purchase prices because they are used as a tender parameter to select suppliers. A major challenge to define reference prices is how to set prices that are fair and compatible with the local reality, i.e., neither above nor below the market prices, so that good suppliers are not deterred by mistaken reference prices.

increases the dependence of the public procurement efficiency on the officers' ability to find and set good reference prices (Coviello et al., 2018; Decarolis et al., 2020a; Lotti and Spagnolo, 2022). The asymmetry of information between procurement officers and potential sellers regarding market reference prices (Grennan and Swanson, 2020) is an important concern. Asymmetric information diminishes officers' bargaining power to negotiate competitive quotations. Furthermore, in the case of products with few providers, it can lead to collusion and result in relatively high prices. Asymmetric information is more likely to exist in products for which procurement officials have less prior experience and those with few potential providers.

To mitigate this type of asymmetric information, the State Secretariat of Finance of Rio Grande do Sul (SEFAZ/RS) created an algorithm that leverages e-invoicing data about local business-to-business (B2B) transactions to calculate market-based reference prices. Since July 2014, the reference prices produced by the algorithm are provided for SES procurement officers to be used as inputs in the tendering of some pharmaceutical products. This is done before they set the tendering parameters, which potentially reduces the asymmetry of information related to prices prevalent in the local B2B market and may also allow procurement officers to negotiate better prices with potential suppliers. Importantly, the production of the market-based reference prices is performed by SEFAZ/RS officials as part of their routine workload, with large economies of scale in terms of the number of products for which prices can be generated on a regular basis, and follows a well-established process of communication with SES.

We investigate whether access to market-based reference prices translates into efficiency gains. Our analysis is based on a high-granularity and novel data set that combines two administrative records from different sources. The first is the electronic invoice data (*Nota Fiscal Eletrônica, NF-e*), administered by SEFAZ/RS, which contains detailed information about the universe of medicines acquisitions made by both state and municipal procuring agencies of Rio Grande do Sul between 2013 and 2015. The second is a data set that contains a list of products tendered by SES during 2013-2015. Products are uniquely identified by the triple substance-dosage-dosage form. These data sets are not public and were obtained through partnership agreements with SEFAZ/RS.

To estimate the causal effects of interest, we propose a triple difference design that compares SES with public buyers other than SES before and after the market-based reference prices were provided for SES procurement officers. There are two sets of products: the treated and the control ones. The former are those for which the market-based reference prices were provided to SES, while the latter are those for which SES had no access to that information during the study period. The use of public rather than private buyers as the comparison group for SES is motivated by the fact that they potentially operate with similar conditions for delivery, payment, and monitoring of suppliers. By Brazilian law, public procuring bodies are obliged to require invoices for the purchase of medical drugs and, since 2013, the invoices whose buyers are public agencies in Rio Grande do Sul are mandatorily registered using the NF-e system. Therefore, we have access to the universe of pharmaceutical acquisitions made by such public agencies.

The main results suggest that the provision of market-based reference prices before procurement officers set the tendering parameters lowered the purchase prices of treated products. We show that the effect is entirely driven by a subset of treated products. They are characterized by having: (i) higher ex-ante unit price, (ii) a smaller number of suppliers, and (iii) a smaller number of public buyers. The purchase prices of such products were reduced, on average, by 13.2 percentage points. We argue that the findings are consistent with a theoretical model in which the provision of up-to-date reference prices from the local B2B transactions would be effective to improve efficiency, especially for products that are purchased with lower frequency by public institutions. Arguably, procurement officers have less information on products bought in smaller overall quantities, and for which there is a lower number of potential suppliers and governmental buyers.

We find no impacts on the buyer-supplier relationships or volume purchased, which suggests that the unit price reduction is not due to supplier substitution or changes in bundle composition. Anecdotal evidence from SES officials suggests that increased bargaining power and the possibility of setting lower reference prices in the tender process are important drivers of our findings. In fact, the average purchase prices for SES of products other than the most expensive treated ones have increased at least 5 percent in the post-treatment period, while for the latter the observed variation was -8.7 percent. This indicates that in the absence of the intervention the prices for such products would probably increase. This also reinforces the argument that the results are not driven by structural reforms implemented by SES in the study period, but relate to specific circumstances made possible by the provision of market-based reference prices for these products.

Our findings show that technological interventions can improve procurement outcomes by providing actionable information for frontline officers at a relatively low cost and under limited maintenance investment. Such a result is especially important given the recent empirical evidence that regulation and enforcement do not necessarily translate into better practices and purchasing prices (Gerardino et al., 2017; Bucciol et al., 2020; Bosio et al., 2022), and that bureaucratic competence is a major driver of procurement efficiency (Decarolis et al., 2020a; Best et al., 2022). Therefore, more cost-effective ways of combating passive waste (Bandiera et al., 2009) are needed to improve the performance of governments in public procurement.

This paper contributes to the literature in at least three strands. First, it is, to the best of our knowledge, the first paper to evaluate the impacts of providing reference prices based on local B2B transactions to a large public buyer of medicines. Despite the importance of how reference prices are determined in public tenders, there is to date little empirical evidence about the impacts of public procurement interventions related to the use of market-based reference prices. In the context of US hospitals, Grennan and Swanson (2020) shows that access to information on purchasing by peer hospitals leads to reductions in the prices that hospitals negotiate for supplies. Therefore, the provision of information on the prices that private firms and peer buyers are paying locally for the same products can improve one's procurement efficiency. Second, this paper innovates by providing a highly reliable identification of off-the-shelf goods. As in Best et al. (2022), we consider invoicing data relative to winners of electronic auctions to study the performance of distinct public organizations. However, differently from Best et al. (2022), who rely on text analysis and machine learning techniques to identify products and classify them into homogeneous bins, we take advantage of a global unique product identifier, the Global Trade Item Number (GTIN).[3] The use of GTIN ensures that we are comparing exactly the same products in terms of dosage, dosage form, manufacturer, and brand (from the pharmaceutical company that markets the medicine) across distinct buyers. This is particularly relevant considering that the dosages, dosage forms, brands, and manufacturers of the same active principle of a given medicine - which are associated with distinct GTINs - have a major influence on the purchase price. Such identification is made possible because pharmaceutical suppliers are required by Brazilian law to include the GTIN of each medicine in the transaction invoices.

Third, our setting allows us to precisely measure the purchase unit price because SEFAZ/RS developed a procedure - which is part of the pricing algorithm - that provides a reliable calculation of the purchase unit price of each medicine in the transaction invoices. This means that we have access to an estimate of the price of the smallest unit of a medicine, such as an individual tablet, even though the medicine is sold at a fixed higher quantity of its smallest unit (e.g., a blister pack with a number of tablets). In particular, the SEFAZ/RS pricing algorithm is able to calculate the unit price of a given pharmaceutical by comparing its purchase price in the invoice with the quantity of the smallest unit of such pharmaceutical that is marketed in Brazil. The information about the quantity of the smallest unit of a given medicine that is typically sold comes from a database maintained by a private firm that registers this information for all pharmaceuticals marketed in Brazil.

This paper also contributes to the literature in economics that studies the effects of using new technologies in the public procurement process (Lewis-Faupel et al., 2016; Schøll and Ubaydi, 2017; Kovalchuk et al., 2019). As Fazekas and Blum (2021) highlight, the available evidence on the role of technology in improving procurement efficiency is still scarce. The evidence base not only for interventions focused on enhancing procurement officers' capacity, but also others

[3] The Global Trade Item Number (GTIN) is a globally unique identifier number for commercial items, which is widely used in physical retail and managed by a global organization called GS1. It takes the form of a barcode and identifies the product, the manufacturing company, and the company's country of origin.

such as e-procurement, pay-for-performance schemes, and watchdog portals is, surprisingly limited. Procurement digitization, however, has been shown to result in increased competition, higher quality suppliers, and improved public works by improving access to information and reducing face-to-face interactions with potentially corrupt public officials (Lewis-Faupel et al., 2016). Digitization also appears to decrease the time and resources required for the contracting process (Schøll and Ubaydi, 2017; Kovalchuk et al., 2019).

This paper also relates to the literature on public procurement and government efficiency. Available evidence in this area indicates that public procurement in many countries suffers from manipulation (Porter and Zona, 1993), overpricing (Singh et al., 2011), and favoritism (Arvate et al., 2016; Titl and Geys, 2019). Studies also highlight other factors such as limited competence among public buyers to find competitive prices (Lotti and Spagnolo, 2022) and the existence of collusion between politicians and private companies (Coulomb and Sangnier, 2014; Baltrunaite et al., 2021). In particular, the literature points out that corruption is higher when contracts are renegotiated and the forms of contracting allow for greater discretion on the part of the purchasing body (Decarolis, 2014; Lichand et al., 2016; Palguta and Pertold, 2017; Coviello et al., 2018). Literature also shows that introducing contractual mechanisms that take into account providers' procurement reputation can prevent malfeasance (Calzolari and Spagnolo, 2006; Decarolis and Palumbo, 2015; Decarolis et al., 2020b).

The remainder of the paper is organized as follows. Section 2 presents the institutional background. Section 3 presents the data, the record linkage procedures, and the sample selection. Section 4 shows the descriptive statistics and outlines the empirical strategy. Section 5 is reserved for the main results, and Section 6 concludes.

