

Should the Government Sell You Goods? Evidence from the Milk Market in Mexico*

Diego Jiménez-Hernández

Enrique Seira

April 19, 2022

[\[click for latest version\]](#)

Abstract

Governments spend considerable resources providing goods directly. We show that such behavior may increase welfare when private suppliers have market power. We do this by studying the staggered rollout of hundreds of government milk “ration stores” in Mexico using a proprietary panel of household food purchases. The rollout lowered the price per liter of privately supplied milk by 2.4% and increased household consumption. To compare direct provision with budget-neutral alternatives, we develop and estimate an equilibrium model of the market that accounts for quality differences. Direct provision generates larger consumer surplus than milk vouchers and unrestricted cash transfers.

JEL: H42, L33, L44, L66, O15.

Keywords: private and public provision; market concentration; milk.

*Jiménez-Hernández (corresponding author): Postdoctoral researcher, Microsoft Research; diego.j.jimenez.h@gmail.com. Seira: Professor, Centro de Investigación Económica, ITAM. Diego is extremely grateful for the guidance, mentorship, and helpful comments provided by Pascaline Dupas, Liran Einav, and Melanie Morten. We wish to thank Tim Bresnahan, Camilo Castillo, Arun G. Chandrasekhar, Jose Ignacio Cuesta, Yiwei Chen, Marcel Fafchamps, Matt Gentzkow, Aprajit Mahajan, Neale Mahoney, Seema Jayachandran, Brad Larsen, Peter Reiss, Mauricio Romero, Meredith Startz, Paulo Somaini, Ali Yurukoglu, and seminar participants at the ITAM Alumni Conference, NEUDC, the Development Tea, and the IO Workshop for helpful comments. We also wish to thank Adalberto Gonzalez Hernandez, James Peter McCarthy, and Rodolfo Camacho Estrada for numerous discussions of the milk market in Mexico. Margarita Khvan, Paola Mejía, and Ricardo Olivares provided outstanding research assistance. Diego gratefully acknowledges funding provided by the Ronald I. McKinnon Memorial Fellowship, Shultz Graduate Student Fellowship in Economic Policy, Stanford King Center on Global Development’s Graduate Student Fellowship, and Leonard W. Ely and Shirley R. Ely Graduate Student Fellowship. Some of the computing for this project was performed on the Sherlock cluster. We would also like to thank Stanford University and the Stanford Research Computing Center for providing both computational resources and support that contributed to these research results. Any errors are our own.

1 Introduction

Governments across the world expend considerable resources providing or subsidizing essential goods and services, from food and public housing to schools and healthcare. In many cases, they provide goods directly, often despite the existence of privately supplied alternatives. An example of this are the “ration stores” widely used in low- and middle-income countries, at which households can purchase staple foods or basic household consumption goods at subsidized prices.¹ Economists often find direct-provision systems of this kind puzzling, as it seems unlikely that the government would have any comparative advantage over the private market in procuring, producing, and distributing goods such as milk or staples, especially when the private sector is already providing them in the area.

Explanations of this phenomenon include better targeting by screening out consumers with lower needs (Blackorby and Donaldson, 1988; Besley and Coate, 1991; Nichols and Zeckhauser, 1982; Dupas et al., 2016; Alatas et al., 2016; Gadenne, 2020); price volatility reduction (Gadenne et al., 2020); paternalism (Cunha, 2014); and political economy considerations (Epple and Romano, 1996; Barse et al., 2000). Its opponents argue that such programs constrain choice and may be weakly dominated by either vouchers (if the goal is to maximize consumption) or unrestricted cash (if the goal is to maximize consumer surplus).² An often ignored aspect is the market power of private providers. Firms in low- and middle-income countries often exercise their ability to price above costs consumed disproportionately by the poor (Goldberg, 2018; Ivaldi et al., 2016). In such cases, the government option acts as an extra competitor, potentially decreasing prices. Cunha et al. (2019) show experimentally that in-kind provision leads to lower prices than in-cash provision, although it could be the case that the opposite occurs if the most elastic consumers substitute toward the government good (Atal et al., 2020).

In this paper, we study the extent to which subsidizing via direct government provision disciplines the market power of private sellers. We use the nationwide provision of subsidized milk via ration stores in Mexico. These stores sell government-produced milk at around one-third of the price of private brands. The setting is ideal, as we have variation induced by the opening of new government stores in areas served by private providers, which we combine with cost shifters that allows us to identify demand. We were able to obtain high-quality proprietary data on household milk purchases for thousand of households (and millions of transactions) and to observe prices, quantities, brands, and product attributes.

Our analysis proceeds in two steps. We begin by using a widespread, staggered introduction

¹ Ration stores exist in Bangladesh, Ethiopia, India, Indonesia, Mexico, Pakistan, and Sri Lanka (Gadenne, 2020). The Indian Public Distribution System serves 65 million individuals, and spends 1% of GDP (Bhattacharya et al., 2017; Nagavarapu and Sekhri, 2016). In Mexico, 7 million individuals benefit from the ration stores we study, and the government spends 0.7% of the GDP on them.

² Atkinson and Stiglitz (1976) show, for instance, that redistribution through income taxes dominates commodity taxation when income is observable, firms are price takers, and consumer preferences are homogeneous .

of more than 1,500 new government stores to document substitution toward government milk and a decrease in the price of private milk in an event study framework. A year after opening, households are 10 percentage points more likely to purchase from the government and 3 percentage points less likely to purchase private milk.³ The increase in purchases is concentrated among poorer households. Turning to prices, we show that the average price per liter for *private* suppliers decreases by 25¢ (2.3% of the mean), and since households partially substitute toward the (cheaper) government good, the price per liter over all suppliers falls by 40¢ (3.8%). We find no evidence of private supplier exit or changes in the number of products offered.

We also exploit the distance of households to the government store—while controlling for variables that are inputs to the government store opening rule—to document being one standard deviation closer in physical distance to the government store is associated with stronger effects: a 13% increase in the monthly consumption of milk, a 3% decrease in the price of private brands, and a 7% decrease in the overall price, which bolsters the argument that these effects are due to store opening. Using other consumption goods (e.g., sodas) as placebos, we show that the above-described patterns are only present in the milk market. That prices in the milk market decrease and quantities increase suggests that consumers may benefit from such an intervention.

In the second step, we quantify the impacts on welfare using a supply and demand model for the milk market that takes into account quality differences across brands of milk as perceived by consumers, differences in production costs between the government and the private sector, as well as supply-side price responses. We compare the direct provision policy with alternative “budget-neutral” policies—counterfactual exercises that, in equilibrium, generate the same government spending as in our observed government-as-a-seller arrangement. The alternatives include (i) giving the equivalent expense as a direct lump-sum cash transfer to consumers and (ii) vouchers to buy milk in the private market.

We follow standard discrete choice literature and model households as having heterogeneous preferences for milk and choose which milk to buy each month, trading off utility from product characteristics against prices. Purchasing the government good entails an additional utility cost that increases with distance to the nearest ration store. On the supply side, we extend the standard static model of oligopolistic price competition with differentiated goods to a setting in which profit-maximizing firms coexist with a government firm that sells subsidized milk up to a budget constraint. We estimate this model using household consumption data, variation induced by the staggered store introduction, and additional data on within-firm heterogeneity in supply costs—cost shifters—to separately identify demand parameters from supply-side parameters.

We use the model to perform three sets of counterfactual exercises. In the first, we study what happens when we remove the direct-provision system and transfer the government budget to households via unrestricted cash transfers. In the absence of the government good, we find

³ Thus we find evidence that ration stores indeed cause an increase in total milk consumption, which is a stated government objective.

that the private market price would increase by 2.4% and consumer surplus *decrease* by 8¢ (or 12% of the per household-month program cost, 68¢). Even though cash weakly dominates direct provision when prices are fixed as a result of a larger choice set, this result does not hold when private providers charge price markups. Furthermore, direct provision is redistributive toward poorer households: surplus of lower income households increases by 18¢ (26%), while richer households are worse off (−7¢ or −10%) when compared with providing the same amount of money via cash. The decrease in prices partially offsets the decline in surplus among richer households. In the absence of the price effect, the consumer surplus of the rich would fall by 27¢.

The second set of counterfactuals focus on milk-specific vouchers that households can redeem in the private market, with no direct government provision. We find that vouchers induce consumption levels similar to direct provision, but at the expense of some pass-through of the subsidy (26%) to increased milk prices. This result is in line with what has been found in the US for the Supplemental Nutrition Assistance Program (Hastings and Washington, 2010) and the Women, Infants, and Children (WIC) Nutrition Program (Meckel, 2020). Compared with direct government provision, consumer surplus *decreases* by 13¢ (20% of the per household-month program cost). The benefit of lower prices with ration stores more than offsets the loss in the surplus from the reduced variety in subsidized alternatives with direct provision (i.e., the government good is the only subsidized option). We also experimented with vouchers that subsidize a subset of brands of milk. These generate even larger price increases and allow the subsidized firms to exert greater market power.

In the third set of counterfactuals, given that we established that ration stores can increase consumer surplus, we turn to the question of how to improve the direct-provision system. For instance, the government could increase the price of the government good, trading off price spillovers to the private market against having to use less money to fund direct provision (and thus the ability to transfer cash to households). We show that the (average) consumer surplus obtained from the current intervention is close to what would be feasibly achievable by slightly increasing the price of the government good and reducing the quantity sold to keep the budget balanced.⁴ However, increasing the price of the government milk would transfer surplus from the poor to the rich. Overall, our results highlight the importance of market power in the design of redistributive policies.

This paper connects with several strands of the literature. A first set of papers looks at how to optimally redistribute resources to households in need and focuses on the cash versus in-kind debate (Aker, 2017; Hidrobo et al., 2014; Cunha, 2014; Cunha et al., 2019; Filmer et al., 2021).⁵ From these papers, we learn that both cash and in-kind transfers increase the consumption of goods (a finding we replicate in our setting), but that in-kind and cash transfers may have differential

⁴ An alternative way to evaluate this question is not to keep the budget balanced and trade off the cost of raising taxes against a larger or smaller budget for this program. We are refraining from such a comparison.

⁵ See Currie and Gahvari (2008) for a review of the previous literature.

price effects.⁶ We add to this literature by quantifying the welfare consequences on consumers of having the government provide the good directly as well as the distribution of these welfare effects across different types of consumers, while allowing for supply side responses. Importantly, we study a precise mechanism and quantify how much of the welfare gain is due to the mitigation of private suppliers' market power when the government acts as an extra competitor.

This paper is also, to our knowledge, the first to study how ration stores coexist with private market alternatives and to consider their effect on prices. [Besley and Kanbur \(1988\)](#) describe the theoretical impact of ration stores on poverty. [Tarozzi \(2005\)](#) uses a subsidy rate change to document no nutritional changes in India. [Gadenne \(2020\)](#) and [Gadenne et al. \(2020\)](#) document how ration stores can improve welfare by redistributing via taxes and subsidies and by insuring the poor against price fluctuations. None of these papers, however, provide an estimate of welfare in equilibrium or counterfactual comparisons.

The second set of papers deals with the effect of public provision on private market outcomes. A small but nascent literature has shown that suppliers may reoptimize their prices to capture part of the subsidies originally intended for consumers ([Hastings and Washington, 2010](#); [Meckel, 2020](#); [Collinson and Ganong, 2018](#); [Leung and Seo, 2019](#)). Since competition drives prices down, our paper (similar to [Handbury and Moshary, 2020](#)) documents positive price spillovers to other consumers in the market. We add to the literature by estimating consumer preferences, directly modeling and quantifying supply-side responses, and by quantifying the size of the markups of private suppliers.

The final strand of literature we contribute to studies the effect of increased competition on retail prices. A subset of these papers examines the entry into the retail market of large, private, for-profit competitors ([Atkin et al., 2016](#); [Moshary and Ilanes, 2020](#)), or small retail firms ([Busso and Galiani, 2019](#)). [Banerjee et al. \(2018\)](#) study entry by soliciting more offers in the procurement process. In contrast, our paper focuses on government-sponsored competition, a literature with renewed focus. [Handbury and Moshary \(2020\)](#) study the effect of school lunch programs on grocery sales. [Atal et al. \(2020\)](#), who document how the entry of public pharmacies increased the prices of their private counterparts in Chile. [Dinerstein et al. \(2020\)](#) study the effect on tuition of a public school construction program in the Dominican Republic. We consider an intervention in a food market in which the entrant is a government competitor with substantially lower prices than those charged by the privately provided alternatives. We complement this literature by measuring consumer surplus and welfare.

The rest of the paper is organized as follows. Section 2 provides the institutional background of our setting. Section 3 documents the data sources and provides a description of supply and demand in the market for milk. Section 4 provides motivating evidence on the effect of direct provision on market outcomes. Section 5 introduces a supply and demand model, which

⁶ While other papers have compared cash transfers and direct provision in rural villages ([Cunha et al., 2019](#); [Filmer et al., 2021](#)), our setting is the urban part of the country (65% of the Mexican population).

we estimate in Section 6. Section 7 evaluates counterfactual policies in order to make welfare assessments of the direct-provision system, and Section 8 concludes. Additional tables and figures relevant to the finer details of the paper are included in the Appendix and are prefixed by “OA-” in the main text of the paper.

2 Setting: Direct Provision of Milk Through Ration Stores

This section describes the intervention through which the government sells milk to households in Mexico. To ensure access to milk, the government operates their own company called Liconsa. According to government records, this company is considered to be a key element in their mission to achieve food security in Mexico (*Diario Oficial de la Federación*, 2014). Through this firm, the government sells milk in government-run ration stores at a low, subsidized price.⁷ In 2010—the first year in our observational period—the program had 5.8 million recipients and the government ran 7.8 thousand ration stores nationwide.

The flagship government brand is a low-cost, fortified milk with the same name as the company. In our sample period, “the government milk”—the name we use for ease of exposition—accounted for 13% of the total retail value in the urban milk market.⁸ It is sold in plastic bags, and is only available in ration stores as a 2-liter option (see Figure OA-1(a) for an example). Government milk is quite cheap: its price represents roughly one-third of the average price of private brands, and about half the price of the cheapest brands available. The government funds between 40% and 50% of the total production and distribution cost, and keeps the nominal price to consumers constant.⁹ The program’s budget is often under attack, with opponents arguing that these funds should be reallocated to other anti-poverty schemes such as cash transfer programs.¹⁰

While fortified and cheap, government milk tastes worse than private alternatives. The government production process has as a main objective to produce inexpensive fortified milk, with taste being a lower priority.¹¹ To keep production costs low, one-fourth (24%) of the total government milk distributed in the country is produced by rehydrating powdered milk, which is cheaper to import than raw milk (see Appendix B for more information on procurement). Government milk is often regarded by consumers as having lower quality relative to its private counterparts, which we confirmed through a double-blind tasting study (see Appendix C).

⁷ The program started circa 1944. The current version of this company dates from 1994. For a complete account of its history and the different interventions in the market through time, see Losada et al. (2000); Yúñez-Naude (2003).

⁸ Historically, the government exclusively sold this flagship milk. In recent years, Liconsa has introduced 10 additional brands of milk. Even so, the flagship product represents 94% of the retail value of government sales in our data.

⁹ In 2015, Liconsa’s budget was 514 million USD—roughly 12% of the budget of one of the largest anti-poverty programs in the world, the cash transfer program (*Progresas/Oportunidades/Prospera*).

¹⁰ Before 2003, Liconsa had *tortivales* (tortilla vouchers), which were eliminated to relocate funds to cash grants. In 2003, public officials debated whether to end the government milk program (Aragón Mladovich and Gómez-Ibáñez, 2004).

¹¹ In the production process, the government adds iron, zinc, folic acid, and vitamins. The government milk has been shown to be more effective, relative to non-fortified milk, in reducing anemia and preventing iron deficiency in children using RCTs (Villalpando et al., 2006; Rivera et al., 2010).