2 Institutional Background

2.1. Use of Electronic Invoicing for Improving Procurement Outcomes

Since 2008, all Brazilian registered firms, subject to taxation on the circulation of goods and services (*ICMS*), are obliged to issue the electronic invoice (NF-e) in their transactions.[4] The NF-e is the digital version of the tax document issued in a given taxable transaction. It has the same purposes as the paper invoice and complies with the same principles, i.e., authenticity, integrity, and readability (Barreix and Zambrano, 2018). Its legal validity is guaranteed by the issuer's digital signature, and its use is authorized by the tax authority. The NF-e project in Brazil was developed in 2005 by the Brazilian Internal Revenue Service (*Receita Federal do Brasil, RFB*) in partnership with the state secretariats of finance of the subnational governments. At that time, the adoption of the NF-e system was voluntary for state governments, and took place early in states such as Rio Grande do Sul and São Paulo (Santos et al., 2015).[5,6]

As the digitization of invoices proceeded in Brazil, the information about goods transactions has become part of large-scale datalakes administered by the state secretariats of finance. In this context, SEFAZ/RS developed a methodology to calculate market-based reference prices for products commonly purchased in procurement processes. The goal was to provide relevant information to guide procurement officers in both administrative and judicial proceedings. As part of their habitual workload, officials of the Treasury Department of SEFAZ/RS in 2014 started to develop a pilot of the pricing algorithm. Once the algorithm was operational, it started to be experimentally used in the calculation of representative prices. In July 2014, SEFAZ/RS staff proceeded with the pilot implementation of the tool restricted to medicines.

---

[4] ICMS (*Imposto sobre Circulação de Mercadorias e Prestação de Serviços de Transporte Interestadual e Intermunicipal e de Comunicação*) is a value-added tax whose collection is under the responsibility of state tax authorities in Brazil. For most Brazilian states, it is the most important component of the state tax revenue.

[5] Rio Grande do Sul adopted the e-invoice system for B2B transactions in 2006.

[6] In 2010, to continue the modernization process of the local tax authority, the state of Rio Grande do Sul requested a loan amounting to USD 60 million from the Inter-American Development Bank (IDB). The plan for the use of resources involved a strong investment in technology and process re-engineering to increase state revenue, but also improving the efficiency of its operations and providing better services to citizens.

2.2. *Technical Aspects of the Pricing Algorithm*

The elaboration of the market-based reference prices using the NF-e data follows a workflow consisting of three major steps: (i) product identification, (ii) data treatment, and (iii) calculation of prices. The first step consists of the identification of the product to which a given invoice refers in the NF-e database. Due to regulations by the National Health Surveillance Agency (*Agência Nacional de Vigilância Sanitária, Anvisa*), the suppliers of medicines are obliged to provide the GTIN code of each medicine purchased in a given transaction with the public sector, which facilitates this step. In order to map the barcodes (GTIN) of each product in a given invoice to the usual product name in Brazil, SEFAZ/RS retrieves information from the Guia Brasíndice. This is a database administered by a private firm and distributed bi-weekly to subscribers. The database, which was created in 1960, registers the prices of all pharmaceuticals sold in Brazil. It is considered the first price guide for medicines and hospital supplies in the country (Dezolt and Muñoz, 2020). It is a reference in the area and presents information that assists in the price search, such as the name of the medicine, substance, dosage, and dosage form.

The second step is characterized by a procedure called *fractionation*. Its main goal is to create variables that indicate whether the products to which an invoice refers are priced per unit of presentation (such as blister packs) or per sets of units (such as an individual tablet). This procedure involves comparing the price registered in the invoice with (i) the price of the medicine per *fractioning* unit and (ii) the price of the medicine per presentation set, according to the Guia Brasíndice database. The Guia Brasíndice database contains both manufacturer and final consumer prices for the same product with granular information for each Brazilian state. Therefore, if a given medicine is sold as a blister pack with, for example, 30 tablets and the medicine's purchase price is BRL 30, the pricing algorithm will be able to impute the purchase unit price as BRL 1. Additionally, the pricing algorithm compares the price estimate of the smallest unit of the medicine with that from the Guia Brasíndice database and performs a statistical analysis that removes outliers. The result of the *fractioning* is a list of invoices for which the algorithm confidently determines the purchase unit price of the products registered.

The third step is the calculation of the final unit reference price. For each product, the prices generated in the second step are compared with several indicators regularly published by Anvisa such as the Maximum Price for Sale to Government (*Preço Máximo de Venda ao Governo, PMVG*) and the Manufacturer Price (*Preço de Fábrica, PF*). In the case of the SES, for each item whose pricing was requested, SEFAZ/RS provided a price range with several values (median and quartiles) to characterize the price distribution based on the local B2B transactions. Besides the calculation of the reference price, an analysis of economy of scale is performed, which seeks to identify whether for the volume purchased that the SES wants to bid, a gain of scale can be expected. The presentation of the reference prices is formatted in such a way as to prevent the identification of individual business transactions. In particular, to guarantee tax secrecy, the market-based reference prices are provided only if there are at least three price quotations from different sellers for each product.

2.3. The Procurement Process at SES

In the study period, SES was the largest individual buyer of medicines in Rio Grande do Sul, with approximately 1,000 products and more than 9 million units purchased every year between 2013 and 2015. Such products used to be distributed to local basic health units and state hospitals.[7] At the time to which this paper refers, SES had eight auctioneers working full-time in the procurement processes, and all of them were public servants with specific training for developing this activity. The procurement was conducted through electronic auctions in an online platform, and the reference price for each specific item tendered was determined manually by each auctioneer. The process typically involved consulting price parameters based on historical purchases made by other public procuring entities in Brazil or making direct requests of price quotations to potential suppliers. This was done through access to comprehensive data sets such as the Health Price Bank (*Banco de Preços em Saúde*, BPS), administered by the Brazilian Ministry of Health. Like SES, there were several other public buyers of pharmaceuticals products in Rio Grande do

[7] In Brazil, the public provision of health services is performed by the federal, state, and municipal government through Brazil's unified health system (*Sistema Único de Saúde, SUS*).

Sul, including health intermunicipal consortia, municipal administrations, and state hospitals, with varying levels of structure in terms of staff, capacity, and purchase scale.

3 Data, Record Linkage, and Sample Selection

In this paper, we use two administrative data sets, which are described in detail below. Both are confidential and were provided by SEFAZ/RS.

3.1. Data Sources

Products purchase data (NF-e). The e-invoice data from SEFAZ/RS contains detailed information about the universe of invoices relative to the business-to-business market that were electronically issued in the state of Rio Grande do Sul in a given period. Each invoice corresponds to a given transaction and may refer to the purchase of more than one product. The variables of the NF-e database include (i) the unit price of each product; (ii) the total quantity of each product (in units); (iii) the year-month in which the invoice was issued; (iv) the barcode (GTIN) of each product; and (v) the product description, which is a text input field limited to 120 characters filled by the invoice issuer.[8]

In addition, the NF-e database registers the taxpayer identification number (CNPJ) of both the buyer and seller in each transaction.[9] Importantly, the data set contains information covering the universe of medicines acquisitions made by public entities located in Rio Grande do Sul.[10] We were given access to a subset of the NF-e data relative to the purchases of pharmaceuticals made by public procuring agencies in Rio Grande do Sul in the period between January 2013 and

[8] In the case of medicines, sellers are obliged to provide the Global Trade Item Number (GTIN) code when the buyer is a state procuring agency of Rio Grande do Sul. Suppliers of state procuring agencies have become required to indicate the GTIN code in the e-invoices of medicines as of February 2013. For details, see State Decree No. 51,200/2013. Although the formal obligation for state procuring agencies was established in 2013, the inclusion of the GTIN code was already enforced by Anvisa for municipal, state, and federal public entities.

[9] The tax identification number of firms in Brazil is a 14-digit sequence called *Cadastro Nacional da Pessoa Jurídica* (CNPJ) and is issued by the Brazilian Internal Revenue Service (RFB).

[10] This is warranted by the fact that the invoices relative to purchases of medicines by public agencies in Rio Grande do Sul are issued electronically since 2013.

December 2015. Since the original NF-e data contains only the specific presentation purchase price of each pharmaceutical and possibly not the price of the smallest unit of the drug, we asked SEFAZ/RS to undertake the *fractioning* procedure in our data set. Therefore, the unit purchase prices we refer to throughout the paper are the prices of the smallest possible unit of a given medicine.

List of treated products. SEFAZ/RS provided a list of 84 pharmaceuticals whose reference prices were required by SES for public tenders that took place as of July 2014. Each product in the list is uniquely identified by three variables: (i) substance, (ii) dosage, and (iii) dosage form (e.g., capsule, tablet, or liquid). As such, the products of the list are the general form of a given product, and not exactly a specific product of a given manufacturer/supplier.[11] For example, one of the products is Carvedilol (substance) 25 milligrams (dosage) in the form of a tablet (dosage form). In addition to the three variables, the list contains the SES tender number and the date on which the tender notification was made public for each product.

3.2. Record Linkage

Given that there are no barcodes in the list of treated products, we resorted to a fuzzy record linkage to identify the set of treated products in the NF-e data. We exploited the text description of each product in both the e-invoice and the list of treated product data to perform the merge. In the latter, the description of each product refers to the triple substance-dosage-dosage form.

3.3. Sample Selection and Data Preparation

Our sample is composed of e-invoices relative to 1,028 medicines purchased between January 2013 and December 2015 by both SES and 370 other public procuring agencies located in Rio Grande do Sul. The list of other procuring agencies includes any state or municipal public bodies that bought medicines in the period 2013-2015, such as state hospitals, intermunicipal health

[11] The list does not refer to specific manufacturers, brands, or product barcodes to avoid favoring any particular supplier.

consortia, municipal health funds. and municipal pharmacies.[12]

Our main data set is a panel of product-buyer combinations over 36 year-months. Each product is uniquely identified by its barcode and each buyer by a different taxpayer identification number. Since the identification of products relies on barcodes, we only considered the invoices for which the suppliers informed the GTIN of the medicines purchased. The original NF-e data set, in contrast, is such that each observation corresponds to a given invoice relative to a purchase of medical drugs by a local public body and contains all product-buyer-supplier combinations for each year-month. It may thus contain more than one invoice issued in the same year-month for a given product-buyer combination.

To transform the original NF-e data into our main data set, we proceed as follows. First, we calculate a weighted average of the unit price of each product-buyer combination in each year-month. This average puts more weight on purchases whose total quantity is larger for the same product-buyer combination. This gives us a monthly average of the unit purchase price of a given product-buyer combination. The main rationality for using the weighted average is that larger acquisitions are more likely to reflect real trends and are also a more faithful representation of the purchasing behavior of public buyers. Moreover, since we cannot identify whether a given invoice transaction refers to a purchase relative to an administrative or judicial proceeding, we believe that larger acquisitions are also more likely to reflect the former, which is exactly the type of purchase that is more appropriate to consider for impact evaluations purposes.[13] Then, we separately calculate a simple average of the total number of purchased units of each product in each year-month. Second, we perform a top and bottom within-product 5 percent winsorization in the unit price to deal with both extremely high and low values.[14] The winsorization aims to avoid having outliers drive the analysis results. At this step, the resultant NF-e data set contains all product-buyer combinations relative to each year-month. It thus includes some buyers who

[12] The list also includes some low frequent buyers of medicines such as the state police.

[13] Judicial proceedings typically result in higher unit prices because procurement officers are forced to accelerate the purchase process in view of deadlines established by judicial courts.

[14] For each product, we winsorized the unit price at the 5th and 95th percentile across different year-months, buyers, and suppliers.

procured a given product only once in the period 2013-2015, and some medicines acquired solely by buyers other than SES.

To ensure a proper comparison and mitigate the risk of compositional changes, we only kept product-buyer combinations with at least one observation in both the pre- and post-treatment periods and which are relative to products acquired by both SES and at least one control buyer. This condition does not lead us to have a balanced year-month panel of product-buyer combinations because the public procuring agencies usually do not purchase the same medicine every month, but rather in lots during some periods of the year. However, and importantly for our analysis, pharmaceuticals are frequently purchased due to storage-, logistic-, and expiration date-related issues. Therefore, even though we do not observe a given product-buyer combination at every year-month, such product-buyer combination will probably be in our data set both in the pre- and post-treatment period if the product is frequently acquired by the corresponding buyer.

4 Descriptive Statistics and Empirical Strategy

4.1. Descriptive Statistics

Table 1 presents some descriptive statistics for the period 2013-2015. Panel A shows the summary statistics relative to the unit price and total quantity, while panel B presents the characteristics of the product-buyer combinations. Panel C depicts the number of relevant groups in our sample. There are 1,028 products, out of which 84 are treated, 371 different buyers including SES, and 278 distinct suppliers. The number of product-buyer combinations is 9,305.

The average unit price is BRL 30.6 with a median of less than BRL 1. The difference between the minimum (BRL 0.02) and the maximum price (BRL 9,076) is huge, as demonstrated by the large standard deviation. The average of the monthly total quantity is 3,949 with a median of 30, a minimum of one, and a maximum of 824,000 units. This massive dispersion in the distribution of both the monthly average unit price and total quantity reflects both the heterogeneity across and within (i.e., over time) the product-buyer combinations that make up our sample. The average number of suppliers of the product-buyer combinations is 2.1, with a median of 3 and a maximum of 9. On average, each product-buyer combination is observed in nearly three pre- and post-treatment year-months.

Table 1 — Summary statistics of the sample

	Mean	Standard Deviation	Median	Min	Max	N
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Characteristics of the main data set						
Unit price	30.58	272.72	0.96	0.024	9,076	71,043
Total quantity	3,949	35,756	30	1	824,000	71,043
Panel B. Characteristics of the product-buyer combinations						
Average no. of suppliers	2.10	1.35	2	1	9	9,305
Average no. of pre-treatment year-months	3.08	3.92	3	1	18	9,305
Average no. of post-treatment year-months	3.55	3.54	2	1	18	9,305
Panel C. Number of relevant groups in the study						
Number of products	1,028					
Number of treated products	84					
Number of buyers	371					
Number of suppliers	278					

Notes: prices are expressed in Brazilian reais (BRL). In panel A, the summary statistics were calculated using the product-buyer-year-month level data set, while in panel B the summary statistisc were calculated considering the product-buyer level data set (which is collapsed from the product-buyer-year-month level data set). The number of products is calculated as the number of distinct barcodes (GTINs).

Table 2 shows some characteristics of the sample separately by buyer group and year. We separate buyers into four groups: (i) SES; (ii) municipalities, which encompass all municipal procuring units located in Rio Grande do Sul; (iii) hospitals, which gathers all public procuring units linked to local hospitals; and (iv) other state buyers, which is a set of five state procuring units that purchase medicines less frequently than other groups: state police, the justice court, the superintendency of penitentiary services, the state special protection fund, and the state socioeducational fund.

Table 2 — Characteristics of the study sample by buyer group and year

	Treated buyer	Control buyers		
	SES	Municipalities	Hospitals	Other state buyers
		2013		
Number of buyers	1	355	9	5
Number of products	1,023	962	177	82
Number of treated products	83	79	20	13
Number of suppliers	60	160	40	24
Total quantity	4,201,798	85,428,106	355,896	142,618
Observations	4,364	18,634	1,240	182
		2014		
Number of buyers	1	353	8	5
Number of products	919	962	183	75
Number of treated products	67	79	22	10
Number of suppliers	59	153	50	29
Total quantity	4,585,456	111,838,867	387,734	318,179
Observations	4,040	20,861	1,373	171
		2015		
Number of buyers	1	343	9	5
Number of products	697	935	172	74
Number of treated products	67	79	22	10
Number of suppliers	66	143	47	25
Total quantity	2,395,074	70,386,259	400,587	143,012
Observations	2,083	16,901	1,164	166

Notes: the summary statistics for each variable were calculated using the product-buyer-year-month level data set, except for the number of suppliers, which was calculated using the invoice-level data set. Total quantity indicates the number of units purchased. Municipalities includes intermunicipal health consortia and other municipal institutions. Other state buyers is a set of five state buyers: the state police, the justice court, the superintendency of penitentiary services, the state special protection fund, and the state socioeducational fund.

SES is, by far, the largest individual buyer of medical drugs in Rio Grande do Sul. Large municipalities and intermunicipal groups, which cover approximately 350 cities in our sample - out of the 497 in the state - are the most similar to SES in terms of the number of purchased products. Table 3 presents the 10 largest buyers of treated products in the period 2013-2015. Not surprisingly, SES is the largest with 94 treated products, while the second-largest is the IPAM pharmacy. IPAM is linked to the municipality of Caxias do Sul, the second-largest city in the state. Out of the other eight buyers, four are specific municipal administrations, three are intermunicipal health consortia, and one is a hospital located in the state capital.[15] Such buyers are very different from SES and the IPAM pharmacy, as they operate on a much smaller scale. The third-largest buyer place, for example, procured 29 treated products, while the tenth-largest one purchased only 15 treated products.

Table 3 — Top 10 buyers of treated products in Rio Grande do Sul in the period 2013-2015

ID	Buyer	No. of treated products	No. of products
1	State Secretariat of Health (SES)	84	1,028
2	IPAM Pharmacy of Caxias do Sul	64	820
3	Intermunicipal Health Consortium of the Northwest of RS	29	275
4	Municipal government of Marau	19	135
5	Intermunicipal Health Consortium of the Cai River Valley	18	134
6	Intermunicipal Health Consortium (CONISA)	17	99
7	Municipal government of Caxias do Sul	17	146
8	Clinical Hospital of Porto Alegre	16	164
9	Municipal government of Rio Grande	16	102
10	Municipal government of Carlos Barbosa	15	95

Notes: the number of products corresponds to the number of distinct barcodes. Caxias do Sul, Marau, Porto Alegre, Rio Grande, and Carlos Barbosa are municipalities of the state of Rio Grande do Sul. 'RS' stands for Rio Grande do Sul.

Table 4 presents the mean and standard deviation of some variables separately for treated and control products in the pre-treatment period. Out of the four variables, three are different among treated and control products: the unit price, the number of buyers, and the number of suppliers. Treated products are, on average, cheaper, purchased by a larger number of distinct buyers, and sold by a larger number of suppliers. Such results indicate that treated products are common medicines among public buyers in Rio Grande do Sul. In this direction, Appendix Table A1 shows that the most purchased treated products are well-known medicines such as Fluoxetine (for depression), Depakene (for bipolar disorder), Diazepam (for anxiety), and Isosorbide (for heart diseases). The average monthly unit price for such products is BRL 10.6, while that for control products is BRL 53.4. From Table 1, we know that the average price of the products is BRL 30.6, which is substantially higher than 10.6 because the treated products account for a minor share of the overall number of products (8.1 percent). The average number of buyers, in turn, is

[15] In Brazil, health intermunicipal consortia are frequently used for the centralized procurement of drugs to enhance public buyers' bargaining power, as combining their demand may allow them to extract lower prices.

18.7 against 8.2 among control products. The average number of suppliers is slightly higher for treated products: 7.8 versus 6.7 for control ones. The differences are all statistically significant at the 5 percent level.