Despite the existence of eligibility criteria, and because of the composition of households in Mexico, access to government milk is quasi-universal. Per 2010 census data, more than 90% of households have at least one eligible individual (e.g., children or elderly adults).¹² On a potential customer's first visit to a ration store, they fill up a questionnaire that details the household's characteristics, backed up by proper documentation (i.e., birth certificate and proof of address). If eligible, customers get a ration card that can be used to buy milk once a week, on a specific store, on a specific time schedule (see Panels (b) and (c) of Figure OA-1). Customers are allowed to buy 4 liters of milk per week, per recipient, up to 24 liters (6 beneficiaries) per household.¹³

Instead of restricting based on income, the government uses the inconvenience of having to stand in line at odd times in the morning (e.g., 6:30 to 7:00 am) to purchase lower quality milk—a “micro-ordeal” cost à la Dupas et al. (2016)—as a mechanism to screen consumers. Income-based screening is often tricky in developing countries such as Mexico. Large fractions of the population work in the informal sector and have no proof of income; even for those who work in the formal sector, underreporting is frequent (Kumler et al., 2020). Figure OA-2 shows an example poster found in stores to encourage households to get their ration card; it details the eligibility criteria with requirement of proof of income. Except for the micro-ordeal cost, purchasing is not severely restricted—which may explain why resale is minimal.¹⁴

The location of ration stores is a useful mechanism for the government to screen consumers. Stores tend to be located closer to poorer households, which makes it easier for them to access government milk (see Appendix B for more information on store locations). In Figure OA-3(a), we use one cross-sectional slice of our sample of households (January 2010) and divide them by quintiles of the distance to the closest ration store. With the caveat that distance is not randomly assigned, we see that households closest to ration stores are four times more likely to purchase the government milk than those farther away.

Consumption of government milk varies with household demographics. In January 2010, 18% of households purchased at least one liter of government milk. Consumption of government milk is negatively correlated with income. Nonetheless, Figure OA-3(b) shows that even households with high socioeconomic status purchase government milk, with 15% purchasing at least one liter that month.¹⁵ Similarly, Figure OA-3(c) shows that households with children are more likely to purchase (25%) than those without (17%). Both consumption patterns suggest that despite

¹²The share of recipients by targeting criteria is: children ≤ 12 years old (54%); girls 13-15 years old (7%); women 45-59 years old (12%); adults ≥ 60 years old (23%); and other criteria (e.g., pregnancy) (4%) according to Liconsa (2015).

¹³Households that purchase government milk source close to 85% (on avg.) of their milk from the government.

¹⁴Two pieces of evidence indicate that reselling is rare. First, reselling is low in our household panel, as only 0.83% of the total purchases of the government milk happen outside of ration stores (involving third parties). Second, in a national survey, only 4% of recipients claimed to know someone who sells any portion of their government milk (Soto Romero et al., 2004, Figure 2.11).

¹⁵This result is not unique to our data. Figure OA-6 divides households by income quintiles using the 2012 National Health Survey (*Encuesta Nacional de Salud y Nutrición*) and shows that 6% of households in the upper quintile purchase government milk.

near-universal eligibility requirements, the “micro-ordeal” mechanism helps to select households in need.

3 Data and Descriptive Statistics

In this section, we first introduce the three primary data sources that, together, allow us to draw a detailed picture of the milk market in Mexico from January 2010 to December 2014. We focus on the downstream market that arises from firms selling liquid milk to consumers. After describing each source in detail, we describe the relevant features of the private market for milk in the country.

3.1 Data

Consumer Panel Microdata The first source of data is a rotating household panel maintained by Kantar World Panel.¹⁶ Households are visited twice a week to obtain a complete consumption diary about all of the products purchased by the household. We observe between 6,100 and 8,000 households per month and a total of 15,751 unique households from all 32 states of Mexico.¹⁷ The sample is designed to represent metropolitan areas (i.e., collection of contiguous municipalities) in Mexico with more than 50,000 inhabitants.

The panel provides information on all packaged food purchases made through time. For each purchase, we observe the transaction date, the item description at the barcode level, the price, the units purchased, the type of store where the purchase was made, whether the product was subject to special promotions, and the payment method. For milk purchases, the item description includes product characteristics such as the lactose content, the cream content (e.g., skim milk), the pasteurization process (e.g., ultra-pasteurized), and the package type (e.g., carton).

The panel also contains demographic and geographic information about each household. We observe demographic variables at the yearly level. They include a socioeconomic status (SES) categorical variable and household composition variables.¹⁸ We observe three values for the SES: low (38% of households), medium (43%), and high (18%).¹⁹ Geographic variables include the neighborhood of residence, which we use to compute a notion of distance to government stores.

Government Ration Store Data The second source is a census of government stores obtained from the Ministry of Social Development (*Secretaría de Desarrollo Social*). An observation in this

¹⁶Kantar World Panel is an international company that operates in 50+ countries. They specialize in the collection of household consumption data for marketing and sales strategy purposes. For more information on the data-sharing agreement, see Aguilar Esteva, Gutierrez, and Seira (2019).

¹⁷On average, each household stays in the sample for 2.3 years. We observe 2,964 households in all 60 months.

¹⁸The socioeconomic status maps from household asset ownership (e.g., appliances and vehicles), household utilities (e.g., internet usage), and the head of household’s schooling to purchasing power. It is computed using Consumer Expenditure Surveys (*Encuesta Nacional de Ingresos y Gastos de los Hogares*), and is widely used by the industry for marketing purposes (AMAI, 2018).

¹⁹Formally, low SES represents the D/E segments; medium SES represents C/C-/D+; high SES represents A/B/C+.

data is a store-month. For each store, we observe the exact address down to the longitude and latitude coordinates. With information at the store-month level, we construct a panel detailing the dates on which each store opens/closes.

Government Production Cost Data We collected data on the cost per liter the government pays to get milk into the hands of the final consumers. The cost per liter, computed yearly through the firm’s financial statements, is the sum of operating and production expenses divided by the total liters of milk. Operating expenses include administrative expenses and ration store costs. Production expenses include expenses for raw milk, expenses on powdered milk, manufacturing (e.g., adding vitamins, water, and pasteurization in the case of raw milk), and collection centers. We observe costs per liter on a yearly basis.

Private Market Distribution Cost Data We also collected data on the location where private firms source their milk. We use the National Statistical Directory of Economic Units (DENUE) from Mexico’s National Statistical Agency (INEGI) to collect information on 75 distribution centers of the 10 biggest firms. When excluding the government, these firms account for 91% of the total retail value in our consumer panel data. Smaller firms—those without distribution centers—source their milk from their processing centers directly. For 45 small firms, we hand-collected information on the location of all of their processing centers. These firms account for 7% of the total retail value in our consumer panel.²⁰

3.2 Sample Definition

We restrict our sample to simplify the computational complexity and report the exact details in [Appendix D](#). We restrict to the 36 most populated metropolitan areas in Mexico to ensure we observe at least 50 households per metro area when constructing market shares.²¹ The set of households we use represents 89% of the total population living in metropolitan areas (according to the 2010 census) and 51% of the total population in the country. In total, our sample restrictions leave us with 11,748 different households in 151 municipalities, 36 metropolitan areas, 60 months (2010 to 2014), and 302,644 metropolitan area-months.

We also restrict our sample in terms of products, which we define as a combination of a brand name, a pasteurization process (pasteurized or ultra-pasteurized); lactose content (lactose or lactose-free); package type (carton, plastic bottle, or plastic bag); and cream content (whole milk, skim milk, or semi-skim). To focus on the main source of milk for households, we restrict to two product sizes: 1 liter and 2 liters. This product restriction is useful because we cut by 40% the number of products but keep most of the retail value transacted (91%) and total liters of milk purchased (also 91%). We drop all household-months in which individuals “purchase” at a price

²⁰The remaining 2% of the private market are firms whose distribution centers’ locations are no longer public. These are small firms that serve only one metropolitan area per firm. Throughout the paper, when using this data we control for producer fixed effects and thus absorb the variation that would arise if we were to have these additional observations.

²¹In this process, we drop a small set of households that live in urbanized municipalities, but not in metropolitan areas.

of zero (0.48% of total transactions). In total, we are left with 528 products by 112 firms.

3.3 Descriptive Statistics

Household Demand for Milk Milk is an important component of the dietary habits of Mexican households. Monthly expenditures in milk, on average, rank third among food products—only behind tortillas and sodas. All of our observed households purchase milk at least once in the data, and, on an average month, 93% of households consume at least one liter of milk. As shown in Panel (a) of Table 1, there is little variation in the extensive margin of consumption across demographics, further evidence that milk is a widely popular consumption item.

The variation in consumption comes from the choice of products, in which preference heterogeneity is reflected in household choices. In Panel (b) of Table 1, we classify products based on their observable characteristics and present market shares (as a percentage of the total retail value) for selected demographics. Poorer households tend to purchase cheaper products: Relative to richer households, poorer households purchase more whole milk, with lactose, regular pasteurization, and sold in plastic bags. Products with these four characteristics are, on average, the cheapest among the available alternatives. In a similar pattern, households with children tend to consume cheaper milk than households without children.

Concentration and Prices in the Private Market Throughout the paper, we refer to firms using pseudonyms for ease of exposition. We refer as Private j to the j -th biggest private firm in the market, in terms of retail value of the milk sold nationally.

The milk the government offers coexists with privately supplied alternatives. One feature of the private milk market is the presence of two large, vertically integrated, nationally available milk brands that compete against smaller, regional competitors. These two firms, Private 1 (LALA) and Private 2 (Alpura) account for 58% of the total sales of milk. The remaining private firms are regional competitors, most of which serve a subset of contiguous states.²² In Panel (c) of Table 1, we describe market shares at the producer level. The uneven competitive landscape translates to us needing only five firms to capture 87% of the total retail value in the country.

Despite the large number of products in metropolitan areas (102 products by 25 firms, on average), the disparity in the size of producers means that most of the business takes place through a small number of sellers. In Figure 1 we plot concentration indices for the 15 biggest metropolitan areas, ordered from left to right in terms of population.²³ In all of the markets except for two, three private competitors capture 90% of the private market retail value. The concentration of sales is particularly relevant, because absent the government product, a small

²²For instance, Private 4 (Leche Guadalajara) is the biggest seller in the Guadalajara Metropolitan Area, located in the state of Jalisco. Products offered by Private 4 retail in other metropolitan areas located contiguous to Jalisco as well. In Figure OA-4, we show that the remaining firms among the biggest 20 follow the same pattern of geographic concentration of their sales in the country, with most of the sales taking place in a set of contiguous states.

²³The c - X index is defined as the retail share from the top X biggest competitors (excluding the government) in the metropolitan area, measured by their retail value market shares.


Table 1: Products, Producers, and Market Shares by Selected Demographics

	Overall	Socioeconomic status			Children		Price per liter (2010 constant MXN)			
		Low	Med	High	No	Yes	(7)			
	(1)	(2)	(3)	(4)	(5)	(6)				
<i>Proportion of household - months</i>										
<i>(a) – Milk consumption</i>										
Any milk	93	94	94	93	94	93				
Government milk	18	20	19	12	13	20				
Only government milk	9	10	9	5	6	10				
<i>Aggregate market shares (%)</i>										
<i>(b) – Market by product characteristics</i>										
All products	100	100	100	100	100	100	8.71	(3.48)	10.81	(1.67)
Cream content										
Whole milk	68	69	70	65	62	70	7.88	(3.52)	10.59	(1.58)
Semi-skimmed milk	17	14	16	23	23	14	11.87	(2.02)	11.87	(2.02)
Skimmed milk	15	17	15	13	14	15	10.56	(1.13)	10.56	(1.13)
Lactose content										
Lactose-free	6	4	5	10	10	5	13.34	(1.14)	13.34	(1.14)
Lactose	94	96	95	90	90	95	8.52	(3.42)	10.66	(1.58)
Bottle type										
Plastic bag	19	23	20	12	12	21	4.55	(1.66)	8.25	(1.07)
Plastic bottle	17	14	18	21	19	17	10.32	(1.41)	10.32	(1.41)
Carton	64	64	63	67	68	62	11.25	(1.49)	11.25	(1.49)
Pasteurization process										
Pasteurized	58	58	60	54	55	59	7.52	(3.53)	10.50	(1.67)
Ultra-pasteurized	42	42	40	46	45	41	11.17	(1.61)	11.17	(1.61)
<i>Aggregate market shares (%)</i>										
<i>(c) – Market by producer</i>										
Government	13	17	14	8	9	15	3.82	(0.12)		
Private 1	46	46	46	47	49	45	11.24	(1.25)		
Private 2	12	11	11	15	14	11	12.26	(1.28)		
Private 3	7	7	8	7	9	7	10.64	(1.48)		
Private 4	5	4	5	5	6	5	9.76	(0.98)		
Private 5	3	3	3	3	3	3	9.15	(1.12)		
Other firms	13	12	13	14	12	14	9.40	(1.89)		
Proportion of households	100	38	44	18	28	72				

Notes: Authors' own calculations using the household panel (see Section 3). The aggregate market share is the retail value of the selected products divided by the total retail value of the whole market. Columns (2) to (6) show results disaggregated by household demographics. Column (7) shows the mean and standard deviation of the average retail price per liter. The retail price represents sales at end price to consumer, including retailer and wholesaler mark-ups and sales tax. Price averages are taken across all transactions, weighting each transaction by the number of liters involved. Standard deviations are shown in parentheses.

number of firms would decide prices in the market, and could exercise their market power to increase them.

Figure 1: Concentration Indices (Excluding the Government Product) by Metropolitan Area



Notes: This figure plots concentration indices for selected metropolitan areas using our household panel data described in Section 3. We define the cX index as the market share (of retail value, when excluding the government retail value) of the top X firms in the metropolitan area. Metropolitan areas are shown in descending order in terms of their total population, starting with the biggest (Valle de Mexico) on the left.

The milk market in Mexico is not among the most concentrated milk industries in the world, which suggests that market power is a feature to consider in many countries. Figure OA-5 shows, using data from Euromonitor International and computing concentration at the parent company per country level, that the Herfindahl-Hirschman Index (HHI) for Mexico is close to that of comparably sized economies in Latin America.

4 Motivating Evidence

Having introduced the milk market, we now discuss the effect of direct government provision on household consumption and prices. In this section, we show descriptive evidence that the government direct provision scheme alters the quantities and prices at which private firms sell milk.


4.1 Evidence from Ration Store Openings

To study the effect of ration store access on household consumption and prices, we exploit a large staggered introduction of stores across the country between 2013 and 2014. They were part of the “National Crusade Against Hunger” (*Cruzada Nacional Contra el Hambre*) — a series of reforms aimed to reduce food insecurity.²⁴ Within our sample municipalities, 437 stores opened and 64 closed in the two-year period (see Table OA-1). As shown in Figure 2, the municipalities selected have wide geographic coverage, and the staggered introduction of stores generated month and

²⁴The government prioritized municipalities in extreme poverty and/or with little access to food as per the 2010 census. See Appendix E.1 for selection criteria and a comparison between selected and non-selected municipalities. 151 out of the 367 municipalities in our main dataset were selected for the Crusade. These selected municipalities encompass 85% of the total population in metropolitan areas.

geographic variation in the easiness of access and availability of government milk for households.

Figure 2: Staggered Store Introduction Overview



Notes: Panel (a) plots the store introduction across metropolitan areas. An observation is a metropolitan area - month. We normalize the number of stores in a given month by dividing by the number of stores in the given metropolitan area in January 2010. Selected metropolitan areas with large store increases are highlighted, but we plot all metropolitan areas in gray below. Panel (b) plots the location of ration stores opened in 2013 and 2014. Each mark on the map represents the location of a government store. Hollow gray circles represent stores that opened before 2013. Blue marks represent stores that closed within the two years. Red dots represent ration stores that opened between January 2013 and December 2014. Black lines represent the main highways.

The introduction of new stores was done in two steps. First, the federal government selected municipalities by prioritizing those with larger numbers of individuals in extreme poverty and those with low food security (see Appendix E.1). Once the municipality was selected, the government milk company was in charge of deciding where to open stores within municipalities. According to government officials, the exact within-municipality location of new stores depended on the number of eligible households living within the given location. In Appendix E.2 we show that, for municipalities with at least one store opening, the probability that a store opens in the census tract with the highest number of eligible individuals in the municipality is 0.40. This probability rapidly decays for the tract with the second-highest count (0.18), third (0.08), and fourth (0.07).

Within-location Variation in Access Over Time Our household panel’s rotating nature limits our ability to use within-individual variation in access to the government good. Instead, we divide the country using a hexagonal grid of 5 miles of edge length and collapse our outcomes at the hexagon level (see Appendix D.2 for additional details on hexagon construction). We restrict to the 230 (out of 289) hexagons with a store opening and regress


$$Y_{ht} = \vartheta_h + \vartheta_t + \sum_{j \neq -1} \beta_j \times \mathbb{1}_{(t-\tau(h)=j)} + \varepsilon_{ht}, \quad (1)$$

where h denotes hexagons, t denotes months, $\tau(h)$ is the month in which a ration store opens within hexagon h , and $\mathbb{1}_{(\cdot)}$ are indicator variables of the events in parentheses. Using this

specification, in Figure 3 we plot outcomes within the same hexagon before ($j < 0$) and after ($j > 0$), relative to one month before ($j = -1$) the ration store opens.