Table 4 — Mean and standard deviation of selected variables in the pre-treatment period, separately for treated and control products

	Treated products	Control products	Difference/p-value
	(1)	(2)	(3)
Unit price	10.62	53.47	-42.85
	[32.890]	[384.938]	[0.001]
Total quantity	355,044	129,149	225,895
	[2,127,217]	[2,230,016]	[0.351]
Number of buyers	18.69	8.19	10.49
	[38.564]	[16.166]	[0.013]
Number of suppliers	7.87	6.74	1.12
	[4.411]	[4.432]	[0.024]
Observations	84	944	

Notes: prices are expressed in Brazilian reais (BRL). To produce the numbers relative to the unit price, total quantity, and number of buyers, we first collapsed the product-buyer-year-month level data set to another one in which each observation represents a given product. In particular, we calculated the average of the corresponding variable for each product across the different year-months of the pre-treatment period, and then calculated the average separately for treated and control products. In contrast, for the number of suppliers, we had to use the invoice-level data set. We first calculated the number of suppliers for each product and then calculated the average separately for treated and control products.

4.2. Empirical Strategy

Triple difference design. Ideally, the only difference between treated and control products should be their treatment status. Nevertheless, in our setting, the selection of treated products was not random. Rather, SES officials chose the set of treated products. Although they aimed to cover different types of products, the characteristics of treated and control products are not similar, as shown in Table 4. Therefore, our identification strategy will not rely on a baseline balance, but on the construction of a difference-in-differences estimator that compares the same product-buyer combinations over time.

To estimate the causal effect of interest, we propose a triple difference (DDD) design that compares the double difference in the average purchase prices between treated and control products for SES with the same double difference for control buyers.[16] The main identification assumption of the DDD is that the difference in the average purchase prices between treated and control products for SES was moving in parallel with the same difference for control buyers in the pre-treatment period. Appendix Figure A1 provides evidence that supports the validity of this assumption.[17] Our main outcome variable is the unit purchase price, but we are interested in estimating the percentage change in the average unit purchase price over time for different groups.[18] Therefore, we use the inverse hyperbolic sine (IHS) transformation of the unit price as an approximation for the percentage change. Another important reason for using the IHS transformation rather than the price in levels is that the former provides easy-to-compare point estimates. In particular, we estimate the coefficients of the following equation

$$\begin{aligned} Y_{i,b,t} \quad = \quad & \beta_0 + \beta_1 (SES_b \times Post_t \times Treated_i) + \beta_2 Post_t \times Treated_i + \\ & + \beta_3 SES_b \times Post_t + \alpha_{i,b} + \gamma_t + \varepsilon_{i,b,t} \end{aligned} \tag{1}$$

where i represents a given product, b a given buyer, and t a given year-month. $Y_{i,b,t}$ is one of the outcomes of interest, SES_b is a dummy variable for whether the buyer is SES, $Post_t$ is a dummy variable for the post-treatment period (i.e., between July 2014 and December 2015), $Treated_i$ is a dummy variable for whether the product is treated, $\alpha_{i,b}$ is a product-buyer fixed effect, γ_t is a year-month fixed effect, and $\varepsilon_{i,b,t}$ is an error term. The product-buyer and year-month

[16] Our strategy allows us to discard potential effects of concomitant initiatives affecting all procurement processes at SES. This is the main reason to avoid implementing a difference-in-differences strategy comparing only across products and time.

[17] Appendix Figure A1 exhibits the point estimates and both the 90 percent and 95 percent confidence intervals associated with the triple interaction coefficient for different placebo starting year-months for the treatment, considering the data relative to the pre-treatment period. We restrict the placebo starting year-months to those between July and December 2013 in order to have at least six months before and after each placebo starting year-month (recall that the original pre-treatment period is between January 2013 and June 2014). The figure provides some evidence that the differences between treated and control products for SES and control buyers were moving in parallel trends before SES started to receive the data-driven reference prices for treated products. The point estimates for the placebo year-months are all positive and statistically not significant, except for the one that considers September 2013 as the start of the treatment period (which is marginally significant at the 10 percent level).

[18] We do not correct the prices for inflation because a uniform correction would not change the differences across products and buyers.

fixed effects play a major role in our identification strategy. The former ensures that we are comparing only the same product-buyer combinations over time, while the latter allows us to control for price seasonality. Standard errors are clustered at the product-buyer level because this is the level at which the variation in the treatment takes place (Abadie et al., 2023). Our interest is identifying β_1, the causal effect of providing the market-based reference prices for SES procurement officers under the triple difference assumptions. The triple difference specification allows us to determine whether the intervention had an average impact on all treated products.

Heterogeneous effect design. To explore whether the impacts of the intervention are greater for a subset of products, we propose a quadruple difference specification. We are interested in determining whether the effect is different among those products with a greater ex-ante unit price. These are the products for which the asymmetry of information is likely to be larger. To this end, we estimate the following equation:

To this end, we also estimate the following quadruple difference specification

$$\begin{aligned} Y_{i,b,t} = \; & \beta_0 + \beta_1(SES_b \times Post_t \times Treated_i) + \\ & + \beta_2(SES_b \times Post_t \times Treated_i \times Above50th\, percentile_i) + \\ & + \sum_{i=3}^{5} \beta_i \cdot (3\, Triple\, Interactions) + \sum_{j=6}^{11} \beta_j \cdot (6\, Double\, Interactions) + \\ & + \alpha_{i,b} + \gamma_t + \epsilon_{i,b,t} \end{aligned} \tag{2}$$

where $Above50th\, percentile_i$ is a binary variable for when the value of product i in the pre-treatment period is greater than the value associated with the 50th percentile of the distribution of the monthly average unit price, and $\epsilon_{i,b,t}$ is an error term. Our main interest in equation 2 is β_2, which gives the difference in the effects between the 50th most and least expensive treated products. Standard errors in equation 2 are also clustered at the product-buyer level.

5 Results

This section shows the results of the estimation of equations 1 and 2 and is divided into three parts. The first part is reserved for the main results, the second part for some additional outcomes, and the third part for robustness checks.

5.1. Main Results

Panel A of Table 5 shows the estimates for the triple difference specification, while panel B shows the results for the heterogeneous effect by ex-ante price group. In column 1, we present the estimates with no control for outliers. In column 2, in contrast, we trim the product-buyer combinations that were found to be the top and bottom 1-percent outliers in terms of the percentage change in price (considering the pre- and post-treatment periods). The point estimates in panel A are all negative and statistically significant. This suggests that SES has benefited from the intervention when we look at the average purchase prices of all treated products. Panel B indicates, nonetheless, that the effect is entirely driven by the most expensive treated products. It is interesting to note that the point estimate for the least expensive treated products is positive (although statistically not significant).

Panel A (Table 5) shows that the provision of market-based reference price has led SES to an average reduction of 4.1 to 4.3 percentage points in the purchase price of treated product. Interestingly, as we can see in Panel B (Table 5), the intervention has led to an average reduction of 13.2 (p-value = 0.003) to 17.2 percentage points (p-value = 0.004) in the purchase price of the 50% most expensive treated products in the study period, but had no impacts on the 50% least expensive treated products. We can also see that the average price of the 50% most expensive treated products (BRL 15.7) is more than twice as large as the average price of treated products (BRL 6.2) in the pre-treatment period. Importantly, the results also indicate that outliers are not driving our findings.

Appendix Table A2 complements the analysis by showing that the difference in the average

Table 5 — Effects on purchase prices

Dependent variable:	IHS(Price)	IHS(Price)
	(1)	(2)
Panel A. Impact on all treated products		
SES x Post x Treated	-0.041*	-0.043**
	[0.023]	[0.018]
Number of clusters	9,305	9,117
R2	0.996	0.997
Observations	71,043	69,991
Panel B. Heterogeneous impact		
(a) SES x Post x Treated	0.015	0.001
	[0.017]	[0.015]
(b) SES x Post x Treated x Above 50th percentile	-0.187***	-0.132***
	[0.060]	[0.048]
Number of clusters	9,305	9,117
R2	0.996	0.997
Observations	71,043	69,991
Avg. price of treated products for SES	6.312	6.195
Avg. price of the 50% most expensive treated products for SES	15.985	15.709
(a) + (b)	-0.172	-0.132
p-value of test (a) + (b) = 0	0.003	0.004
Trimming the top and bottom 1% outliers	No	Yes

Notes: the dependent variable in columns 1 and 2 is the inverse hyperbolic sine (IHS) transformation of the monthly average unit purchase price. In column 2, we removed the product-buyer combinations of the top and bottom 1 percent outliers for the percentage change in price over time. To do so, we first calculated the pre- and post-treatment averages of the monthly average unit price for each product-buyer combination, and then calculated the percentage change considering such averages. All regressions include fixed effects for product-buyer and year-month combinations. Robust standard errors in brackets are clustered at the product-buyer level. Above 50th percentile is a binary variable for when the value of a given product in the pre-treatment period is greater than the value associated with the 50th percentile of the distribution of the monthly average unit price. The average prices of treated products depicted in the table refer to the pre-treatment period. *** p<0.01, ** p<0.05, * p<0.1

price between treated and control products has reduced for both SES and control buyers in the post-treatment period. The difference diminished by 9.7 p.p. for SES and 5 p.p. for control buyers. Moreover, we observe in Appendix Table A3 that the reduction in the price of approximately 17 p.p. for the 50% most expensive treated products is a result of two distinct movements: (i) an 8.7-percentage point reduction (p-value = 0.127) for the most expensive treated products and (ii) an 8.2-percentage point increase (p-value < 0.001) for the most expensive control products. In particular, the only negative point estimate in Appendix Table A3 is the one relative to the most expensive treated products. It is also worth mentioning that the price of both the least expensive treated and control products increased more for SES than control buyers: 7.7 p.p. (p-value < 0.001) and 5.9 p.p. (p-value < 0.001), respectively. This explains the positive estimate of 0.015 in column 1 of Table 5.