Using hexagons as a unit of analysis has some advantages over using municipalities. First, unlike municipalities, hexagons are equal-sized except for small geodesic adjustments. Second, hexagons, on average, are smaller than municipalities in terms of area, which allows for a more tightly parameterized specification. Third, hexagons allow us to capture the fact that households close to the border of a municipality can visit contiguous municipalities to go to a ration store. Finally, these hexagons are spacious enough that even with our panel’s rotating nature, we can include hexagon fixed effects to control for time-invariant differences across locations.

Figure 3: Event Studies with Ration Store Openings



Notes: These figures plot the estimates of β_j in Equation 1 for different values of j on the x-axis. Each line is a separate regression. An observation in this graph is a hexagon-month. Confidence intervals at the 90% statistical level are based on robust standard errors two-way clustered at the hexagon and month levels. All regressions include hexagon and month fixed effects.

Figure 3(a) estimates equation 1 where the dependent variable is an indicator for household h purchasing any milk in that period. It shows that the increase in the consumption of the government good around the time the government store opened coincided with a simultaneous decrease in the consumption of privately provided milk. There are parallel trends before the ration store opens, which suggests that unobservables are not driving the correlation between entry and consumption changes. One year after the store’s introduction, the probability that a household purchases milk from the government (shown in red) is 10 pp higher and stabilizes at that level. The probability that households purchase from private brands (shown in gray) decreases, but not by the same magnitude as the increase in the government good’s consumption probability, which suggests an imperfect substitution from the private market onto the government good. Figure OA-7(a) shows that the increase in consumption from the government is concentrated among low-SES households.

Turning to prices, Figure 3(b) shows in blue that the average price per liter within the hexagon decreases by 40¢ (3.8% of the unconditional mean) two years after the ration store opens. The decrease in the price per liter is a combination of the increase in the government good’s consumption and a decrease in the price of privately provided alternatives. For private goods, the price

per liter decreases by 25¢ (2.3% of the unconditional mean) two years after the store opening, shown by the gray line in our graph.²⁵ As a sanity check for our data, the price per liter of the government good (in red) is the same before and after the ration store opens.

One concern with the event study regressions is that the timing of entry may coincide with negative aggregate demand shocks that drive prices down. In such cases, we would expect the demand shocks to generate qualitatively similar patterns for other goods. Nonetheless, we find that new stores' introduction does not affect consumption within the market for sodas, yogurt, or cereals or their transacted prices (Figure OA-8).²⁶

Within Municipality-month Cross-sectional Variation in Distance The existing store arrangement, along with the opening of new stores, introduced variation in the distance to the closest ration store for households. In our metropolitan area sample, 43% of households faced a change in the distance to their closest ration store. Figure OA-10(a) plots the cross-sectional variation in the distance to the closest ration store for households in our metropolitan area sample. Figure OA-10(b) plots the average distance to the closest ration store over time for those with a change in the distance.

To assess the effect of ration stores on outcomes at the household-month level, we use the following specification for household i in municipality m in month t :

$$Y_{it} = \vartheta_{mt} + \beta \times \text{distance}_{it} + \gamma \cdot X_i + \varepsilon_{it}, \quad (2)$$

where Y_{it} is an outcome of interest; we measure the distance to the closest ration store active in period t from the centroid of the neighborhood where household i resides. Our coefficient of interest is β . We control for municipality-month fixed effects (ϑ_{mt}), as well as individual-level observables (X_i) that include their sociodemographic information and census-tract characteristics.²⁷ The identifying variation for Equation 2 exploits within-municipality-month cross-sectional variation, since we have households closer to or farther from ration stores, and within-individual across-month variation, since the store expansion decreased the distance to the closest ration store.

Conditional on observables, including a rich set of census tract level controls and municipality-month fixed effects, our identification strategy assumes that the residual variation in distance (arising from changes in store locations) is as good as random. Even after controlling for a host of

²⁵ Moreover, Figure OA-7(b) shows that the decrease in prices is mostly concentrated among poorer households.

²⁶ In addition to the placebos, we also show that changes in the composition of products (or producers), and specifically product exit, do not drive our results. To test this hypothesis, Figure OA-9 runs an event study in which the dependent variable is the number of unique products (producers) traded within the hexagon-month. We find flat event studies.

²⁷ Household demographic characteristics include an indicator for whether a given household has at least one child living in the family and SES fixed effects. Census tract characteristics are the share of households that live within the area that are eligible for government milk, the share of the population who cannot read or write, share of population without middle school or above, share of households with no access to laundry machines, share of households without access to running water, and share of households without a refrigerator.

observables, our specification could be picking up unobservable differences between households closer to ration stores and those farther away. We address the most important confounders in our robustness section, which we detail after showing the results.

Table 2: Staggered Store Introduction Results

Source:	At least one liter purchased			Total liters per household			Price per liter (2010 constant MXN)		
	Govt (1)	Priv (2)	Any (3)	Govt (4)	Priv (5)	Any (6)	Govt (7)	Priv (8)	Any (9)
Distance to closest govt. store	-0.014 (0.006) **	0.008 (0.005)	-0.000 (0.004)	-0.259 (0.110) **	-0.007 (0.164)	-0.266 (0.166)	0.000 (0.001)	0.055 (0.022) **	0.129 (0.052) **
One std. dev. increase distance ($\sigma \times \beta$)	-0.08	0.04	-0.00	-1.50	-0.04	-1.54	0.00	0.32	0.75
Mean of dependent variable	0.19	0.85	0.94	2.74	8.08	10.82	3.85	10.96	10.62
Municipality-month FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Household-type controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Household census tract controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	362,847	362,847	362,847	362,847	362,847	362,847	53,156	313,470	339,244
R-squared	0.353	0.135	0.103	0.294	0.298	0.280	0.770	0.417	0.376

Notes: Robust standard errors shown in parentheses. All regressions include month and municipality fixed effects and probability weights using our household panel. An observation is a household-month. The dependent variable for Columns (1) to (3) are categorical variables equal to one if the given household purchases at least one liter from the government, any private brand, or any brand, respectively. The dependent variable for Columns (4) and (5) is the total liters per household purchased from the government and from any private brand, respectively. The dependent variable for Columns (6) to (8) is the weighted average price per liter paid for the government, any private brand, or any brand, respectively.

Our estimation results show that ration-store access alters the equilibrium prices and quantities in the market. We present our results in Table 2. The table includes the regression estimates of β multiplied by one standard deviation of our distance, 5.8 kms. Columns (1) to (3) use as the dependent variable an indicator equal to one if the household purchased at least one liter of milk that month from the specific source. Column (1) shows that households located one standard deviation closer to ration stores have an 8 pp (50% of the mean) increase in the probability that they purchase from the government.²⁸ Column (2) shows that households located one standard deviation closer to ration stores are 4 pp less likely to purchase in the private market, but this is not statistically significant at conventional levels. Analogous results hold when the dependent variable is total liters (Columns 4 to 6).

Turning to prices, we see that households closer to ration stores pay lower prices per liter of milk consumed. In Columns (7) to (9), the dependent variable is the weighted price per liter, computed as the ratio between the total expenditures on milk of a certain type divided by the total liters of milk consumed of the same type. As expected, households farther away pay the same prices for government milk as households closer to ration stores. Interestingly, the total

²⁸ Table OA-5 shows that households with children living one standard deviation closer to ration stores are 19 pp. more likely to consume than households with children living farther away. Similarly, low socioeconomic status households living one standard deviation closer are 7 pp. more likely to consume than those farther away.

price per liter of milk *decreases* by 75¢ (7.4% of avg. price per liter) for households closer to ration stores. This is a combination of them purchasing more from the government, but also from a 32¢ decrease in the prices paid for private brands.²⁹

The main concern of our regression framework is that the location and opening/closing of ration stores (and thus our minimum distance measure) may be correlated with unobservables. For instance, if the government opens ration stores closer to more price-sensitive households (even after controlling for household characteristics), we would expect that those closer to ration stores would purchase more from the government (the cheapest product) and pay lower prices in the private market (since they are likely to purchase cheaper goods). Table OA-3 uses the same specification (Equation 2) to show that the distance to the closest ration store is not correlated with soda, yogurt, or cereal consumption or their prices. Thus the location selection explanation would have to be specific to household milk preferences. We find this implausible as the government did not use or have this information when choosing store locations.

To further study the price effect of ration stores, we turn to data at the transaction level and restrict to purchases from private brands. For transaction k by household i in month t , we regress

$$Y_{ikt} = \vartheta_{mt} + \beta \times \text{distance}_{it} + \gamma \cdot X_i + \phi \cdot W_{ikt} + \varepsilon_{ikt} \quad (3)$$

and show the results in Table OA-4. We weight observations appropriately to replicate the price results in Table 2, as shown in Columns (1) and (2). We then define product groups by combining product characteristics (e.g., ultra-pasteurized, lactose-free, skimmed milk in a plastic bottle) and add them as fixed effects in Column (3). Note that the distance coefficient when focusing on the within-product-group variation is about one-fourth the size of the distance coefficient without these fixed effects. Column (4) shows that products closer (in product space) to the government milk are those whose prices are the most affected. We define products as “close substitutes” if they share at least three out of the four product characteristics listed in Table 1. Close substitutes consumed by households one standard deviation closer to government stores are 53¢ cheaper (5% of the avg. price per liter of close substitutes) than close substitutes farther away. Meanwhile, far substitutes closer to government stores are only 12¢ cheaper (1% of the avg. price per liter of far substitutes), and this difference is not statistically significant.

4.2 Evidence from a Price Increase in the Government Good


On rare occasions, the Mexican government increases the price at which the government milk is sold. This typically happens when there is a budgetary restriction. In this section, we exploit one such change that occurred to fall during our sample period. In June 2013, the government

²⁹We obtain comparable results using (as regressors) the logarithm of distance, or with municipal-level regressions using store density (i.e., the number of stores per eligible households), as shown in Table OA-3.

increased the price per liter of raw milk paid to farmers by 50¢ in nominal terms.³⁰ To keep the subsidy per liter (and total government expenditures) somewhat constant, the government announced a 38¢ increase in nominal terms, from \$4.12 to \$4.50, for the price of the government good. The announcement took place in July and came into effect in August.

Figure 4 plots the average price per liter over time for the government (gray) and private brands (red). Dashed vertical lines represent the month of the increase in the cost and price, respectively.³¹ In constant MXN, the change represented a 33¢ increase in the price from one month to another—roughly an 8% increase in the price of the government good. There was a one-to-one pass-through of the government price change onto *private* prices: The 33¢ increase in the price per liter of the government coincided with a 34¢ increase in the price of private brands.

Figure 4: Price per Liter (Constant MXN) Across Consumer-Months (Jan 2012 - Dec 2014)



Notes: Authors' own calculation with household panel data using population weights. An observation is a consumer-month. Each dot represents the average price per liter across all transactions for the selected firm. Error bars denote the standard error of the estimation.

A simple explanation for the pass-through is that the government's increase in the price per liter paid to farmers increases costs to private producers, which even in a perfect competitive world would then be passed through to consumers.³² This is unlikely for several reasons. The government's increase in the amount of money paid to farmers for their raw milk happened two months before the retail sales price increase, but the increase in private suppliers' prices happened either the same week the government's *sales* price increased or a week later. More importantly perhaps, a large part of the milk the government sells comes from international sources, making it unlikely that the price the government pays farmers influences the price these farmers charge to private sellers. Moreover, large private suppliers are vertically integrated and their cost should not change as a result of this government action, yet they also respond in prices (see Figure OA-11). Finally, it is not the case that products derived from milk experienced price changes, as

³⁰Farmers in the Jalisco region were threatening to destroy their milk unless the government increased the price paid for raw milk (Euromonitor International, 2015).

³¹The difference between the red and gray lines is not drawn to scale, but each dotted line represents 20¢. Variation over time in the price of the government good arises from variation in the consumer price index.

³²In a vastly different context, Clemens and Gottlieb (2016) show that Medicare payments may have spillovers to private market prices by increasing the physicians' bargaining power.

would be expected from a cost-side explanation.

4.3 Discussion

The results so far show that direct government provision through ration store alters equilibrium outcomes by increasing consumption and lowering prices in the market. This may operate, on the one hand, by decreased demand of the private good as households substitute toward the government good by virtue of having a nearer government store. It may also operate more directly as the lower government good price limits the prices private providers can charge. Because there seems to be little price dispersion, there are positive spillovers to other households from lower prices, even if they do not live close to a government store.

Even if prices decrease, the welfare implications of direct provision are ambiguous for at least two reasons. Government milk has different characteristics than non-government milk—but also because taxpayers money is used to fund this endeavor. Moreover, there are alternatives for government involvement in this market, such as giving vouchers for milk. We need a framework to evaluate how individuals trade off lower prices versus the cost of direct provision. Ideally, such a framework would incorporate market power and supply-side responses to allow us to compare the outcomes and welfare implications of counterfactual government interventions. The next section develops a parsimonious model that allows us to do just this.

5 A Model of Consumer Choice with Private and Public Supply

We introduce a discrete choice, static model of oligopolistic price competition with differentiated products. In this model, profit-maximizing firms coexist with a government firm whose product is available through ration stores and sells up to a budget constraint.

Throughout the model, a geographic unit is a metropolitan area and a time unit is a month. A market is a combination of a geographic and time units, and is denoted by m to simplify notation. Products are denoted by j , and the collection of all available products in a market is denoted by $\mathcal{J}(m)$. We denote by $j = 0$ the outside option and by $j = g$ the government product. There are F firms, each indexed by f . Firm f sells $|J(f)|$ products and $J(f, m) \subset J(f)$ of these products are available in market m .

5.1 Demand

We follow the literature on discrete-choice demand with substitution patterns that potentially differ based on household characteristics (Berry, 1994; Berry et al., 2004; Petrin, 2002). For simplicity, we make the assumption that households buy a single product, or the outside option, per month.³³ We separate households into seven types based on their observable characteristics. The

³³When households purchase more than one product in a given month, we define the chosen product as the one that represents the highest amount of liters purchased for the given household-month. In Appendix G.1 we provide additional details on our implementation.

first six household types are eligible to purchase from the government (e.g., those with children or elderly adults), and are defined by a combination of their SES (low, medium, or high) and whether they have children. The remaining household type is ineligible to purchase government milk.

Household i of type t buying product j in market m derives utility given by

$$u_{ijm} = \mathbf{x}_j \cdot \boldsymbol{\beta}_t - \alpha_t \cdot p_{jm} - \gamma_t \cdot d_{im} \cdot \mathbb{1}_{(j=g)} + \xi_{jm} + \varepsilon_{ijm}, \quad (4)$$

where we normalize the outside option utility ($j = 0$) to zero. The vector of observable covariates (\mathbf{x}_j) includes a constant (i.e., an indicator variable for inside products), and product characteristics fixed effects based on combinations of the product's lactose content, cream content, bottle type, and pasteurization process, as described in Table 1. Our main specification also includes producer fixed effects, including a government producer dummy, to allow the valuation of products of seemingly the same observables to vary with the producer and across types.

Distance to the closest government ration store for household i is given by d_{im} and varies across households within a metropolitan area and over time. Households pay the utility distance cost only if they purchase the government good (denoted by the indicator $\mathbb{1}_{(j=g)}$). This normalization indicates that households have to go out of their way (and pay a micro-ordeal cost) to purchase the government product, while all other products can be purchased in the supermarket. Implicit in this assumption is that households visit the supermarket anyway, even if they do not purchase milk that month.³⁴

The retail price per liter of milk is given by p_{jm} and is constant across households within a market. Prices can be correlated with product-market-specific preference shocks (ξ_{jm}), which are also constant across households within a market. These are common knowledge to households, private sellers, and the government, but are unobserved by the econometrician. These shocks may reflect both unobserved product characteristics across markets (e.g., advertising) and/or unobserved variation in tastes across markets. Finally, ε_{ijm} is an idiosyncratic preference shock that is observed by households and is assumed to be iid extreme value type I error.