Table 6 shows the results of the heterogeneous analysis in multiple dimensions besides the ex-ante price level. In particular, we test for each variable of Table 4 - namely the number of suppliers, number of buyers, and total quantity - whether the effects are larger for the above-median products in terms of their position in the corresponding pre-treatment variable distribution. All groups are defined using transaction characteristics relative to the ex-ante period. To facilitate the comparison of patterns across characteristics, we repeat in column 1 the estimate relative to the price from Panel B of Table 5. Column 2 shows the estimates for the number of suppliers, column 3 for the number of public buyers, and column 4 for the total quantity.

Interestingly, we observe that the three tested features have an important influence on the treatment effects. In column 2, for example, we observe that the intervention effects are statistically indistinguishable from zero for products sold by a larger number of suppliers (p-value = 0.969), but are statistically significant for those with a number of suppliers below the median (p-value = 0.019). The same pattern applies to columns 3 and 4. Thus, we conclude that the intervention effects are concentrated on treated products that are sold by a smaller number of suppliers and purchased by a smaller number of public buyers, but also on a smaller scale. The heterogeneity analysis suggests that the gains in efficiency are driven by having access to unit price information on products that are bought less frequently.

5.2. Additional Results

We now turn to other procurement dimensions, such as the volume purchased, the share each product represents in the total expenditure of each buyer, and the buyer-supplier relationships.

Volume purchased and share in the total expenditure. Table A4 shows the estimation results

Table 6 — Heterogeneous effects on purchase prices by product characteristic in the pre-treatment period

Dependent variable:	IHS(Price)	IHS(Price)	IHS(Price)	IHS(Price)
Product characteristic:	Price	No. of suppliers	No. of buyers	Total quantity
	(1)	(2)	(3)	(4)
(a) SES x Post x Treated	0.015	-0.104**	-0.117***	-0.240***
	[0.017]	[0.044]	[0.043]	[0.082]
(b) SES x Post x Treated x Product Characteristic (Above median)	-0.187***	0.103**	0.113**	0.243***
	[0.060]	[0.050]	[0.050]	[0.084]
Number of clusters	9,305	9,305	9,305	9,305
R2	0.996	0.996	0.996	0.996
Observations	71,043	71,043	71,043	71,043
(a) + (b)	-0.172	-0.001	-0.004	0.002
p-value of test (a) + (b) = 0	0.003	0.969	0.857	0.893

Notes: the dependent variable in all columns is the inverse hyperbolic sine (IHS) transformation of the monthly average unit purchase price. All regressions include fixed effects for product-buyer and year-month combinations. Robust standard errors in brackets are clustered at the product-buyer level. Product Characteristic (Above median) is a binary variable for whether the value associated with a given product is above the median of the corresponding characteristic in the corresponding distribution in the pre-treatment period. The product characteristics refer to those listed in the column headers. *** p<0.01, ** p<0.05, * p<0.1

for the first two outcome variables. Panels A and B follow the same logic as before. In column 1, we show the estimates for the volume purchased, and in column 2 those for the share that each product represents in the total expenditure of each buyer. Panel A indicates no difference between SES and control buyers in terms of the total number of units purchased. The point estimate for the most expensive treated products is negative, while that for the least expensive treated products is positive, but both are not statistically significant. Conversely, the point estimates in column 2 for the most and the least expensive treated products are not different from zero, but the difference between them is statistically significant (p-value = 0.075). This result reflects a mechanical effect: the share reduced relatively more for the most expensive treated products than for the least expensive treated ones because the price was reduced more for the former while the volume purchased remained the same.

Buyer-supplier relationships. Table A5 extends the investigation by focusing on the creation and destruction of relationships with suppliers. Column 1 restricts the sample to relationships that existed in the pre-treatment period and indicates whether such relationships persisted in the post-treatment period. The dependent variable is a dummy variable that assumes 1 if the relationships continued and 0 otherwise. Therefore, a negative point estimate points to the destruction of buyer-supplier partnerships. Similarly, column 2 restricts the sample to relationships that existed in the post-treatment period and tests whether they also existed in the pre-treatment period, which means that a positive point estimate suggests the creation of new buyer-supplier relationships. Panel A of column 1 indicates that the intervention induced SES to destruct more relationships with suppliers of treated than control products and that such destruction was 10.3 p.p. (p-value = 0.025) more likely to happen with SES than control buyers. Panel B demonstrates that this effect is mainly driven by the least expensive treated products. Similarly, the estimates from column 2 indicate that the creation of new supplier relationships was higher for SES than control buyers for the treated products, and was also concentrated in the least expensive treated products. Altogether, these findings suggest that the observed reduction in the price of the most expensive treated products for SES is not a consequence of changes in the supplier-buyer relationships.

5.3. Robustness Checks

In this last part of Section 5, we present a series of robustness checks that test the sensitivity of our results to variations related to the data aggregation level, the choice of the Nth moment of the pre-treatment price distribution, the use of placebo buyers or placebo treatment starting period, and the use of distinct sets of fixed effects in our regressions.

Data aggregation level. Table A6 presents estimates considering different data aggregation levels: year-month, year-quarter, year-semester, and before-after (i.e., only comparing the pre- and post-treatment periods). Panel A is relative to the triple difference design, while panel B refers to the quadruple difference one. The point estimates for the year-monthly data set are the more conservative in terms of statistical significance in both panels. Moreover, the magnitude seems to increase when we consider more aggregated levels. Therefore, our results do not seem to be driven by the granularity level of the data.

Choice of the Nth moment of the pre-treatment price distribution. One might consider that the results are sensitive to the choice of the Nth moment of the distribution of the monthly average purchase price in the pre-treatment period. To shed some light on this issue, we present in Figure A2 both the point estimates for the quadruple difference design and the corresponding 90 and 95 percent confidence intervals considering different moments of such distribution. The results indicate that the effects for the Nth most expensive treated products are even larger in magnitude for the percentiles between 51 and 80, which is consistent with the assumption that the asymmetry of information is greater for such products.

Placebo buyers. Figure A3 shows the point estimates and the 90 and 95 percent confidence intervals of the triple interaction coefficient for different placebo buyers, considering the top nine control buyers presented in Table 3. Each coefficient corresponds to a different regression that excludes SES and considers as the unique placebo-treated buyer one of the nine buyers other than SES. There are only two placebo buyers for which the point estimates are more negative than that for SES: CONISA (ID number 6) and the municipal government of Rio Grande (ID number 9). Nevertheless, none of them are statistically significant (p-values equal to X and Y, respectively), differently from the effect for SES (p-value = 0.078).

Placebo treatment starting year-months. Appendix Figure A4 presents the point estimates and both the 90 and 95 percent confidence intervals for the coefficients of the triple difference design when we set the treatment starting year-months to year-months other than July 2014. As explained before, the tenders for the first products for which the market-based reference prices were provided for SES took place in July 2014. Therefore, one could ask whether the conclusions would change should a different treatment starting year-month have been chosen. The findings indicate that even more negative point estimates by the end of 2014 and the beginning of 2015, which is consistent with the fact that the invoices from bid winners of auctions held close to July 2014 are issued only after the tenders are concluded. Further, the point estimates from regressions whose treatment starting period is mid-2015 remain negative and statistically significant at the 5 percent level.

Different sets of fixed effects. Finally, we also test whether the negative impacts on price are sensitive to the set of fixed effects we consider in the regressions. Appendix Table A7 demonstrates that the results remain robust when we replace the year-month fixed effects by either year-quarter or year-semester ones. The same applies if we include buyer-year-month, buyer-year-quarter, or buyer-year-semester fixed effects instead of year-month ones. The point estimates are all negative and statistically significant across specifications, and the magnitude varies only slightly. We thus conclude that our results are not driven by an arbitrary choice of the relevant fixed effects.

5.4. Back of the Envelope Calculation of Intervention Benefits

Access to up-to-date local market information led to a reduction of 4 percentage points in purchase prices, as shown in Table 5-Panel A. As a result, the average change in purchase prices for treated products was 4 percentage points lower after the intervention was implemented in comparison to the trend of control products. Table A3 demonstrates that instead of increasing by 7.1%, SES-treated products only increased by 3.2%. The annual savings generated by SES in each product is equal to 4.1% of the average ex-ante price multiplied by the average quantity purchased for each product in the period 2013-2015. Finally, SES savings is the sum of the savings in all the purchased products.[19]

In the period of the study, SES generated savings of US$ 11.4 thousand in the purchase of treated products. Assuming that 4 percentage points is a conservative extrapolation for all products bought by SES, an envelope calculation leads to potential annual savings of USD 723,9 thousand. On average, the State Secretariat of Health (SES) in Rio Grande do Sul spent USD 18.5 million annually on pharmaceuticals in the period 2013-2015.[20] Thus, savings are 4% of the

[19] Note that 4 percentage points is a lower bound for the effect of the intervention. Both, Table A6 and Figure A4 show that the effect could be up to 8 percentage points. Thus, the benefits generated by the program are also a lower bound.