Given the above assumptions, the probability that household i of type t purchases product j in market m is given by $\mathbb{P} [u_{ijm} \geq u_{ikm} \forall k \in \mathcal{J}(t, m)]$, where $\mathcal{J}(t, m)$ is the set of all products available in a market for an individual of a given type. For eligible types, it includes all products in the market, whereas it excludes the government product for the ineligible type. Suppressing the arguments of the demand function, and using the functional form of the logit error and our

³⁴The proposed model is equivalent to one in which we assume that there is a constant distance cost for all other products (d'_{im}), and that households have to pay this distance cost even if they do not purchase milk (and thus the distance cost renormalizes the utility of the outside good to a different scalar).

household utility decomposition, this probability collapses to

$$\sigma_{ijm} = \frac{\exp(\mathbf{x}_j \cdot \boldsymbol{\beta}_t - \alpha_t \cdot p_{jm} - \gamma_t \cdot d_{im} \cdot \mathbb{1}_{(j=g)} + \xi_{jm})}{1 + \sum_{k \in \mathcal{J}(t(i), m)} \exp(\mathbf{x}_k \cdot \boldsymbol{\beta}_t - \alpha_t \cdot p_{km} - \gamma_t \cdot d_{im} \cdot \mathbb{1}_{(k=g)} + \xi_{km})}, \quad (5)$$

which varies across households, products, and markets.

We aggregate the individual-level choice probabilities to construct both the type-level market share and the aggregate market share for product j in market m . Denote by $\mathcal{J}(t, m)$ the set of all households in market m of type t . Then the type-level share for product j is given by the average of the choice probability of all type- t households within the market

$$\sigma_{jm}^{(t)} = \mathbb{E}_i [\sigma_{ijm} \mid i \in \mathcal{J}(t, m)] \quad (6)$$

where the expectation operator, $\mathbb{E}_i[\cdot]$, denotes the average across individuals. The aggregate market share for product j in market m is given by

$$\sigma_{jm} = \mathbb{E}_i [\sigma_{ijm}], \quad (7)$$

and is an important primitive for producers as described below.

5.2 Private Supply

Firms set prices following a Nash-Bertrand model with differentiated products within a market. We drop all market subscripts to simplify notation. Each private firm f sells $|J(f)|$ products in the given market. Firms sell directly to households, so we use firms and sellers interchangeably. To sell milk, each seller produces brand j at a per liter marginal cost of $c_j > 0$. These costs vary across products and (for a given product) across markets. Sellers know the distribution of types in each market and are constrained to selling each of their products at a single price per market. For simplicity, we focus on the choice of prices taking products (and product characteristics) as given. Following the literature, we assume that products that are not sold in a market were never offered.

Firms choose the prices that maximize their profits conditional on the remaining firms' prices. We denote by p_j the per-liter price of product j in the given market and by \mathbf{p}^f the vector of prices of those products sold by firm f . Similarly, denote by \mathbf{p}^{-f} the remaining prices. The seller-maximization problem for firm f in the given market is to choose a price vector \mathbf{p}^f to maximize their own profits, taking the remaining firms' prices as given:

$$\max_{\mathbf{p}^f > 0} \sum_{j \in J(f)} (p_j - c_j) \times \sigma_j(\mathbf{p}^f, \mathbf{p}^{-f}), \quad (8)$$

where σ_j is the aggregate demand for a given product in a given market, as in Equation 7. The

first-order condition associated with Equation 8 with respect to price p_j is given by

$$\sigma_j(\mathbf{p}) + \sum_{k \in I(f)} (p_k - c_k) \times \frac{\partial \sigma_k(\mathbf{p})}{\partial p_j} = 0, \quad (9)$$

where firms internalize that increasing the price of product j potentially affects demand for the other products they offer in the same market. Given our demand estimates, the observed prices, and market shares, the set of price first-order conditions for all firms in a given market pinpoint a single set of cost estimates that rationalize our observed data (Bresnahan, 1981).

5.3 Government Supply

We assume the government's objective function is to sell as much milk as possible up to a previously defined budget constraint using a prespecified price per liter. The government deems milk an essential good whose availability should be guaranteed to a set of eligible households. Consequently, the government does *not* maximize revenue. Instead, for a fixed per liter budget B , the government firm satisfies

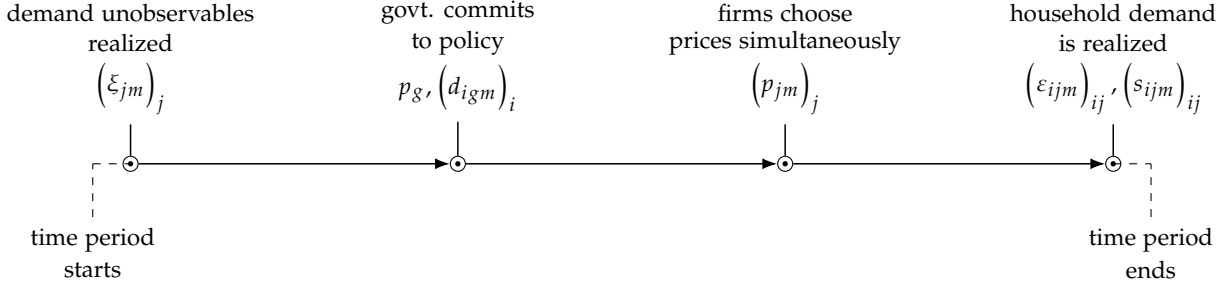
$$c_g \times \sigma_g(p_g, \mathbf{p}^{-g}) \leq B + p_g \times \sigma_g(p_g, \mathbf{p}^{-g}), \quad (10)$$

where c_g is the per liter cost and p_g is the per liter price. Note that when B is positive, the government may sell at a price lower than its cost. These modeling assumptions resemble two key features from the actual intervention (as explained in Section 2). First, the government budget is set long before, up to a year in advance. Second, the government firm is constrained to setting a single price for all geographic areas within a given month, and thus the government's pricing behavior does not respond to local demand shocks. Through our data we observe prices, quantities, and the government cost per liter, and thus can back out the per-market budget for every market.

5.4 Equilibrium

Figure 5 introduces the timing of our model within a given market. At the beginning of the time period, demand unobservables (to the econometrician) are realized. These unobservables are common knowledge to households, private firms, and the government firm. Next, the government firm chooses and credibly commits to a price for the government good and a location of all ration stores in the country. After the government sets up their desired price and locations, the remaining (private) firms choose prices simultaneously, taking every other firms' prices as given. Only after both of the price-setting stages do households make choices.

Figure 5: Model Timeline for a Given Market



Notes: This figure shows a timeline for a given time period in a given market (metropolitan area-month).

5.5 Model Discussion

5.5.1 Benefits

The model provides a framework to study the effects of the direct provision on market outcomes. It accommodates a set of key trade-offs to evaluate direct provision. Moreover, it is flexible enough to capture the main market features (e.g., price dispersion, consumer heterogeneity, and varying degrees of market power) we documented in Sections 2, 3, and 4. We highlight the most important model features and the main parameters for each of these features below.

Screening Mechanisms Absent price effects, direct provision can increase consumer welfare by improving the targeting of transfers when income is costly to verify. Direct provision can allow for “self-targeting” by incentivizing only the intended recipients to participate while others opt out (Blackorby and Donaldson, 1988; Besley and Coate, 1991; Gahvari and Mattos, 2007; Gadenne, 2020). In such instances, the government can direct more resources to intended recipients than with direct cash transfers, since these would have to be universally distributed. In our model, the screening of recipients can arise both because of the perceived quality of the government good across types and because of the cost of traveling to the ration store.³⁵ The type-specific perceived quality is parameterized through the type-specific government producer fixed effect; the household-specific distance to travel to the the closest ration store is given by $\gamma \times d$.

General Equilibrium Responses An important advantage of the model is that it enables us to assess equilibrium changes in welfare and tease out their sources. Absent price effects, direct provision has the drawback that households would rather obtain cash-equivalent transfers (Atkinson and Stiglitz, 1976; Currie and Gahvari, 2008). However, such a relationship is no longer obvious when we consider that prices may change in response to the government intervention. From the consumer’s ex ante perspective (i.e., before making choices), the effect of direct provision on con-

³⁵The empirical literature has shown that both targeting through quality (Lieber and Lockwood, 2019) and targeting by imposing a small cost to consumers (Sylvia et al., 2015; Dupas et al., 2016) are effective as screening tools to select households in need.

sumer welfare is ambiguous and depends on (i) the slope of demand with respect to prices with and without provision, (ii) the quality of the government good, and (iii) the cost of production for the government, relative to private alternatives. In [Appendix F](#), we use a simplified model with one private firm, one government firm, and one consumer to show the intuition.

5.5.2 Limitations

For parsimony, our model makes certain simplifying assumptions that leave out some aspects of the market we consider to be second order in our context. We discuss the main limitations of the model below, with a particular emphasis on why we think that abstracting from them is unlikely to dramatically affect the main conclusions.

Non-binding Quota We decided not to model the government good quota because it seems not binding for most households in terms of their most consumed alternative. We do not directly observe the quota at the household level, but two pieces of evidence support our assertion. First, for 86% of the household-months with government purchases, the government good is the product with the highest number of liters for a household. Second, conditional on purchasing from the government in a given month, households source 85% of their total consumption from the government on average.

No Income Effects One of the discrete choice model's known features is that it can be derived from a quasilinear utility function which, by design, is free of wealth effects ([Nevo, 2000](#)). Because of that, demand would not change with the size of the cash transfer in our counterfactuals, despite its being a channel theorized by [Cunha, de Giorgi, and Jayachandran \(2019\)](#) through which prices may increase. We consider killing this channel to be a reasonable assumption, given our industry and the size of our transfers. Furthermore, our assumption, if anything, means that price increases absent the government intervention would be higher with income effects. Thus the consumer surplus gain with direct provision would be underestimated.

Entry and Exit Throughout our analysis we focus on short-run pricing competition, omitting the entry and exit of retailers from markets. Endogenizing entry and exit seemed beyond the scope of this paper, both because of the potential multiplicity of equilibria and because there is little variation in this margin in our data (as shown by [Figure OA-9](#)), which suggests that entry and exit are not first-order concerns. Nonetheless, in [Appendix H](#) we describe how (a somewhat stylized version of) the entry of new competitors would change our results through simulations.

Cost Changes with Government Provision The model does not take a stand on how the direct-provision system changes private providers' cost structure. However, costs may decrease because, without government procurement, farmers would sell milk exclusively to (non-vertically integrated) firms. We consider the cost pass-through of direct provision to be, if anything, small. First, powdered milk, which represents one-fourth of the government's milk, comes from inter-

national sources where the government is a price taker. Second, the largest sellers (Private 1 and 2) in the market are vertically integrated, which implies that they produce their milk. For these reasons, we leave the cost structure unchanged in the main body of the paper. Through simulations, in [Appendix H](#) we describe changes to the cost structure and their impact on our counterfactual results.

6 Empirical Implementation

The end goal of this section is to identify and estimate the parameters of household utility (Equation 4) and marginal costs (Equation 9). In estimating demand, we face a common identification challenge: Firms’ pricing decisions can be correlated with demand unobservables.³⁶ To deal with such correlation, we make use of additional data on supply costs that do not directly correlate with choices. Our preferred solution is to estimate the model using the generalized method of moments (GMM) with the additional “cost-shifters” that we introduce below. In this section, we first introduce the instrument for prices, then discuss the additional identifying moments and finish by showing our demand and supply estimation results.

6.1 Econometric Specification

Parameterization of Product-Market Unobservables For product j in market m , corresponding to calendar month n , calendar year y , and metro area a , we decompose our unobserved product valuations as

$$\xi_{jm} = \zeta_a + \zeta_k + \zeta_y + \Delta\xi_{jm}, \quad (11)$$

where ζ_a is a metropolitan area fixed effect that captures systematic variations in taste across locations, relative to the outside option; ζ_n and ζ_y are calendar month and calendar year fixed effects intended to capture variation in aggregate demand over time. The remaining structural error $\Delta\xi_{jm}$ represents deviations across products within a market after controlling for the above described fixed effects.

Cost-shifter Instrument for Prices We trace out the distribution network that firms use to deliver milk to final retailers, capturing variation in costs across producers and markets. In particular, we exploit the geographic distance between metropolitan areas and the location at which firms source their products.³⁷ To compute our distance measure, we obtain the driving kilometers between each metropolitan area centroid and the nearest distribution center using Google Maps (see [Figure 6](#) for an example).³⁸ We then multiply the driving distance by the diesel price per

³⁶Firms set prices with knowledge of these unobservables and may, for instance, decide to increase prices for products with (unobservably) good taste shocks. For instance, such correlation may arise from advertising—unobservable in our data but observable to consumers—whereby firms may decide to advertise and introduce price discounts ([Petrin and Train, 2010](#)).

³⁷A similar approach has been used in the markets for beer ([Miller and Weinberg, 2017](#)) and yogurt ([Villas-Boas, 2007](#)).

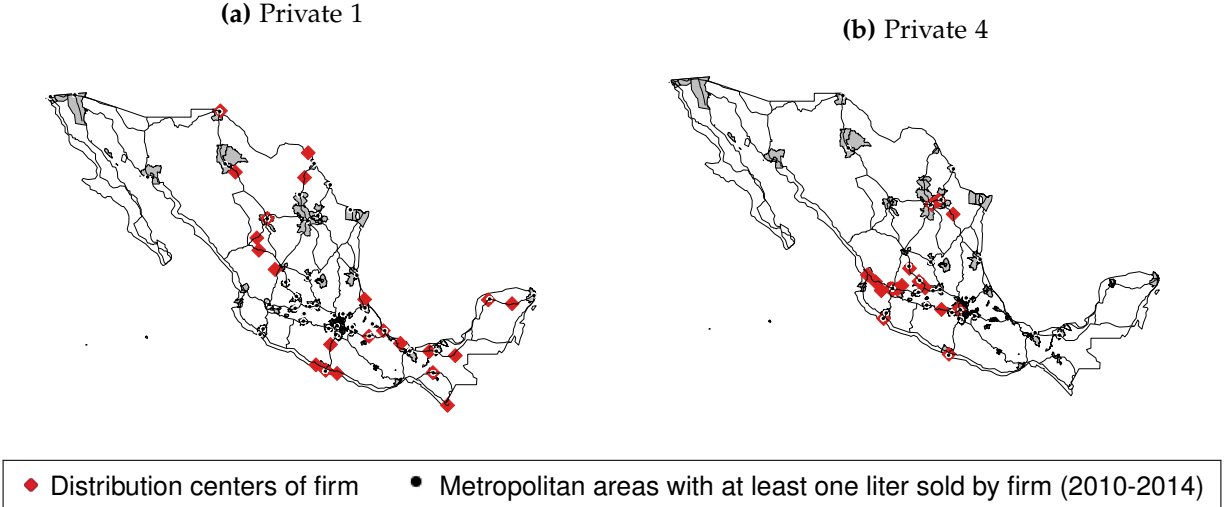
³⁸Some small firms do not have a distribution center. We compute the distance from the closest processing center.

liter. Formally, for product j , sold by firm $f(j)$, in market m , and month $\tau(m)$, the instrument is

$$z_{jm} = \text{distance to metro area}_{f(j),m} \times \text{diesel price}_{\tau(m)-3}, \quad (12)$$

which varies at the producer-market level. The cross-sectional variation comes from driving kilometers; time variation comes from real price changes in the diesel cost.³⁹

Figure 6: Location of Distribution Centers for Selected Firms



Notes: This figure shows the location of distribution centers for two sample firms. Red diamonds represent all of the locations of distribution centers, while black dots represent the centroids of the different metropolitan areas in which the given firm has at least one liter sold in our data. For reference, the solid dark lines are the main highways in the country.

To measure the effect of distance on prices, we run the following regression in which j and m index products and markets, respectively:

$$p_{jm} = x_j \cdot \pi + \kappa \cdot z_{jm} + \eta_{jm}, \quad (13)$$

where x_j are the product characteristics fixed effects of the given good. The covariates also include producer, calendar month, calendar year, and metropolitan area fixed effects. Table OA-6 shows that the coefficient on our distance instrument is positive and statistically significant, which implies that increasing our notion of distance increases the price per liter. Keeping diesel price constant, an increase of 1 std. deviation in the distance increases the price per liter by 11¢. The F statistic is 89, which suggest the instrument is not weak. The magnitude of our coefficient indicates that the variation in prices that scales with the diesel price times the distance accounts for 2.3% to 3.5% of the total variation in prices—similar in order to those found in other industries

³⁹We used the three-month lagged diesel price as opposed to contemporaneous prices because, based on the pasteurization process, a liter of milk may take up to three months into get to the hands of the final consumer. By using the lagged diesel price, we isolate the variation that affects producers' costs and not consumers directly (since the contemporaneous effect of transport costs may have a direct impact on household demand).

using the same type of instruments.⁴⁰

Market Shares by Demographics Besides the cost-shifter instruments, we use additional moments from our data. The first set of moments we include minimizes the distance of the model’s shares predictions at the product-type level ($\sigma_{jm}^{(t)}$) to those in the data ($s_{jm}^{(t)}$). The moment condition to minimize is given by

$$\sigma_{jm}^{(t)} - s_{jm}^{(t)} = 0 \text{ for all } j = 0, 1, \dots, \mathcal{J}(m) \text{ and } m = 1, 2, \dots, \mathcal{M} \text{ and } t = 1, 2, \dots, 7, \quad (14)$$

where we use our micro data to construct the observed shares.

Staggered-store Introduction Moments The second set of moments is somewhat less standard and uses variation introduced by the staggered-store introduction. These moment restrictions choose parameters to match the cross-sectional distance regression results from Section 4. Equation 2 (restated here to match the model notation for readability) in Section 4.1 computes the effect of the distance to the closest ration store on market outcomes. For household i in market m , municipality $\mu(i)$, and month $\tau(m)$, the equation is

$$Y_{i,\mu(i),\tau(m)} = \vartheta_{\mu(i),\tau(m)} + \lambda \times \text{distance}_{i,\mu(i),\tau(m)} + \rho \cdot \mathbf{X}_i + \varepsilon_{i,\mu(i),\tau(m)}, \quad (15)$$

where λ is the coefficient of interest. In the estimation procedure, we include as a targeted moment the difference in the value of λ in the data and model-based regressions for the probability of consuming the government good. We match the aggregate distance gradient (Column 1 in Table 2) and the differential distance effect across household demographics (Column 1 in Table OA-5).

6.2 Identification Intuition and Estimation

Identification Intuition Having introduced the specification of the errors in the market, the price instrument, and the additional moments that we posit to identify demand, we provide the intuition for the identification of our full demand model. Our identification assumption is that the residual variation in product-market unobservables ($\Delta\xi_{jm}$) is orthogonal to the product characteristics (x_j), the location of all ration stores (and consequently the distance to the ration store for households, d_m), the above-defined distance instrument (z_{jm}), and the calendar month, calendar year, and metropolitan area fixed effects (ζ_m).

First, to identify the price coefficients (α_t), we use the correlation between the cost-shifter-induced variation in prices and actual household choices. The identification assumption requires that our instrument satisfy the exclusion and relevance restrictions, after conditioning for the battery of fixed effects. For instance, our fixed effects absorb geographic and time unobservable taste differences (e.g., advertising). The instrument generates variation in costs across products

⁴⁰ A similar instrument accounts for 2% to 3% of the total variation in costs in the beer industry (Miller and Weinberg, 2017). Note, however, that the magnitude of distribution costs may be bigger, but part of it is absorbed by producer and metropolitan area fixed effects.

of different producers (for the same market) and markets (for the same product). Thus, it is reasonable to expect higher prices for firms for whom the above defined distance notion is higher. The exclusion restriction requires that product-market taste shocks are not correlated with the residual variation in the instrument. Even though firms may observe these cost shocks, it seems unlikely that they can respond to them by altering the unobservable taste shifters of their products accordingly.

Second, to identify consumer preference parameters (β_t), we use the correlation between the product characteristics of the alternatives available to consumers and the actual observed choices. The assumption is that product characteristics (x_j) are not correlated with the taste shifters ($\Delta\xi_{jm}$), conditional on the set of included fixed effects. One advantage of the microdata is that, by observing multiple households of the same type within a market, we can use the variation in their choices to identify consumer preferences while holding the unobservable taste-shifters fixed (Berry and Haile, 2020). In addition, the existence of regional firms that sell in a small set of contiguous states means that a given product is offered along with those of different competitors across markets.

Third, to identify the distance parameter coefficients (γ_t), we use the correlation between a household’s distance to the ration store and their likelihood of purchasing the government brand. As discussed in Section 4.1, the staggered-store introduction is unlikely to be correlated with the taste-shifters after we control for the set of controls in Equation 2. The data also allow us to distinguish the distance sensitivity across household types, since holding the distance to the ration store fixed, poorer households are more likely to purchase from the government. By using the model to replicate the staggered-store introduction results—loosely speaking—we use the “exogenous part” of the cross-sectional correlation between distance and choices.

Estimation We estimate the model through two-step GMM following Berry, Levinsohn, and Pakes (2004), taking advantage of our micro data using a slightly tweaked algorithm to handle a large number of observations (3.9 million household-product-months, 140 thousand type market share equations, and 21 thousand market share equations). The estimation procedure is, for the most part, standard. We describe it in detail in Appendix G.2 and summarize here the two main steps. In the first step, we randomly select 80% of the markets (metropolitan areas-months). We use this sample to estimate the “mean valuation” for each product in each market and the type-idiosyncratic taste shifters by using the type-market share moments and the staggered-store introduction moments (Equations 14 and 15). In the second step, we use the estimated type taste shifters and return to the full data to estimate the mean taste parameters (α, β) using the instrumental variation provided by our cost shifter instruments (Equation 12).

6.3 Supply and Demand Estimates

Table 3 shows the demand estimation results of our preferred specification; we discuss additional specifications and how our parameters depend on the data moments in G.5. The first

panel reports the coefficients of the type-specific utility (Equation 4). Each column represents a household type. Our estimates show statistically significant differences in preferences and imply heterogeneous substitution patterns across household types. We discuss the main drivers of choice and its heterogeneity below.

First, the price coefficient estimates imply that paying an extra peso for the same good hurts poor households more than rich ones in terms of utility (e.g., 0.90 vs. 0.78 for households with at least one child), which induces them to consume cheaper goods. Similarly, households with at least one child are more reactive to prices, conditional on socioeconomic status. The implied weighted average own-price elasticity is 5.27.⁴¹ This elasticity is higher than those found in prior studies of the milk industry in other contexts without instrumental variation,⁴² and in the range of values for other papers with similar cost-shifter instruments.⁴³

Second, the distance coefficients show that richer households are less reactive (in utility terms) than poorer ones to the distance to the closest ration store. However, the true total cost of visiting the ration store requires us to consider that richer households live farther away from these stores. To interpret the distance parameters, we divide them by the price coefficient and multiply by the distance to their closest store—which translates to the utility cost of traveling to the ration store in monetary terms. Despite richer households living farther away from ration stores, the cost of traveling is 2.6 times *lower* for richer households with children than for poor households without children. Based on these estimates, the distance to the ration stores disincentivize poorer households more than richer ones. Thus, the location of stores serve as a *negative* screening mechanism, selecting out poorer households.

Third, the constant denotes the utility value for pasteurized, lactose, whole milk in a plastic bag container by Private 1 and sold in Mexico City in January 2010 at a price of zero. In comparison, the government producer fixed effect is the difference in utility for an equal-characteristic good, but sold by the government (without considering the distance cost). Negative government producer fixed effects show that government goods are deemed to have lower quality. Richer households find the government good to be worse than poorer households, as can be seen from the magnitude (in absolute terms) of these coefficients. The government fixed effects include all differences in quality that are constant to all individuals within a given type. In contrast to the cost of traveling to ration stores, the perceived quality differences serve as a *positive* screening mechanism, selecting

⁴¹The weighted own-price elasticity is computed first by computing the elasticity at the product-type-market level with respect to the price of the given product $\left(\frac{\partial \sigma_{jm}^{(t)}}{\partial p_{jm}}\right) \times \left(p_{jm} / \sigma_{jm}^{(t)}\right)$ and then averaging across type-market shares using population weights.

⁴²Comparable estimates for the elasticity of milk in the US depend on the product specifics and range between $[-0.80, -0.51]$ (Gould, 1996); $[-0.59]$ (Andreyeva et al., 2010); $[-2.04, -1.16]$ (Davis et al., 2009); and $[-2.317]$ (Schmit et al., 2002). In Mexico, the only computation available is $[-1.65]$ (Colchero et al., 2015) based on an almost ideal demand system with no instrumental variation. Interestingly, our demand estimation produces comparable estimates if we do not instrument for prices, or account for heterogeneity in demand, with an elasticity of -0.8 .

⁴³Miller and Weinberg (2017) find an elasticity of -4.74 in the market for beer and Villas-Boas (2007) of -5.9 in the market for yogurt.

out richer households.

Table 3: Demand Estimation Results

	Ineligible households	No children			With children		
		Socioeconomic status					
		Low	Med	High	Low	Med	High
(1)	(2)	(3)	(4)	(5)	(6)	(7)	
<i>(a) – Demand estimation coefficients</i>							
Price (α)	-0.73 (0.02)	-0.74 (0.01)	-0.68 (0.01)	-0.63 (0.02)	-0.90 (0.01)	-0.84 (0.01)	-0.78 (0.01)
Distance to ration store (γ)	-	-0.16 (0.09)	-0.11 (0.06)	-0.05 (0.05)	-0.16 (0.05)	-0.10 (0.03)	-0.05 (0.05)
Government producer FE	-	-1.07 (0.08)	-1.32 (0.09)	-1.57 (0.11)	-2.60 (0.07)	-2.85 (0.07)	-3.10 (0.09)
Constant	9.00 (0.19)	9.23 (0.14)	8.60 (0.15)	7.97 (0.18)	10.78 (0.12)	10.15 (0.11)	9.52 (0.13)
<i>(b) – Selected statistics</i>							
Own-price elasticity	-5.44 (0.04)	-4.95 (0.09)	-4.74 (0.07)	-5.01 (0.06)	-5.39 (0.12)	-5.33 (0.10)	-5.51 (0.07)
Uniform price increase elasticity	-0.06 (0.00)	-0.06 (0.00)	-0.05 (0.00)	-0.05 (0.00)	-0.07 (0.00)	-0.06 (0.00)	-0.06 (0.00)
Utility cost of traveling to ration store (\$)	-	0.16 (0.00)	0.11 (0.00)	0.06 (0.00)	0.12 (0.00)	0.09 (0.00)	0.05 (0.00)

Notes: Robust standard errors clustered at the market level computed via bootstrap shown in parentheses. This table summarizes the most important demand-side parameters. Panel (a) shows raw coefficients from Equation 4 for the most important parameters. Tables OA-7 and OA-8 show the remaining parameters. Panel (b) shows selected moments computed using these coefficients.

Model Fit Figure OA-12 plots selected moments to show that our model appropriately fits the data. Panel (a) shows a binned scatterplot where an observation is the market share of a given product among households of a given type in a given market. The locally weighted regression, shown in red, closely resembles the 45 degree line, and the differences in the scatter plot are small. In addition, Panel (b) zooms into the share of households that purchase from the government (across all markets), an untargeted moment. Quantitatively, the model closely follows the data, and is able to recover the same patterns of consumption (e.g., that richer households consume less government milk).

We also manage to match the distance regression coefficients from the staggered-store introduction moments shown in Equation 15. The mean absolute difference between the model- and data-based regression coefficients is 0.010. Panel (c) in Figure OA-12 separates households by quintiles of the distance to the closest government store and computes the average choice probability of consuming the government good in the model and the data. Again, the model closely follows the correlation between purchases and distance.

Supply Costs and Markups We next turn to costs and markups and report our estimation results in Table 4. The mean marginal cost across products in the private market is 7.54 pesos per liter. The average retail markup $((p - c)/c)$ is around 27%. We also report our supply estimates for the two large national retailers, the remaining regional firms, and products similar in product characteristic space to the government good.

Table 4: Supply Estimation Results

	Producer Type (All Firms)		Type of Producer (Private Market)		Closeness to Govt. Product	
	Private (1)	Government (2)	National (3)	Regional (4)	Close Subs. (5)	Far Subs. (6)
<i>(a) – Cost estimates and implied consumer utility (constant MXN)</i>						
cost per liter	7.54	6.89	7.25	8.03	7.26	7.86
	[2.66, 10.64]	[6.03, 7.33]	[1.21, 10.78]	[6.18, 9.98]	[2.26, 10.13]	[3.65, 11.08]
utility at $p = c$ in monetary terms	4.72	-0.03	5.95	2.66	4.97	4.43
	[0.55, 12.22]	[-1.40, 1.48]	[0.97, 14.77]	[0.17, 5.11]	[0.71, 14.43]	[0.31, 10.28]
<i>(b) – Implied margin and markups</i>						
margin $(p - c)$	3.23	-3.03	4.08	1.81	3.32	3.14
	[1.36, 8.93]	[-3.47, -2.06]	[1.54, 10.24]	[1.30, 2.67]	[1.34, 9.49]	[1.46, 7.01]
markup $(p - c)/p$	0.27	-0.79	0.32	0.18	0.27	0.26
	[0.12, 0.61]	[-0.91, -0.52]	[0.13, 0.78]	[0.12, 0.27]	[0.12, 0.71]	[0.12, 0.55]

Notes: This table shows the average values along with the 10th to 90th percentiles of selected moments for our supply estimation results. We compute averages across household-product-markets, weighting observations by the number of consumers. Panel (a) shows the implied cost per liter (in constant MXN) and the implied utility (measured in MXN) consumers would receive if products were sold exactly at their marginal cost. Panel (b) shows the margin (measured as price minus cost) and markups (as a percent of prices). Columns (1) and (2) disaggregate by type of firm. Columns (3) and (4) disaggregate by type of retailer, where Private 1 and 2 are the national retailers and remaining firms are regional. Columns (5) and (6) divide products by their closeness to the government good in product space, where close substitutes are those that share at least three product characteristics with the government good, and far substitutes are the complements of those.

Note that our estimates suggest that without adjusting for quality, private goods are more expensive to produce than the government good. This observation is in line with anecdotal evidence: Because of the powdered milk, the government milk is cheaper to produce. To adjust for quality, we compute the utility that households would get from different products if they were sold exactly at their marginal cost and report results (relative to the outside option) measured in monetary terms. For household i of type t consuming product j , the utility at a price equal to their marginal cost is


$$\widehat{u}_{ijm}(c_{jm}) = \frac{1}{\widehat{\alpha}_t} \cdot \left(\mathbf{x}_j \cdot \widehat{\boldsymbol{\beta}}_t - \widehat{\alpha}_t \cdot c_{jm} - \widehat{\gamma}_t \cdot d_{im} \cdot \mathbb{1}_{(j=g)} + \widehat{\xi}_{jm} \right), \quad (16)$$

which we compute using our estimated cost parameters. The utility from the government good is 3 cents lower than the outside option, and is negative for 53 percent of household-months. Our estimates suggest that consumers are almost ambivalent between consuming nothing (absent idiosyncratic variance) than the government good sold at its actual cost instead of being subsidized.

On the other hand, products in the private market would yield positive utility for consumers on average. The utility from purchasing in the private market at marginal cost is negative (i.e., worse than the outside option) only for 5.8 percent of the markets.

To further decompose costs, we first show that national sellers (which are vertically integrated) have lower production costs and higher margins and markups than regional producers. Products from these big national retailers are higher quality since selling them at their marginal cost yields higher utility than those of regional firms. We then separate products using their observable characteristics (Table 1). We define two broad groups of products based on their comparability with the government milk. We define a product as a “close substitute” if it shares three or more (out of four) product characteristics with the government good. We define as “far substitutes” those products that are not close substitutes. Products closer (in characteristics space) have lower costs to produce, generate more utility, and have higher margins and markups than products further away. The latter fact suggests that, in the absence of the government good, firms would be able to extract more resources among those households who buy close substitutes.

Figure 7: Utility of Purchasing at Average Price by Production Cost



Notes: This figure plots estimates of utility and estimates of the cost per liter and shows that the government produces a low-cost, low-quality product. An observation is a product and the size of the dots are the product market shares, summing across all markets in our data. The x-axis is the average (across markets) cost per liter for the given product estimated from the model. To produce the numbers in the y-axis, we first compute the implied model utility (at the individual-product-month) if all products were sold at the same price (the average price per liter in the private market). We divide our estimates by the marginal utility of income (α) and average across individual-months. In red, we show the government good with its actual cost estimated from our data. The black line shows a local polynomial regression.

To measure the government good’s quality, we compute the utility households would get in a hypothetical world in which all products are sold exactly at the average price per liter of milk in Mexico. To do so, we use Equation 16 but instead of measuring the values at their marginal costs, we calculate the utility at $p^* = 10.95$. Figure 7 plots the average cost per liter of a given product against the average implied utility households would get by consuming each of the goods at the same price. There is a positive correlation between costs and utility, which suggests that (holding

prices fixed) consumers benefit more from products of higher cost. This scatterplot also allows us to compare households' utility when all firms sell products at the same price. The government product ranks in the bottom tier in terms of utility, further emphasizing its low-quality nature. In addition, the government product falls below the local polynomial regression, suggesting that an equal cost good produced by the private market would have a higher quality.