[20] In Brazilian Reais, the average annual expenditure is $R 48.1 millions. For conversion to US dollars ($USD), we are using the average exchange rate for the period 2013-2015: $R 2.6 per $USD (World Bank data lake: https://data.worldbank.org/).

average annual total expenditure in the set of medicines considered in our study.[21]

6 Conclusion

This paper estimates the effects of providing reference prices based on local B2B market prices relative to some medical drugs for a large state procuring agency in a developing country. We study the case of the State Secretariat of Health of Rio Grande do Sul, Brazil, between January 2013 and December 2015 using e-invoicing data about the universe of local medicines acquisitions made by state and municipal public bodies. The market-based reference prices were calculated by an algorithm developed by the State Secretariat of Finance and then distributed to the State Secretariat of Health to be used in tenders carried out as of July 2014. Our main interest is to determine whether the intervention has led SES to monetary savings, which is arguably one of the most important dimensions of the public procurement efficiency of off-the-shelf goods such as pharmaceuticals.

We find that the purchase unit price of the relatively most expensive treated products increased less for the SES than other state and municipal buyers of the same products. We find no evidence that the same occurred for the least expensive treated products. We argue that the asymmetry of information between the SES procurement officers and the potential suppliers is more likely to be greater for the relatively most expensive treated products. If this is the case, our results are consistent with the effects being greater and statistically significant only for the products for which one would expect a greater asymmetry of information in the pre-treatment period. When exploring potential mechanisms driving our results, we show that SES broke up more relationships with suppliers than control buyers, and also created more new ones for the treated products. Interestingly, the effects for the least expensive treated products are statistically significant, while those for the most expensive treated ones are not, which raises the question of what should have channeled the effects to the purchase unit prices. Anecdotal evidence

[21] Dezolt and Muñoz (2020) shows that this intervention can lead to savings of up to 10% of total expenditure in the purchase of pharmaceuticals.

from SES officials indicates that the mechanism might be related to increased bargaining of the procurement officers during the price negotiation phase of the tender p rocess. It could also be related to a better establishment of maximum prices in the auction of such products. Unfortunately, our data do not allow us to pin down the underlying drivers of the observed reduction in price, nor to rule out alternative channels.

Our findings suggest, however, that the adoption of new technologies in the context of public procurement can contribute to improving its efficiency through the mitigation of market failures related to asymmetric information. This finding has wide-ranging policy implications given the recent evidence that regulation and enforcement may not translate into efficiency gains. The intervention of interest provided actionable information at a relatively low cost and required minimal effort from procurement officers to make use of it. Interestingly, the savings were large and benefits in terms of lowered purchase prices may continue accruing over time with limited maintenance investment. Although we are unable to gather information on all involved costs, it may be the case that such type of light-touch intervention is cost-effective and even more cost-effective than those focusing on the enhancement of bureaucratic competence.

We believe that our setting is comparable to those of several other developing countries in which the procurement of medicines is centralized and executed by large public procuring agencies. Therefore, in terms of external validity, we consider that our results are not exclusively a consequence of setting-specific features and may apply reasonably well to other contexts.

References

Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge. 2023. "When should you adjust standard errors for clustering?" *The Quarterly Journal of Economics* 138 (1): 1–35.

Arvate, Paulo, Klenio Barbosa, and Eric Fuzitani. 2016. "Party expertise, campaign donation and government contracts: Evidence from an electoral quasi-experiment." *Unpublished working paper.*

Baltrunaite, Audinga, Cristina Giorgiantonio, Sauro Mocetti, and Tommaso Orlando. 2021. "Discretion and supplier selection in public procurement." *The Journal of Law, Economics, and Organization* 37 (1): 134–166.

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

Barreix, Daniel, and Raul Zambrano. 2018. *Electronic Invoicing in Latin America.* Volume 595. Inter-American Development Bank.

Best, Michael Carlos, Jonas Hjort, and David Szakonyi. 2022. "Individuals and organizations as sources of state effectiveness." *Working Paper.*

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

Bucciol, Alessandro, Riccardo Camboni, and Paola Valbonesi. 2020. "Purchasing medical devices: The role of buyer competence and discretion." *Journal of Health Economics* 74: 102370.

Calzolari, Giacomo, and Giancarlo Spagnolo. 2006. "Reputation and collusion in procurement." *Manuscript, University of Bologna and Stockholm School of Economics.*

Coulomb, Renaud, and Marc Sangnier. 2014. "The impact of political majorities on firm value: Do electoral promises or friendship connections matter?" *Journal of Public Economics* 115: 158–170.

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

De Michele, Roberto, and Juan Cruz Vieyra. 2022. "From Fishing to Catching: Developing Actionable Red Flags in Public Procurement to Prevent and Control Corruption." *International Journal of Forecasting*.

Decarolis, Francesco. 2014. "Awarding price, contract performance, and bids screening: Evidence from procurement auctions." *American Economic Journal: Applied Economics* 6 (1): 108–32.

Decarolis, Francesco, Leonardo M Giuffrida, Elisabetta Iossa, Vincenzo Mollisi, and Giancarlo Spagnolo. 2020a. "Bureaucratic competence and procurement outcomes." *The Journal of Law, Economics, and Organization* 36 (3): 537–597.

Decarolis, Francesco, and Giuliana Palumbo. 2015. "Renegotiation of public contracts: An empirical analysis." *Economics Letters* 132: 77–81.

Decarolis, Francesco, Giancarlo Spagnolo, and Riccardo Pacini. 2020b. "Past performance and procurement outcomes." *NBER Working Paper No. 22814*.

Dezolt, Ana Lucia, and Andrés Muñoz. 2020. "NF-e para estimativa de preços de referência para compras públicas do setor saúde." *IDB Discussion Paper IDP-DP-826*.

Fazekas, Mihály, and Jurgen Rene Blum. 2021. "Improving Public Procurement Outcomes: Review of tools and the state of the evidence base." *Policy Research Working Paper 9690*.

Fazekas, Mihály, and Bence Tóth. 2017. "The effectiveness of the European Union in safeguarding competition in public procurement markets."

Ferraz, Claudio, Frederico Finan, and Dimitri Szerman. 2015. "Procuring firm growth: The effects of government purchases on firm dynamics."

Gallego, Jorge, Gonzalo Rivero, and Juan Martínez. 2021. "Preventing rather than punishing: An early warning model of malfeasance in public procurement." *International Journal of Forecasting* 37 (1): 360–377.

Gerardino, Maria Paula, Stephan Litschig, and Dina Pomeranz. 2017. "Distortion by Audit: Evidence from Public Procurement." *NBER Working Paper No. 23978.*

Grennan, Matthew, and Ashley Swanson. 2020. "Transparency and negotiated prices: The value of information in hospital-supplier bargaining." *Journal of Political Economy* 128 (4): 1234–1268.

Huber, Martin, and David Imhof. 2019. "Machine learning with screens for detecting bid-rigging cartels." *International Journal of Industrial Organization* 65: 277–301.

Kovalchuk, Artur, Charles Kenny, Mallika Snyder et al. 2019. "Examining the Impact of E-procurement in Ukraine." *Center for Global Development.*

Lauletta, Maximiliano, Martín Antonio Rossi, Juan Cruz Vieyra, and Diego Arisi. 2019. "Monitoring public investment: The impact of MapaRegalías in Colombia."

Lewis-Faupel, Sean, Yusuf Neggers, Benjamin A Olken, and Rohini Pande. 2016. "Can electronic procurement improve infrastructure provision? Evidence from public works in India and Indonesia." *American Economic Journal: Economic Policy* 8 (3): 258–83.

Lichand, Guilherme, Marcos FM Lopes, and Marcelo C Medeiros. 2016. "Is corruption good for your health." *scholar. harvard. edu/glichand/publications/job-market-paper. Acessado em* 25 (04): 2016.

Lotti, Clarissa, and Giancarlo Spagnolo. 2022. "Indirect Savings from Public Procurement Centralization."

Palguta, Ján, and Filip Pertold. 2017. "Manipulation of procurement contracts: Evidence from the introduction of discretionary thresholds." *American Economic Journal: Economic Policy* 9 (2): 293–315.

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

Santos, Carlos, Meirelles Guerro, Greice Leal, José Primo, Lucídio Beserra, Maria do Socorro Silva, and Miriam Borges Santos. 2015. "Aplicação da metodologia do preço de referência para licitações públicas com base na nota fiscal eletrônica." *Mimeo.*

Schøll, Anja, and Doriyush Ubaydi. 2017. "Impact of technology on corruption: A study of impact of e-procurement on prices of various government purchases." Master's thesis.

Singh, Ram et al. 2011. "Determinants of cost overruns in public procurement of infrastructure: Roads and railways." In *India Policy Forum,* 7: 97–158, National Council of Applied Economic Research.

Titl, Vitezslav, and Benny Geys. 2019. "Political donations and the allocation of public procurement contracts." *European Economic Review* 111: 443–458.

A Appendix Figures and Tables

Figure A1 — Effects on purchase prices considering different placebo starting year-months for the treatment, with data restricted to the pre-treatment period

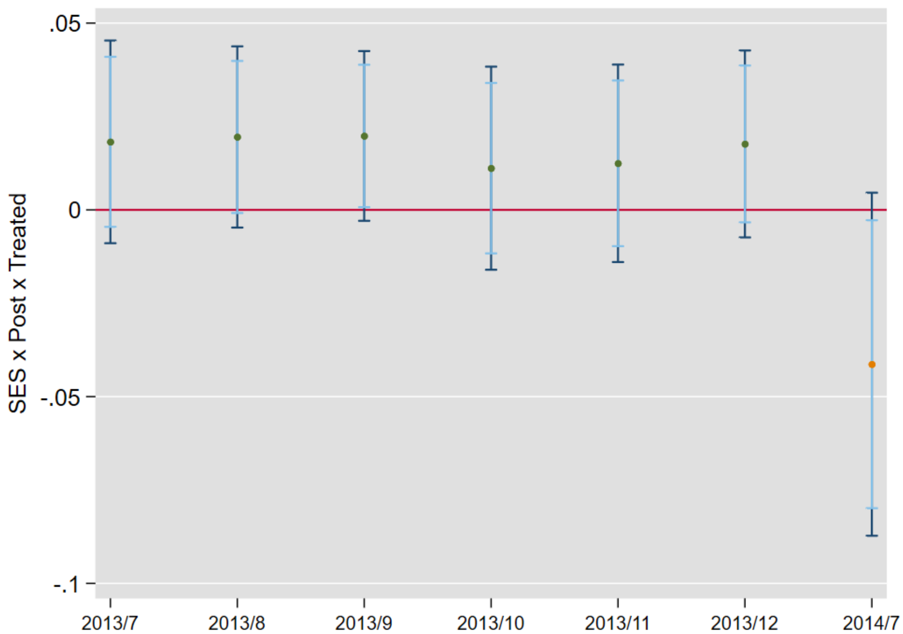

Notes: coefficients are extracted from regressions that include fixed effects for product-buyer and year-month combinations. Each regression considers a different placebo starting year-month for the treatment. The figure should be interpreted as follows. In the x-axis, we show a given placebo year-month, while in the y-axis we show the corresponding point estimate associated with the coefficient of the triple interaction SES x Post x Treated. Post is a different variable for each regression. For example, for the value 2013/7 in the x-axis, Post assumes 1 for year-months between July 2013 and June 2014 (end of the pre-treatment period) and 0 otherwise (i.e., between January 2013 and June 2013). The vertical light and dark blue lines form the 90 and 95 percent confidence intervals. Standard errors are calculated robust and clustered at the product-buyer level. The point estimate in dark blue corresponds to the SES x Post x Treated coefficient from column 1 of Table 5.

Figure A2 — Effects on purchase prices using different moments of the monthly average unit price distribution in the pre-treatment period

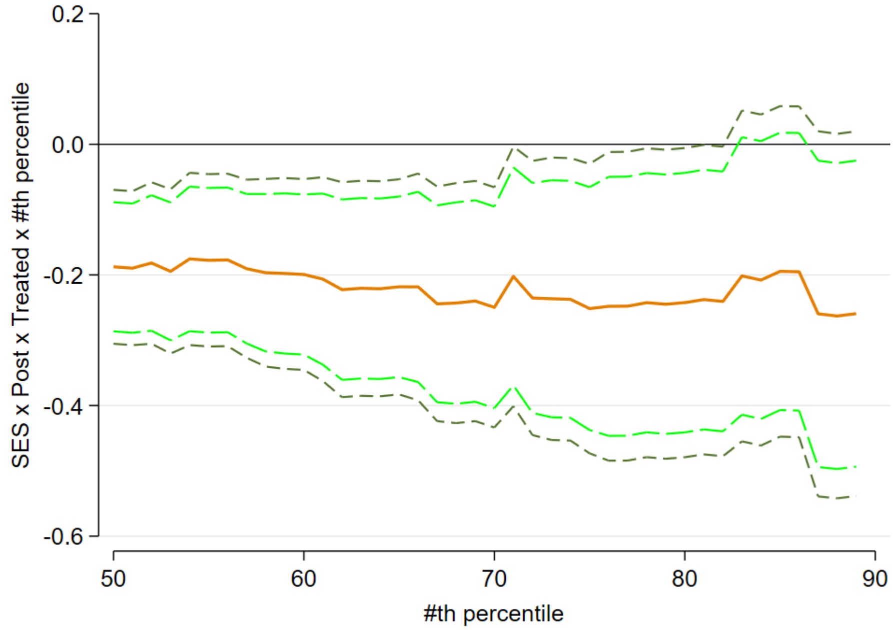

Notes: the solid line shows the point estimates, while the light green and dark green dashed lines form the 90 and 95 percent confidence intervals, respectively. Each estimate is obtained from a specification that follows column 1 of Table 5. #th percentile is a binary variable for when the value of the monthly average unit price of the product in the pre-treatment period is greater than the value associated with the #th percentile where # is an integer value between 50 and 90.

Figure A3 — Effects on purchase prices considering different placebo buyers

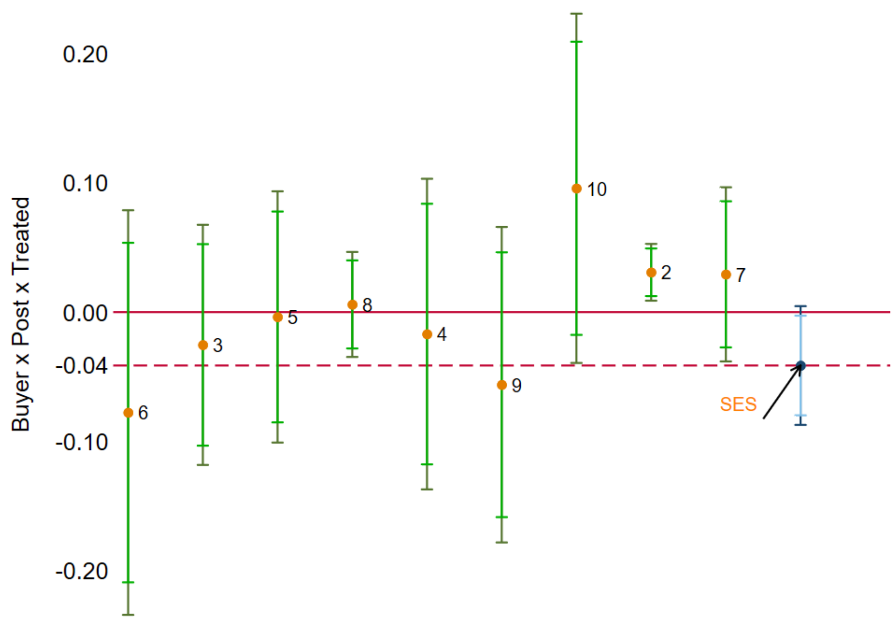

Notes: each point estimate is associated with the coefficient of the triple interaction Buyer x Post x Treated of a regression that considers as the only treated buyer an institution different from SES. The nine buyer institutions other than SES presented in Table 3 were each considered as treated in a separate regression. We removed SES from the sample in all regressions. The labels for each coefficient follow the ID column of Table 3. Coefficients are extracted from regressions that include fixed effects for product-buyer and year-month combinations. Standard errors are calculated robust and clustered at the product-buyer level. The vertical lines form the 90 and 95 percent confidence intervals. The point estimate in dark blue corresponds to the SES x Post x Treated coefficient from column 1 of Table 5.

Figure A4 — Effects on purchase prices considering different placebo starting year-months for the treatment

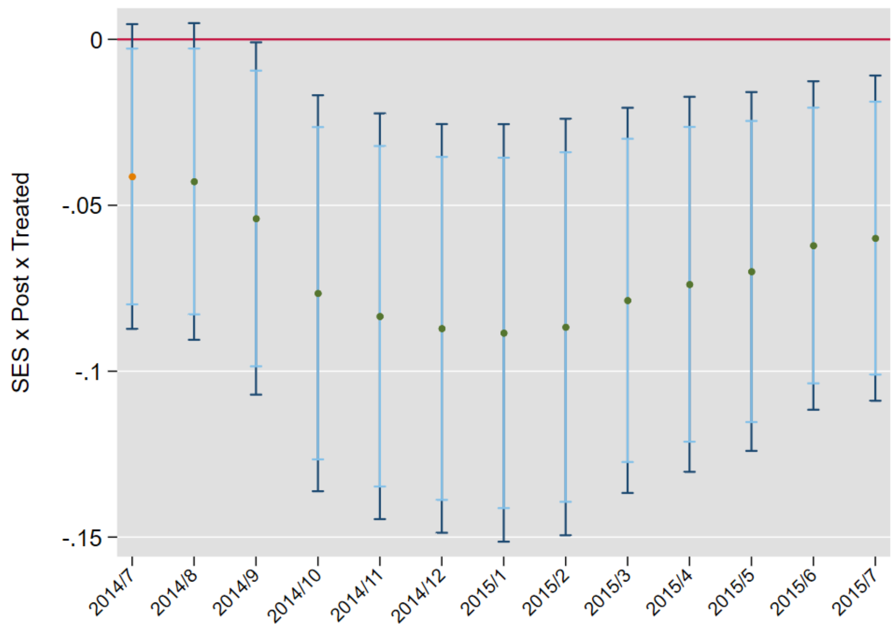

Notes: coefficients are extracted from regressions that include fixed effects for product-buyer and year-month combinations. Each regression considers a different placebo starting year-month for the treatment. The figure should be interpreted as follows. In the x-axis, we show a given placebo year-month, while in the y-axis we show the corresponding point estimate associated with the coefficient of the triple interaction SES x Post x Treated. Post is a different variable for each regression. For example, for the value 2014/8 in the x-axis, Post assumes 1 for year-months between August 2014 and December 2015 and 0 otherwise. The vertical light and dark blue lines form the 90 and 95 percent confidence intervals. Standard errors are calculated robust and clustered at the product-buyer level. The point estimate in orange corresponds to that from column 1 of Table 5.

Table A1 — Most purchased products in the period 2013-2015, separately for all products and treated products only

Substance	Dosage	Total quantity
Panel A. All products		
Omeprazole	20 mg	181,203,600
Metformin	850 mg	87,183,408
Hematofer	40 mg	26,786,892
Amoxicilin	500 mg	23,664,626
Serenata	50 mg	10,270,649
Panel B. Treated products only		
Fluoxetine	20 mg	84,791,232
Depakene	500 mg	18,225,968
Diazepam	5 mg	12,364,701
Prolopa	200 mg	5,426,441
Isosorbide	20 mg	4,581,042

Notes: total quantity is calculated considering all buyers and year-months of the period 2013-2015.