7 The Impact of Alternative Government Interventions

The Mexican government has several objectives when intervening in the milk market. One is to increase the consumption of milk for low-income families. The other is to provide milk at "affordable" prices and improve welfare. However, there are several ways the government could intervene to achieve these objectives. In this section, we discuss the welfare effects of alternative policy arrangements that use the same budget as the observed direct-provision scheme. We first discuss policies that involve removing direct provision altogether and instead provide cash transfers or vouchers to buy milk produced by private companies. We then discuss counterfactual arrangements that modify the direct-provision system, such as changing the government milk price. We focus our exposition on the main findings while holding fixed entry decisions, unobserved product characteristics, and firms' marginal costs. In [Appendix H](#), we provide additional details of the calculations and discuss how relaxing our assumptions would change our results.

To measure consumer surplus, we use the standard formula derived from household utility with logit errors (Equation 4) that defines it as the ex ante utility (i.e., before making choices). For household i in market m , with products A and prices \mathbf{p} , the consumer surplus formula is

$$CS_{im}(\mathbf{p}, A) = \left(\frac{1}{\alpha_{t(i)}} \right) \cdot \log \left[1 + \sum_{j \in A} \exp(\mathbf{x}_j \cdot \boldsymbol{\beta}_{t(i)} - \alpha_{t(i)} \cdot p_{jm} - \gamma_{t(i)} \cdot d_{im} \cdot \mathbb{1}_{(j=g)} + \xi_{jm}) \right] + \tau_i, \quad (17)$$

where we divide by the marginal utility of income (α) to measure surplus in monetary terms, and τ_i allows us to transfer unconditional cash (which we normalize to zero for direct provision). We compute the consumer surplus for each household-month then average across household months. We evaluate the benefit of direct provision relative to counterfactual arrangements by using compensating variations, i.e., the amount of money needed to keep households indifferent between a counterfactual set of prices and alternatives (\mathbf{p}', A') and the ones in the observed provision system (\mathbf{p}, A) . We define the compensating variation as

$$CV_{im}(\mathbf{p}, \mathbf{p}', A, A') = CS_{im}(\mathbf{p}, A) - CS_{im}(\mathbf{p}', A'),$$

with $CV_{im} > 0$ if household i in market m prefers (\mathbf{p}, A) to (\mathbf{p}', A') .

Table 5: Main Counterfactual Results

	Direct Provision				Cash Transfers				Vouchers			
	Overall (1)	Low (2)	Med (3)	High (4)	Overall (5)	Low (6)	Med (7)	High (8)	Overall (9)	Low (10)	Med (11)	High (12)
<i>(a) - Prices (Constant MXN)</i>												
Avg. transacted price per liter	8.33 (0.02)	7.72 (0.03)	8.21 (0.02)	8.95 (0.03)	10.97 (0.02)	10.90 (0.03)	10.96 (0.02)	11.12 (0.04)	10.34 (0.03)	8.79 (0.03)	11.07 (0.02)	11.22 (0.04)
Posted price per liter (private market)	10.64 (0.02)	10.55 (0.03)	10.61 (0.02)	10.79 (0.04)	10.97 (0.02)	10.90 (0.03)	10.96 (0.02)	11.12 (0.04)	11.09 (0.02)	11.02 (0.03)	11.07 (0.02)	11.22 (0.04)
<i>(b) - Market Shares</i>												
Proportion who purchase any brand	0.95 (0.00)	0.96 (0.00)	0.95 (0.00)	0.94 (0.00)	0.92 (0.00)	0.93 (0.00)	0.91 (0.00)	0.91 (0.00)	0.96 (0.00)	0.97 (0.00)	0.96 (0.00)	0.95 (0.00)
<i>(c) - Change in Consumer Surplus Relative to Direct Provision</i>												
Change in consumer surplus (ΔCS)	-	-	-	-	-0.08 (0.01)	-0.18 (0.01)	-0.07 (0.01)	0.07 (0.01)	-0.13 (0.01)	-0.22 (0.01)	-0.12 (0.01)	0.03 (0.01)
Transfer benefit (τ) or Voucher benefit (VB)	-	-	-	-	0.68 (0.00)	0.68 (0.00)	0.68 (0.00)	0.68 (0.00)	0.73 (0.00)	0.74 (0.00)	0.73 (0.00)	0.73 (0.00)
Direct benefit (DB)	-	-	-	-	0.56 (0.01)	0.66 (0.01)	0.55 (0.01)	0.41 (0.01)	0.56 (0.01)	0.66 (0.01)	0.55 (0.01)	0.41 (0.01)
Indirect benefit (IB)	-	-	-	-	0.20 (0.00)	0.20 (0.00)	0.20 (0.00)	0.20 (0.00)	0.30 (0.00)	0.30 (0.00)	0.30 (0.00)	0.29 (0.00)

Notes: This table shows results for selected counterfactuals. We report the simulation details in [Appendix H](#). Columns (1) to (4) report outcomes of the observed direct-provision system. Columns (5) to (8) show outcomes that would arise with direct cash transfers. Columns (9) to (12) show outcomes for vouchers that subsidize all goods for all consumers.

7.1 Alternative Interventions

Unrestricted cash transfers Cash transfers are equivalent, in this model with no income effects, to eliminating the direct-provision scheme, finding the new equilibrium in the private market that would arise, and then transferring the government budget to the six eligible types of households via a lump-sum payment per household-month. We do this as follows. We constrain transfers to be the same among all households-months within a calendar year. Households get a transfer if at least one member of their household would have been eligible to get government milk under the status quo. Households get the transfer regardless of whether they purchase milk. We present our results for cash transfers in columns (5) through (8) of Table 5.

In the absence of the government, the *transacted* price per liter⁴⁴ would be, on average, \$2.57 higher, from \$8.33 to \$10.90. This amounts to an increase of 31%. The difference arises from private companies increasing their *posted* prices by 2.4% (from \$10.64 to \$10.90) as a result of facing less competition (no government good available), but also by being forced to reallocate their purchases to more expensive milk products, since the cheapest good is no longer available. With these higher prices, the share of households that purchase milk within a given month would decrease by 3 pp (or 3.1% relative to direct provision), from 95% to 92%. The poorest set of households, which are the most sensitive to the price increase, are likely to stop consuming, with a decrease of 3.1% in the share of households that purchase milk.

Interestingly, the consumer surplus with cash is 8¢ *lower* than the one with direct provision.⁴⁵ To put this number in perspective, this is 12% of the per household-month cost of the program (68¢). To further explain the changes in surplus, we use an off-equilibrium calculation and rewrite the difference between the surplus under direct provision and cash transfers in three parts. Denote by (p, A) the equilibrium prices and the alternatives available with direct provision, by g the government good, by p_{CT} the prices with cash transfers, and by τ the per household transfer equivalent to the cost of direct provision. The “direct” benefit (DB) from the government good being available, holding prices fixed, is

$$DB_i = CS_i(p, A) - CS_i(p, A \setminus \{g\}).$$

The “indirect” benefit (IB) that arises from the changes in prices—the competition effect—is

$$IB_i = CS_i(p, A \setminus \{g\}) - CS_i(p_{CT}, A \setminus \{g\}).$$

The cash transfer equivalent of the cost of direct provision per person is τ_i . Intuitively, the direct benefit operates in part through preference parameters related to milk product characteristics and distance to the store, while the indirect benefit operates through preference parameters related to price sensitivity, and the supply-side competition effect.

⁴⁴Transacted price is the weighted average of prices for different milk products, weighted by their market share.

⁴⁵In Table OA-9, we show that the increase in consumer surplus is a redistribution from producer profits.

In sum, the change in consumer surplus can be written as

$$\Delta CS_i = \tau_i - DB_i - IB_i \quad (18)$$

where $\Delta CS_i < 0$ if the surplus of household i in a given market under direct provision is higher than that of cash. We drop all market subscripts for simplicity, but perform our computations market by market. Panel (c) of Table 5 displays the four components of Equation 18.

First, note that cash transfers hurt poor households and benefit richer ones. Second, our decomposition shows that the gain in surplus for poor individuals is due to both direct and indirect benefits—each separately would give a lower benefit than cash. In contrast, richer households value the government good less and have lower gains from the lower prices in the private market (since they are less price elastic). This discrepancy highlights the role product characteristics, distance to the store, and price policy play in screening households to ensure that the government targets those with lower incomes.

Naturally, cash transfers may be associated with a lower cost of public funds than direct provision, which implies that for every \$1 through direct provision we could distribute $\$(1+x)$ through cash. In our calculations, we use the average cost per liter for the government, taking into consideration the costs of setting up and administering stores and the costs associated with producing milk. It could well be true that we miss some costs of direct provision. A back-of-the-envelope calculation shows that for the average consumer surplus to be equal in cash and direct provision, we would need $x = 8\%$. In the case of poorer households, we need $x = 6\%$.

Universal Vouchers for All Goods The second counterfactual we study is to use the fixed budget to provide vouchers that households can redeem in the private market for milk. Formally, a voucher v decreases the price of a given milk product by $p - v$. We first focus on completely untargeted vouchers that are distributed to all eligible households. After showing the results for this scenario, we investigate the benefits of targeting by subsidizing a subset of the goods available in a market (e.g., the top products purchased by the poor), a subset of households (e.g., those that, in the direct provision world, purchase from the government), or both. We present results in columns (9) to (12) of Table 5.

Vouchers generate upward pressure on prices. The extent to which vouchers are passed through to higher producer prices depends on the slope of residual demand (Weyl and Fabinger, 2013). The voucher size per household-month, given the fixed government budget, is on the order of 73¢, which implies that a total of 19¢ is passed through to higher producer prices. In consequence, for every \$1 returned to households via a voucher, posted prices increase by 26¢, on average. Table OA-9 shows that most of the voucher gains are captured by producers in terms of profits. Prices would be 24% higher with vouchers than under direct provision.

In terms of consumption, we find that direct provision and vouchers generate similar con-

sumption levels—a common finding in the literature that we replicate in our setting (Aker, 2017; Hidrobo et al., 2014; Cunha, 2014). On average, vouchers induce 96% of households to purchase in a given month. Relative to direct provision, both rich and poor households have comparable levels of consumption, with 96% purchasing milk in a given month. When compared with cash, poor households increase their consumption more than richer ones.

Turning to consumer surplus, we find that households are worse off with vouchers than with direct provision. Consumer surplus with vouchers is 13¢ (or 20% of the per-household program cost) lower than with direct provision. Denote by p_V the equilibrium posted prices with vouchers, and by $p - v$ the consumer prices under the voucher system.⁴⁶ Similar to our decomposition above, we can decompose the change in surplus from direct provision to vouchers in three parts. First, the direct benefit from receiving the government good, holding prices fixed, is

$$DB_i = CS_i(p, A) - CS_i(p, A \setminus \{g\}).$$

Second, the slightly modified indirect benefit that arises from competition, now holding transfers at the voucher level, is

$$IB_i = CS_i(p - v, A \setminus \{g\}) - CS_i(p_V - v, A \setminus \{g\}).$$

Third, the voucher benefit (VB), which is the difference in surplus from receiving equivalent-sized vouchers instead of the government good, holding prices fixed, is

$$VB_i = CS_i(p - v, A \setminus \{g\}) - CS_i(p, A \setminus \{g\}).$$

The change in surplus is given by $\Delta CS_i = VB_i - DB_i - IB_i$.

Using this decomposition, we document that despite using the same budget and inducing somewhat similar consumption levels, the surplus gain from lower prices in the private market more than offsets the loss of the “freedom to choose” when replaced by a government good that is of lower quality. Consumer surplus is lower than that one from the direct-provision system. Again, poorer households are those that benefit the most, while richer households would rather get the vouchers. Note that the indirect benefit from vouchers is higher than with unrestricted cash transfers, since equilibrium prices are even higher with vouchers than with cash.

Vouchers for Selected Goods or Selected Households Unlike cash transfers, vouchers could be made to apply to only a subset of milk products. For instance, this targeting feature could be useful in order to focus assistance on low-income households. As we show next, the problem is that narrowing the scope of the vouchers for a subset of goods gives the sellers of those goods

⁴⁶Note that in the simplest case we can equivalently write utility for the outside good as $u_0 = \alpha y$, and for milk product j as $u_j = \theta_j + \alpha(y - p_j)$, where y is income, p_j is the price of good j , and α is the marginal utility of income. A voucher of v pesos if consumers buy milk amounts to increasing income by that much: $u_0 = \alpha y$; $u_j = \theta_j + \alpha(y - p_j + v)$, or by changing the price from p to $p - v$.

more market power. Consistent with this, [Hastings and Washington \(2010\)](#) and [Meckel \(2020\)](#) documented that in the US, posted prices did indeed increase with the voucher cycle.

We compute the effect of two distinct targeting policies. In the first, we subsidize the top five milk products that low-income households consume in an attempt to target through choices. In the second exercise, we compute the market outcomes that would arise if we managed to target by income directly and distribute vouchers exclusively to low-income households. This second option is likely infeasible in practice; we use it as a comparison benchmark only.

Results for the first exercise are reported in columns (1) to (4) of [Table 6](#). We find that posted prices increase more than they do with universal vouchers, and the average consumer surplus relative to direct provision decreases more than with universal vouchers. Nonetheless, slightly more money goes to the poor. Richer households are still better off with universal vouchers, as many of them would buy the alternatives not subsidized under this counterfactual.

Results for the second exercise are reported in columns (5) to (8). Vouchers with perfect screening on income would generate 9¢ less in consumer surplus than direct provision, but redistribute more of the surplus gains to the poor. Interestingly, the average price paid by poor households would be \$2 higher than that with direct provision. Nonetheless, the broader set of subsidized products implies that low-income households would be better off with perfect targeting than in the direct-provision world.

Table 6: Additional Voucher Counterfactual Results

	Who receives vouchers: Products subsidized:	Eligible Households Top 5 Among Poor				Poor Households All Goods			
		Overall (1)	Low (2)	Med (3)	High (4)	Overall (5)	Low (6)	Med (7)	High (8)
<i>(a) - Prices (Constant MXN)</i>									
Avg. transacted price per liter		10.54 (0.02)	10.44 (0.03)	10.51 (0.03)	10.62 (0.05)	10.31 (0.03)	8.78 (0.03)	11.07 (0.02)	11.21 (0.04)
Posted price per liter (private market)		11.24 (0.02)	11.19 (0.03)	11.24 (0.03)	11.36 (0.05)	11.08 (0.02)	11.02 (0.03)	11.07 (0.02)	11.21 (0.04)
<i>(b) - Market Shares</i>									
Proportion who purchase any brand		0.95 (0.00)	0.96 (0.00)	0.95 (0.00)	0.95 (0.00)	0.95 (0.00)	0.99 (0.00)	0.92 (0.00)	0.92 (0.00)
<i>(c) - Change in Consumer Surplus Relative to Direct Provision</i>									
Change in consumer surplus (ΔCS)		-0.30 (0.01)	-0.39 (0.01)	-0.30 (0.01)	-0.15 (0.01)	-0.09 (0.02)	1.20 (0.02)	-0.83 (0.01)	-0.68 (0.01)
Voucher benefit (VB)		0.59 (0.00)	0.60 (0.01)	0.59 (0.00)	0.60 (0.01)	0.76 (0.02)	2.17 (0.01)	-0.00 (0.00)	-0.00 (0.00)
Direct benefit (DB)		0.56 (0.01)	0.66 (0.01)	0.55 (0.01)	0.41 (0.01)	0.56 (0.01)	0.66 (0.01)	0.55 (0.01)	0.41 (0.01)
Indirect benefit (IB)		0.34 (0.00)	0.33 (0.01)	0.34 (0.00)	0.34 (0.01)	0.29 (0.00)	0.30 (0.00)	0.28 (0.00)	0.27 (0.00)

Notes: This table shows results for selected voucher counterfactuals. We report the simulation details in [Appendix H](#). Columns (1) to (4) report outcomes of subsidizing only the top products among the poor. Columns (5) to (8) show outcomes of subsidizing only poor households.

Price Caps in the Private Market Setting maximum prices can increase consumption without having to go through direct provision. The downside is that price caps seem *ex ante* infeasible, since they require that regulators have information on the cost structure of private providers or risk the exit of some products if producers find them to no longer be profitable. In such cases, the loss in variety may hurt households more than the lower prices in the market.

We focus on capping prices directly via a uniform maximum price (e.g., $p_{jm} \leq x$). In addition, we return the government budget to eligible households through unconditional cash.⁴⁷ In the fully unconstrained world (i.e., large x), outcomes are identical to those that arise with cash transfers. There are two countervailing forces that take a primary role with stricter caps. Holding the available alternatives fixed, decreasing the cap has a positive effect on surplus because households have to pay lower prices for the same products. However, setting a stringent cap may induce sellers to not provide certain goods if the cap is lower than their marginal cost. Such a decrease in the product variety decreases household utility.

We graphically show results for different price cap values in Figure 8. Moving to the left on the x-axis imposes a stricter price cap; moving to the right gets us to the cash transfers counterfactual. These graphs demonstrate that consumption and consumer surplus are not monotonous on the price cap level, since the two countervailing forces detailed above come into play at different values of the price cap. As products exit the market (as shown by the gray line in Figure 8(a)), households substitute toward the remaining alternatives. The lower the number of products, the more likely that households will substitute towards the outside option. Figures 8(a) and 8(b) show that as a tool to incentivize consumption, direct provision and vouchers work better than capping prices.

The direct-provision system generates approximately the same consumer surplus, on average, as a cap around the 80th percentile of the observed private-market price distribution at about \$12.50 (Figure 8(c)). For poor households to be indifferent, we would require a stricter cap of around the 60th percentile at about \$11.75 (Figure 8(d)). Lower caps (than the ones defined above) may induce higher consumer surplus. However, such stricter caps seem unlikely to be implemented without a substantial backlash from private sellers. These caps wipe out profits in the private market.

7.2 Changes in the Observed Direct-provision Scheme


Thus far, our results suggest that direct provision works as a tool to increase consumption and consumer surplus when compared with natural policy alternatives. We now focus on how to improve the current direct-provision system by changing the price of government milk.

Uniform Price Changes in the Government Good We consider changes to the price of the government milk that apply to all markets. For instance, the government may sell at a higher

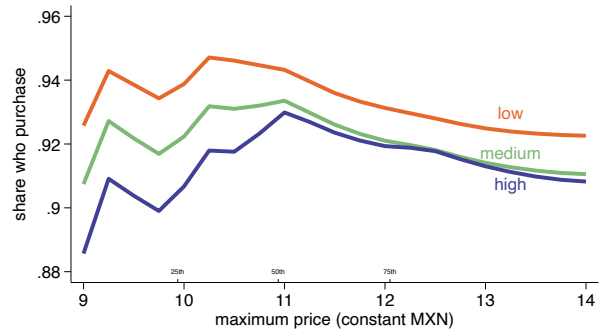
⁴⁷In the Appendix, we discuss caps that limit markups instead via a uniform maximum (e.g., $(p_{jm} - c_{jm})/c_{jm} \leq x$).

Figure 8: Price Caps in the Private Market Counterfactual Results


(a) Share of Households that Purchase (left) and Alternatives Available (right)




(b) Share of Households that Purchase by Socioeconomic Status



(c) Consumer Surplus (Relative to Direct Provision)



(d) Consumer Surplus by Socioeconomic Status (Relative to Direct Provision)



Notes: This figure plots selected market outcomes for the price cap counterfactuals. The x-axis represents the maximum price per liter allowed. Panels (a) and (b) show on the left axis the share of household-months with at least one purchase. Panel (a) also includes in gray the share of alternatives from which households choose, relative to the non-constrained equilibrium. A value of 1 means 100% of existing alternatives today; a value of 0.4 means only 40% of them. Panels (c) and (d) show the consumer surplus relative to the direct-provision world. Positive values show cases in which price caps generate larger consumer surplus gains than those of direct provision.

price, thus liberating part of their budget, which they can then transfer through unrestricted cash to eligible households. At the other extreme, the government may provide milk for free, albeit by selling lower quantities (or raising additional taxes) to keep within the same government budget. In this latter case there will be excess demand, which we allow by raising additional taxes with a cost of public funds of 30%.

Formally, denote by p_g the nominal price per liter for the government good. We study changes of the following form:

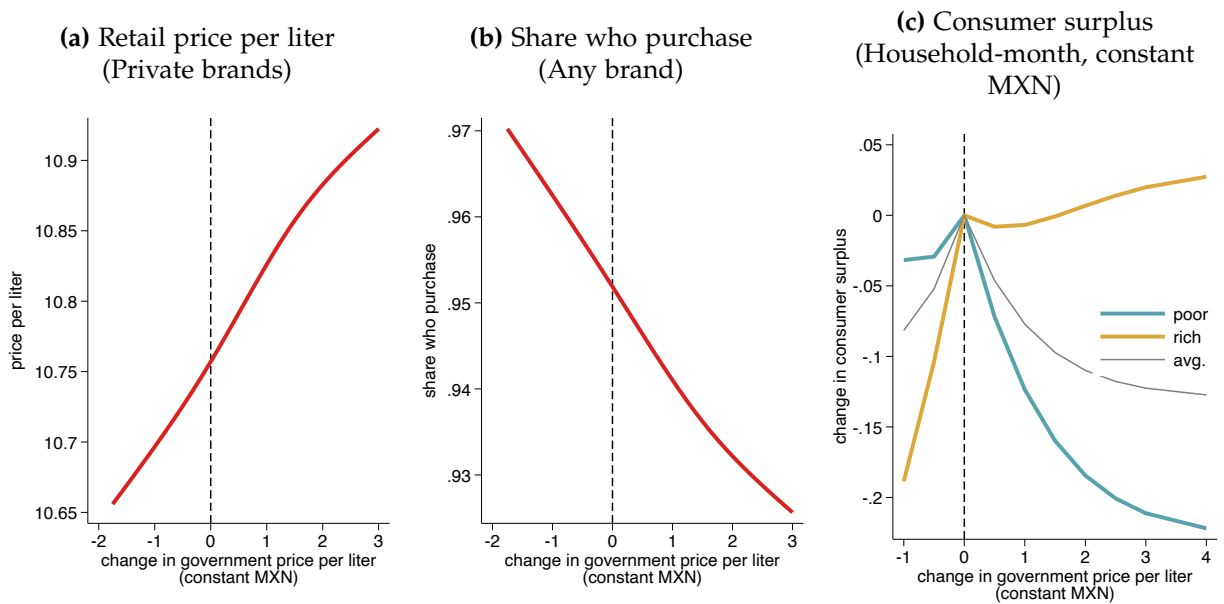
$$p'_g(x) = \max\{p_g + x, 0\}, \quad (19)$$

which we then convert to constant MXN. With $x = 0$, we represent outcomes under the observed direct-provision system. When x increases, we move toward the unrestricted cash-transfers counterfactual world and with $x = -p_g$ we are in the fully subsidized transfers world. We present our results in Figure 9. As these figures show, small increases in the price of the government

good may be beneficial, on average, since they retain most of the direct and indirect benefits we have discussed extensively, while being able to redistribute part of the money through cash.

Interestingly, consumer surplus is not necessarily monotone in prices. Slightly increasing the government milk price implies that households can then receive some cash transfers by using the liberated budget. However, the loss in surplus arising from the higher government prices—and the higher private prices resulting from this—eventually becomes larger than the cash-transfer benefit. Increasing the price of the government good would decrease the average consumer surplus. On top of that, increasing the price of the government good hurts poor consumers, for whom the direct benefit of direct provision is higher, and who would then be hurt by the slightly higher prices. Finally, note that the current system attains close to the total surplus attainable from changing the price of the government good. Decreasing the price of the government, on the other hand, hurts households as the system becomes more difficult to fund.

Figure 9: Optimizing via the government price (uniform price changes in constant MXN)



Notes: This figure shows the estimation results for the counterfactuals that change the price of the government good. The x-axis in each of these graphs is the nominal difference in the price of the government good, relative to the status quo pricing. Panel (a) shows the retail price per liter in the private market. Panel (b) shows the share of households that would purchase any brand. Panel (c) shows the consumer surplus, disaggregating between rich and poor households.

8 Conclusion

In this paper, we study the demand- and supply-side responses to government direct provision of milk in Mexico. We show that there are two channels through which the direct-provision system changes consumer welfare. The first one is a direct channel: government milk, at the set price, is preferred by some households and its availability is therefore beneficial for them. The second

one is indirect: By disciplining the market power of private suppliers, this causes a decrease in market prices and generates spillovers to households that do not purchase the government good. Direct provision serves as a second-best way to transfer money from producers to consumers.

Governments often use the direct provision of goods through ration stores with the sole intention of increasing access. Yet we show that they could have the added benefit of decreasing prices by generating stronger competition, especially in markets with market power like the one we study. We also show that it had the benefit of targeting the subsidy to those most in need through the location of stores and micro-ordeal costs.

This paper abstracts from several aspects of the market. First, we abstract away from any effect that direct government provision can have on exit of private suppliers. Second, we do not take into consideration changes in the profits of intermediate suppliers, such as small milk producers. Nonetheless, we believe we capture essential trade offs in the decision to provide directly or not. We believe that the trade off we study are present in other markets as well, which potentially include development banking, health, among others.

Should the government sell you goods? It depends, but should not be discarded a priori. We show that direct provision of a lower-quality good—even if produced at higher cost—can increase consumer welfare when firms have market power and when it is hard to screen consumers.

References

- Aguilar Esteva, A., E. Gutierrez, and E. Seira (2019). The Effectiveness of Sin Food Taxes: Evidence from Mexico. *Working paper*, 1–46.
- Aker, J. (2017). Cash or Coupons? Testing the Impacts of Cash versus Vouchers in the Democratic Republic of Congo. *World Bank Economic Review* 31(1), 44–70.
- Alatas, V., A. V. Banerjee, R. Hanna, B. A. Olken, R. Purnamasari, and M. Wai-Poi (2016). Self-Targeting: Evidence from a Field Experiment in Indonesia. *Journal of Political Economy* 124(2).
- AMAI (2018). Regla Nivel Socio Económico Asociación Mexicana de Agencias de Inteligencia de Mercado y Opinión. Technical report.
- Andreyeva, T., M. W. Long, and K. D. Brownell (2010). The Impact of Food Prices on Consumption: A Systematic Review of Research on the Price Elasticity of Demand for Food. *American Journal of Public Health* 100(2), 216–222.
- Aragón Mladosich, E. R. and J. A. Gómez-Ibáñez (2004). Liconsa and the Program of Social Assistance for Milk. *Harvard Kennedy School of Government Case Program* (CR14-04-1766.0).
- Atal, J. P., J. I. Cuesta, F. González, and C. Otero (2020). The Economics of the Public Option: Evidence from Local Pharmaceutical Markets. *Working paper*.
- Atkin, D., B. Faber, and M. Gonzalez-Navarro (2016). Retail Globalization and Household Welfare: Evidence from Mexico. *Journal of Political Economy* 126(1), 1–73.

- Atkinson, A. B. and J. E. Stiglitz (1976). The Design of Tax Structure: Direct versus Indirect Taxation. *Journal of Public Economics* 6(1-2), 55–75.
- Banerjee, A. V., R. Hanna, J. Kyle, B. A. Olken, and S. Sumarto (2018). Private Outsourcing and Competition: Subsidized Food Distribution in Indonesia. *Journal of Political Economy* 127(1), 101–137.
- Bearse, P., G. Glomm, and E. Janeba (2000). Why Poor Countries Rely Mostly on Redistribution In-kind. *Journal of Public Economics* 75(3), 463–481.
- Berry, S. T. (1994). Estimating Discrete-Choice Models of Product Differentiation. *RAND Journal of Economics* 25(2), 242–262.
- Berry, S. T. and P. A. Haile (2020). Nonparametric Identification of Differentiated Products Demand Using Micro Data. *Working paper*.
- Berry, S. T., J. Levinsohn, and A. Pakes (1995). Automobile Prices in Market Equilibrium. *Econometrica* 63(4), 841–890.
- Berry, S. T., J. Levinsohn, and A. Pakes (2004). Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market. *Journal of Political Economy* 112(1), 68–105.
- Besley, T. and S. Coate (1991). Provision of Private Goods and the Redistribution of Income. *American Economic Review* 81(4), 979–984.
- Besley, T. and R. Kanbur (1988). Food Subsidies and Poverty Alleviation. *Economic Journal* 98(392), 701–719.
- Bhattacharya, S., V. L. Falcao, and R. Puri (2017). The Public Distribution System in India (1939–2001). In H. Alderman, U. Gentilini, and R. Yemtsov (Eds.), *The 1.5 Billion People Question: Food, Vouchers, or Cash Transfers?*, Chapter 2, pp. 43–105. The World Bank.
- Blackorby, C. and D. Donaldson (1988). Cash versus Kind, Self-Selection, and Efficient Transfers. *American Economic Review* 78(4), 691–700.
- Bresnahan, T. F. (1981). Departures from Marginal-cost Pricing in the American Automobile Industry: Estimates for 1977–1978. *Journal of Econometrics* 17(2), 201–227.
- Busso, M. and S. Galiani (2019). The Causal Effect of Competition on Prices and Quality: Evidence from a Field Experiment. *American Economic Journal: Applied Economics* 11(1), 33–56.
- Clemens, J. and J. D. Gottlieb (2016). Influence on Private Physician Payments. *Journal of Political Economy* 125(1), 1–39.
- Colchero, M. A., J. C. Salgado, M. Unar-Munguía, M. Hernández-Ávila, and J. A. Rivera-Dommarco (2015). Price Elasticity of the Demand for Sugar Sweetened Beverages and Soft Drinks in Mexico. *Economics and Human Biology* 19, 129–137.
- Collinson, R. and P. Ganong (2018). How Do Changes in Housing Voucher Design Affect Rent and Neighborhood Quality? *American Economic Journal: Economic Policy* 10(2), 62–89.
- Cunha, J. M. (2014). Testing Paternalism: Cash versus in-kind Transfers. *American Economic Journal: Applied Economics* 6(2), 195–230.

- Cunha, J. M., G. de Giorgi, and S. Jayachandran (2019). The Price Effects of Cash Versus In-Kind Transfers. *Review of Economic Studies* 86(1), 240–281.
- Currie, J. and F. Gahvari (2008). Transfers in Cash and In-Kind: Theory Meets the Data. *Journal of Economic Literature* 46(2), 333–383.
- Davis, C., D. Blayney, J. Cooper, and S. Yen (2009). An Analysis of Demand Elasticities for Fluid Milk Products in the U.S. In *International Association of Agricultural Economists Meeting*, Number 202, pp. 1–24.
- Diario Oficial de la Federación (2014). Reglas de Operacion del Programa de Abasto Social de Liconsa. Technical report, Secretaria de Desarrollo Social.
- Dinerstein, M., C. Neilson, and S. Otero (2020). The Equilibrium Effects of Public Provision in Education Markets: Evidence from a Public School Expansion Policy. *Working paper*.
- Dupas, P., V. Hoffmann, M. Kremer, and A. P. Zwane (2016). Targeting Health Subsidies Through a Non-price Screening Mechanism: Evidence from a Randomized Controlled Trial in Rural Kenya. *Science* 353(6302), 889–896.
- Epple, D. and R. E. Romano (1996). Public Provision of Private Goods. *Journal of Political Economy* 104(1), 57–84.
- Euromonitor International (2015). Drinking Milk Products in Mexico. Technical report.
- FAO (2019). *The State of Food Security and Nutrition in the World*. FAO.
- Filmer, D., J. Friedman, E. Kandpal, and J. Onishi (2021). Cash Transfers, Food Prices, and Nutrition Impacts on Ineligible Children. *Review of Economics and Statistics*.
- Gadenne, L. (2020). Can Rationing Increase Welfare? Theory and an Application to India’s Ration Shop System. *American Economic Journal: Economic Policy*.
- Gadenne, L., S. Norris, M. Singhal, and S. Sukhtankar (2020). In-Kind Transfers as Insurance. *Working paper*.
- Gahvari, F. and E. Mattos (2007). Conditional Cash Transfers, Public Provision of Private Goods, and Income Redistribution. *American Economic Review* 97(1), 491–502.
- Goldberg, P. K. (2018). Firms in Developing Countries: Can Trade Policy Serve as Competition Policy? In *NBER Development Economics Master Lecture, Summer Institute*.
- Gould, B. W. (1996). Factors Affecting U.S. Demand for Reduced-Fat Fluid Milk. *Journal of Agricultural and Resource Economics* 21(1), 68–81.
- Handbury, J. and S. Moshary (2020). School Food Policy Affects Everyone: Retail Responses to the National School Lunch Program. *Working paper*, 1–40.
- Hansen, L. P. (1982). Large Sample Properties of Generalized Method of Moments Estimators. *Econometrica* 50(4), 1029–1054.
- Hastings, J. S. and E. Washington (2010). The First of the Month Effect: Consumer Behavior and Store Responses. *American Economic Journal: Economic Policy* 2(2), 142–162.