Table A2 — Effects on purchase prices using DD specifications separately for SES and control buyers

Dependent variable:	IHS(Price)
	(1)
Panel A. Difference-in-differences - SES	
Treated x Post	-0.097***
	[0.022]
Number of clusters	1,028
R2	0.996
Observations	10,487
Panel B. Difference-in-differences - Control buyers	
Treated x Post	-0.050***
	[0.006]
Number of clusters	8,277
R2	0.996
Observations	60,556

Notes: the dependent variable is the inverse hyperbolic sine (IHS) transformation of the monthly average unit purchase price. The regression includes fixed effects for product-buyer and year-month combinations. Robust standard errors in brackets are clustered at the product-buyer level. *** p<0.01, ** p<0.05, * p<0.1

Table A3 — Effects on purchase prices using different DD specifications

	All	Most expensive products only	Least expensive products only
	(1)	(2)	(3)
Panel A. Difference-in-differences – Treated products			
SES x Post	0.032	-0.087	0.077***
	[0.023]	[0.057]	[0.018]
Number of clusters	1,570	294	1,276
R2	0.988	0.989	0.908
Observations	10,681	2,163	8,518
Panel B. Difference-in-differences – Control products			
SES x Post	0.071***	0.082***	0.059***
	[0.005]	[0.009]	[0.005]
Number of clusters	7,735	3,009	4,726
R2	0.996	0.994	0.960
Observations	60,362	25,885	34,477
Panel C. Triple difference			
SES x Post x Treated	-0.041*	-0.173***	0.016
	[0.023]	[0.057]	[0.017]
Number of clusters	9,305	3,303	6,002
R2	0.996	0.994	0.948
Observations	71,043	28,048	42,995

Notes: the dependent variable is the inverse hyperbolic sine (IHS) transformation of the monthly average unit purchase price. All regressions include fixed effects for product-buyer and year-month combinations. Robust standard errors in brackets are clustered at the product-buyer level. The most expensive products are those whose average unit purchase price in the pre-treatment period is above median, while the least expensive ones are those below median. *** p<0.01, ** p<0.05, * p<0.1

Table A4 — Effects on total quantity and the share that each product represents in the total expenditure of each buyer

Dependent variable:	IHS(Total quantity)	Share of total expenditure
	(1)	(2)
Panel A. Impact on all treated products		
SES x Post x Treated	-0.051	-0.003
	[0.186]	[0.002]
Number of clusters	9,305	9,305
R2	0.832	0.748
Observations	71,043	18,610
Panel B. Heterogeneous impact		
(a) SES x Post x Treated	0.045	-0.003
	[0.275]	[0.003]
(b) SES x Post x Treated x Above 50th percentile	-0.193	-0.004
	[0.324]	[0.005]
Number of clusters	9,305	9,305
R2	0.833	0.748
Observations	71,043	18,610
(a) + (b)	-0.148	-0.007
p-value of test (a) + (b) = 0	0.388	0.075

Notes: the dependent variable in column 1 is the inverse hyperbolic sine (IHS) transformation of the monthly average total quantity, while that in column 2 is the share that each product represents in the total expenditure of each buyer. Regressions in column 1 include fixed effects for product-buyer and year-month combinations. In column 2, we used a different data set in which each observation is a product-buyer-time combination, where 'time' is either the pre- or post-treatment period. Regressions in column 2 include fixed effects for product-buyer combinations and a binary variable for the post-treatment period. In all regressions, robust standard errors in brackets are clustered at the product-buyer. Above 50th percentile is a binary variable for when the value of a given product in the pre-treatment period is greater than the value associated with the 50th percentile of the distribution of the monthly average unit price. *** p<0.01, ** p<0.05, * p<0.1

Table A5 — Effects on creation and destruction of supplier relationships

Dependent variable: Relationship exists in the period	Destruction of relationships (1)	Creation of relationships (2)
Panel A. Impact on all treated products		
SES x Post x Treated	-0.103**	0.101***
	[0.046]	[0.035]
Number of relationships	14,002	13,867
Number of clusters	9,305	9,305
R2	0.230	0.212
Observations	28,004	27,734
Panel B. Heterogeneous impact		
(a) SES x Post x Treated	-0.114**	0.126***
	[0.058]	[0.041]
(b) SES x Post x Treated x Above 50th percentile	0.062	-0.119
	[0.096]	[0.078]
Number of relationships	14,002	13,867
Number of clusters	9,305	9,305
R2	0.231	0.212
Observations	28,004	27,734
(a) + (b)	-0.051	0.007
p-value of test (a) + (b) = 0	0.506	0.906

Notes: to produce this table, we use the data at the product-buyer-supplier-time level where time refers to either the pre- or post-treatment period. The dependent variable in columns 1 and 2 is a binary variable for whether a given product-buyer-supplier relationship exists in the corresponding time period. In column 1, the sample is restricted to relationships that existed in the pre-treatment period, while in column 2 the sample is restricted to relationships that existed in the post-treatment period. Therefore, in column 1, the dependent variable assumes 1 in the pre-treatment period and might be either 0 or 1 in the post-treatment period, while in column 2, the dependent variable assumes 1 in the post-treatment period, and might be either 0 or 1 in the pre-treatment period. All regressions include fixed effects for product-buyer combinations and a binary variable for the post-treatment period. Robust standard errors in brackets are clustered at the product-buyer level. Number of relationships shows the number of product-buyer-supplier combinations. Above 50th percentile is a binary variable for when the value of a given product in the pre-treatment period is greater than the value associated with the 50th percentile of the distribution of the monthly average unit price. *** p<0.01, ** p<0.05, * p<0.1

Table A6 — Effects on purchase prices using different data aggregation levels

Aggregation level:	Monthly	Quarterly	Semesterly	Before-After
Dependent variable:	IHS(Price)	IHS(Price)	IHS(Price)	IHS(Price)
	(1)	(2)	(3)	(4)
Panel A. Impact on all treated products				
SES x Post x Treated	-0.041*	-0.049*	-0.061**	-0.082***
	[0.023]	[0.026]	[0.027]	[0.028]
Number of clusters	9,305	9,305	9,305	9,305
R2	0.996	0.995	0.995	0.994
Observations	71,043	48,266	34,488	18,610
Panel B. Heterogeneous impact				
(a) SES x Post x Treated	0.015	0.012	0.001	-0.006
	[0.017]	[0.019]	[0.020]	[0.024]
(b) SES x Post x Treated x Above 50th percentile	-0.187***	-0.211***	-0.215***	-0.250***
	[0.060]	[0.063]	[0.063]	[0.066]
Number of clusters	9,305	9,305	9,305	9,305
R2	0.996	0.995	0.995	0.994
Observations	71,043	48,266	34,488	18,610
(a) + (b)	-0.172	-0.198	-0.214	-0.256
p-value of test (a) + (b) = 0	0.003	0.001	<0.001	<0.001

Notes: the dependent variable is the inverse hyperbolic sine (IHS) transformation of the (time) average unit purchase price, where time is either monthly, quarterly, semesterly, or pre- / post-treatment period. All regressions include fixed effects for product-buyer and time combinations, where time is year-month in column 1, year-quarter in column 2, year-semester in column 3, and a binary variable for the post-treatment period in column 4. Robust standard errors in brackets are clustered at the product-buyer level. Above 50th percentile is a binary variable for when the value of a given product in the pre-treatment period is greater than the value associated with the 50th percentile of the distribution of the monthly average unit price. *** p<0.01, ** p<0.05, * p<0.1.

Table A7 — Effects on purchase prices considering different sets of fixed effects

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Impact on all treated products						
SES x Post x Treated	-0.041*	-0.039*	-0.039*	-0.048**	-0.044**	-0.045*
	[0.023]	[0.023]	[0.023]	[0.024]	[0.024]	[0.024]
Number of clusters	9,305	9,305	9,305	8,995	9,222	9,287
R2	0.996	0.996	0.996	0.996	0.996	0.996
Observations	71,043	71,043	71,043	68,624	70,504	70,916
Panel B. Heterogeneous impact						
(a) SES x Post x Treated	0.015	0.017	0.017	0.006	0.011	0.011
	[0.017]	[0.017]	[0.017]	[0.017]	[0.017]	[0.017]
(b) SES x Post x Treated x Above 50th percentile	-0.187***	-0.188***	-0.187***	-0.163***	-0.174***	-0.176***
	[0.060]	[0.060]	[0.060]	[0.062]	[0.061]	[0.061]
Number of clusters	9,305	9,305	9,305	8,995	9,222	9,287
R2	0.996	0.996	0.996	0.996	0.996	0.996
Observations	71,043	71,043	71,043	68,624	70,504	70,916
p-value of test (a) + (b) = 0	0.003	0.003	0.003	0.009	0.006	0.005
Product-buyer FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-month FE	Yes	No	No	No	No	No
Year-quarter FE	No	Yes	No	No	No	No
Year-semester FE	No	No	Yes	No	No	No
Buyer-year-month FE	No	No	No	Yes	No	No
Buyer-year-quarter FE	No	No	No	No	Yes	No
Buyer-year-semester FE	No	No	No	No	No	Yes

Notes: the dependent variable in all columns is the inverse hyperbolic sine (IHS) transformation of the monthly average unit price. Robust standard errors in brackets are clustered at the product-buyer level. Above 50th percentile is a binary variable for when the value of a given product in the pre-treatment period is greater than the value associated with the 50th percentile of the distribution of the monthly average unit price. *** p<0.01, ** p<0.05, * p<0.1