- Hidrobo, M., J. Hoddinott, A. Peterman, A. Margolies, and V. Moreira (2014). Cash, Food, or Vouchers? Evidence from a Randomized Experiment in Northern Ecuador. *Journal of Development Economics* 107, 144–156.
- Ivaldi, M., F. Jenny, and A. Khimich (2016). Cartel Damages to the Economy: An Assessment for Developing Countries. In F. Jenny and Y. Katsoulacos (Eds.), *Competition Law Enforcement in the BRICS and in Developing Countries*, pp. 103–134. Springer.
- Kumler, T., E. Verhoogen, and J. Frías (2020). Enlisting Employees in Improving Payroll-Tax Compliance: Evidence from Mexico. *Review of Economics and Statistics*, 1–45.
- Leung, J. H. and H. K. Seo (2019). How Do Government Transfer Payments Affect Retail Prices and Welfare? Evidence from SNAP. *Working paper*, 1–34.
- Liconsa (2015). Programa Institucional.
- Liconsa (2017). Requisitos para venderle leche a Liconsa.
- Lieber, E. M. J. and L. M. Lockwood (2019). Targeting with In-Kind Transfers: Evidence from Medicaid Home Care. *American Economic Review* 109(4), 1461–1485.
- Losada, H., R. Bennett, J. Cortes, J. Vieyra, and R. Soriano (2000). The Historical Development of the Mexico City Milk Supply System: Local and Global Contradictions. *Habitat International* 24(4), 485–500.
- Meckel, K. (2020). Is the Cure Worse than the Disease? Unintended Consequences of Fraud Reduction in Transfer Programs. *American Economic Review* 110(6), 1821–1865.
- Miller, N. H. and M. C. Weinberg (2017). Understanding the Price Effects of the MillerCoors Joint Venture. *Econometrica* 85(6), 1763–1791.
- Moshary, S. and G. Ilanes (2020). Market Structure and Product Assortment: Evidence from a Natural Experiment in Liquor Licensure. *Working paper*, 1–62.
- Nagavarapu, S. and S. Sekhri (2016). Informal Monitoring and Enforcement Mechanisms in Public Service Delivery: Evidence from the Public Distribution System in India. *Journal of Development Economics* 121, 63–78.
- Nevo, A. (2000). A Practitioner’s Guide to Estimation of Random Coefficients Logit Models of Demand. *Journal of Economics & Management Strategy* 9(4), 513–548.
- Nichols, A. L. and R. J. Zeckhauser (1982). Targeting Transfers through Restrictions on Recipients. *American Economic Review* 72(2), 372–377.
- Petrin, A. (2002). Quantifying the Benefits of New Products: The Case of the Minivan. *Journal of Political Economy* 110(4), 705–729.
- Petrin, A. and K. Train (2010). A Control Function Approach to Endogeneity in Consumer Choice Models. *Journal of Marketing Research* 47(1), 3–13.
- Rivera, J. A., T. Shamah, S. Villalpando, and E. Monterrubio (2010). Effectiveness of a Large-scale Iron-fortified Milk Distribution Program on Anemia and Iron Deficiency in Low-income Young Children in Mexico. *American Journal of Clinical Nutrition* 91(2), 431–439.

- Schmit, T. M., D. Dong, C. Chung, H. M. Kaiser, and B. W. Gould (2002). Identifying the Effects of Generic Advertising on the Household Demand for Fluid Milk and Cheese: A Two-step Panel Data Approach. *Journal of Agricultural and Resource Economics* 27(1), 165–186.
- SEDESOL (2013). Procedimiento de selección de los 400 municipios prioritarios de la Cruzada Nacional Contra el Hambre. Technical report.
- Soto Romero, J. M., E. R. Aragón Mladovich, H. J. Fuentes Castro, M. Á. Gutiérrez Banegas, J. E. Lozano Ortega, E. Ramírez Ramírez, I. Schockaert, and A. Zambrano Osorio (2004). Evaluación de Resultados del Programa de Abasto Social de Leche, a cargo de Liconsa, S.A. de C.V. Technical report, Tecnológico de Monterrey.
- Sylvia, S., X. Ma, Y. Shi, S. Rozelle, and C.-Y. C. Lin Lawell (2015). Ordeal Mechanisms, Information, and the Cost-Effectiveness of Subsidies: Evidence from Subsidized Eyeglasses in Rural China. *Working paper*.
- Tarozzi, A. (2005). The Indian Public Distribution System as provider of food security: Evidence from child nutrition in Andhra Pradesh. *European Economic Review* 49(5), 1305–1330.
- Villalpando, S., T. Shamah, J. A. Rivera, Y. Lara, and E. Monterrubio (2006). Fortifying Milk with Ferrous Gluconate and Zinc Oxide in a Public Nutrition Program Reduced the Prevalence of Anemia in Toddlers. *Journal of Nutrition* 136(10), 2633–2637.
- Villas-Boas, S. B. (2007). Vertical Relationships Between Manufacturers and Retailers: Inference with Limited Data. *Review of Economic Studies* 74(2), 625–652.
- Weyl, E. G. and M. Fabinger (2013). Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition. *Journal of Political Economy* 121(3), 528–583.
- Yúñez-Naude, A. (2003). The Dismantling of CONASUPO, a Mexican State Trader in Agriculture. *World Economy* 26(1), 97–122.

Should the Government Sell You Goods? Evidence From the Milk Market in Mexico

Diego Jiménez-Hernández and Enrique Seira

Appendix - For Online Publication

Contents

Appendix A Additional Tables and Figures	OA - 2
A.1 Additional Tables	OA - 2
A.2 Additional Figures	OA - 8
Appendix B Additional Information of the Government Intervention	OA - 16
Appendix C Pilot Study for the Government Good's Quality	OA - 19
Appendix D Dataset construction	OA - 22
D.1 Main Dataset Construction	OA - 22
D.2 Hexagon Dataset Construction	OA - 24
Appendix E The Crusade Against Hunger and Distance to Ration Stores	OA - 26
E.1 Municipality selection criteria	OA - 26
E.2 Within-municipality selection of ration stores location	OA - 28
Appendix F Stylized Model Description and Intuition	OA - 31
F.1 Model Set-Up	OA - 31
F.2 Parametrization and Initial Results	OA - 32
F.3 Comparative Statics	OA - 33
Appendix G Model Estimation Details	OA - 36
G.1 Implementation details	OA - 36
G.2 Algorithm Description	OA - 37
G.3 Demand Estimation	OA - 41
G.4 Supply Estimation	OA - 45
G.5 Influence of Moment Parameters	OA - 45
Appendix H Counterfactual Estimation Procedure and Additional Results	OA - 47
H.1 Estimation Procedure	OA - 47
H.2 Robustness	OA - 49

Appendix A Additional Tables and Figures

A.1 Additional Tables

Table OA-1: Description of Government Stores (2010 - 2014)

	2010 (1)	2011 (2)	2012 (3)	2013 (4)	2014 (5)
<i>Panel A. All sample</i>					
Active stores	7,837	7,831	7,968	9,093	9,563
Opened stores	197	143	241	1,178	687
Closed stores	117	151	100	63	208
<i>Panel B. Restricting to metropolitan areas in consumer panel sample</i>					
Active stores	3,277	3,296	3,371	3,550	3,740
Opened stores	76	60	103	208	229
Closed stores	26	38	22	26	38

Notes: Authors own calculations using government store data described in Section 3. Each column represents a different year. The number of active stores is the count of stores that were active by December of the corresponding year. Opened (closed) stores denote the count of stores that open (closed) at any point in time between January and December.

Table OA-2: Placebos: Staggered Store Introduction Results in Other Markets

	At least one purchase			Total units per household			Price per unit		
	Sodas (1)	Yogurt (2)	Cereal (3)	Sodas (4)	Yogurt (5)	Cereal (6)	Sodas (7)	Yogurt (8)	Cereal (9)
Distance to closest govt. store	-0.003 (0.002)	0.001 (0.005)	-0.007 (0.006)	-0.217 (0.376)	0.005 (0.017)	-0.004 (0.012)	0.040 (0.070)	0.514 (0.701)	0.673 (0.834)
One std. dev. increase distance ($\sigma \times \beta$)	-0.02	0.01	-0.04	-1.25	0.03	-0.03	0.23	2.98	3.89
Mean of dependent variable	0.97	0.74	0.58	22.92	0.87	0.59	8.27	56.35	63.25
Municipality-month FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Household-type controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Household census tract controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	362,847	362,847	362,847	362,847	362,847	362,847	351,009	264,171	211,747
R-squared	0.116	0.115	0.115	0.286	0.103	0.133	0.144	0.088	0.053

Notes: Robust standard errors shown in parenthesis. All regressions include month and municipality fixed effects and probability weights using our household panel. An observation is a household-month. The dependent variable for Columns (1) to (3) are categorical variables equal to one if the given household purchase at least one unit of sodas, yogurt, and cereal, respectively. The dependent variable for Columns (4) and (6) are the total liters per household in the case of sodas and yogurt, and the total kilograms in the case of cereal. The dependent variable for Columns (6) to (8) is the weighted average price per unit (liters in the case of sodas and cereals, and kilograms in the case of yogurt) paid for the sodas, yogurt, and cereal.

Table OA-3: Robustness of Staggered Store Introduction Results

Source:	At least one liter purchased			Total liters per household			Price per liter (2010 constant MXN)		
	Govt (1)	Priv (2)	Any (3)	Govt (4)	Priv (5)	Any (6)	Govt (7)	Priv (8)	Any (9)
<i>(a) - Using Logarithm of Distance as Dependent Variable</i>									
Log. of min. distance	-0.012 (0.005) **	0.006 (0.004)	-0.001 (0.002)	-0.217 (0.089) **	0.079 (0.102)	-0.137 (0.107)	0.001 (0.001)	0.031 (0.014) **	0.083 (0.040) **
Municipality-month FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Household-type controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Household census tract controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	368,762	368,762	368,762	368,762	368,762	368,762	53,632	318,553	344,672
R-squared	0.352	0.135	0.103	0.294	0.298	0.281	0.769	0.415	0.376
<i>(b) - Using Information at the Municipality-month Level</i>									
Govt. stores per 10k elig. households	0.017 (0.004) ***	-0.005 (0.003)	0.002 (0.003)	0.204 (0.056) ***	0.052 (0.184)	0.256 (0.187)	-0.064 (0.075)	-0.129 (0.033) ***	-0.126 (0.028) ***
One std. dev. increase in govt. stores ($\sigma \times \beta$)	0.043	-0.013	0.005	0.504	0.130	0.634	-0.159	-0.318	-0.312
Mean of dependent variable	0.220	0.843	0.945	3.228	8.186	11.414	4.888	9.806	9.601
Municipality FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	10,217	10,217	10,217	10,217	10,217	10,217	5,329	10,179	10,179
R-squared	0.900	0.694	0.504	0.880	0.736	0.681	0.884	0.792	0.887

Notes: Robust standard errors shown in parenthesis. All regressions include probability weights using our household panel. An observation is a household-month in Panel (a) and a municipality-month in Panel (b). The independent variable in Panel (a) is the logarithm of distance to the closest ration store and in Panel (b) is the number of ration stores divided by the number of households with at least one eligible individual in the census (in tens of thousands). The dependent variable for Columns (1) to (3) are categorical variables equal to one if the given household purchase at least one liter from the government, any private brand, or any brand, respectively. The dependent variable for Columns (4) and (5) is the total liters per household purchased from the government, and from any private brand, respectively. The dependent variable for Columns (6) to (8) is the weighted average price per liter paid for the government, any private brand, or any brand, respectively.

**Table OA-4: Staggered Store Introduction Results
(Additional Transaction-Level Price Regressions)**

	(1)	(2)	(3)	(4)
Distance to closest govt. store	0.055 (0.022) **	0.055 (0.022) ***	0.014 (0.003) ***	0.021 (0.032)
Close substitutes				-0.146 (0.041) ***
Distance × Close subs.				0.071 (0.035) **
One std. dev. increase distance ($\sigma \times \beta$)	0.32	0.32	0.32	0.12 (far) / 0.53 (close)
Mean of dependent variable	10.96	10.96	10.96	10.96
Unit of observation	hh-month	transaction	transaction	transaction
Product FE			✓	
Municipality-month FE	✓	✓	✓	✓
Household-type controls	✓	✓	✓	✓
Household census tract controls	✓	✓	✓	✓
Observations	313,470	2,807,957	2,807,954	2,807,957
R-squared	0.417	0.316	0.734	0.316

Notes: Robust standard errors clustered at the household-month level shown in parenthesis. All regressions include the fixed effects used as in Table 2. An observation is a household-month in Column (1), and a transaction for Column (2) onwards. The dependent variable in all regressions is the weighted average price per liter paid for private brands. Column (1) is Columns (8) of Table 2. Column (2) replicates our regression results using transactions, weighting each transaction by the share of the total liters for the given household-month. Column (3) replicates Column (2) but adding product fixed effects. Column (4) includes an interaction with a close substitutes categorical variable, which are those products who share at least 3 (out of 4) product characteristics with the government good.

**Table OA-5: Staggered Store Introduction Results
(Heterogeneity of Government Purchases with Demographics)**

	(1)
distance to closest govt. store × eligible household	-0.050 (0.003) ***
distance to closest govt. store × children in hh × eligible household	0.004 (0.003) *
distance to closest govt. store × socioeconomic status × eligible household	0.050 (0.003) ***
Municipality-month FE	✓
Household-type controls	✓
Household census tract controls	✓
Observations	362,847
R-squared	0.358

Notes: Robust standard errors shown in parenthesis. All regressions include probability weights, and include the same controls as in Table 2. An observation is a household-month. The dependent variable is a categorical variables equal to one if the given household purchase at least one liter from the government. Children in household is a categorical variable equal to one if there is at least one child in the household. Eligible household is a categorical variable equal to one if the household is eligible to purchase government milk. The socioeconomic status is a 0-1 mapping from our categorical variable with three values: 0 (low), 0.5 (medium), 1 (high).

Table OA-6: Instrumental Variable Results

	(1) Price per liter (constant MXN)
Distance to metro area × lagged diesel price	0.0037*** (0.0004)
Product characteristic groups FE	✓
Metro area FE	✓
Producer FE	✓
Calendar month FE	✓
Calendar year FE	✓
Observations (product - markets)	20,834
R-squared	0.83
Cragg-Donald Wald F statistic	88.74

Notes: Robust standard errors shown in parenthesis. This table shows the estimation results for Equation 13. An observation is a product-market. The dependent variable is the average price per liter for a given product market. The instrument is defined by Equation 12. The Cragg-Donald Wald F statistic tests for weak instruments.

Table OA-7: Demand Estimation Results – Product Characteristics

	Ineligible households (1)	No children			With children		
		Socioeconomic status					
		Low (2)	Med (3)	High (4)	Low (5)	Med (6)	High (7)
<i>Excluded group: pasteurized × whole × lactose × plastic bag</i>							
pasteurized × whole × lactose × carton	2.55 (0.09)	2.90 (0.06)	2.49 (0.06)	2.08 (0.07)	2.79 (0.05)	2.39 (0.04)	1.98 (0.05)
pasteurized × whole × lactose × plastic bottle	1.57 (0.08)	1.76 (0.06)	1.59 (0.06)	1.42 (0.07)	1.43 (0.04)	1.26 (0.03)	1.09 (0.05)
pasteurized × semi-skimmed × lactose × carton	-2.43 (0.13)	-0.20 (0.10)	-0.30 (0.07)	-0.41 (0.08)	-1.70 (0.09)	-1.81 (0.05)	-1.91 (0.07)
pasteurized × semi-skimmed × lactose × plastic bottle	0.45 (0.16)	0.63 (0.07)	1.60 (0.05)	2.56 (0.08)	-0.64 (0.08)	0.32 (0.04)	1.28 (0.06)
pasteurized × semi-skimmed × lactose × plastic bag	1.30 (0.11)	0.14 (0.06)	0.15 (0.07)	0.16 (0.10)	0.20 (0.05)	0.21 (0.02)	0.22 (0.05)
pasteurized × semi-skimmed × lactose-free × plastic bottle	1.74 (0.58)	1.66 (0.19)	2.11 (0.15)	2.56 (0.13)	0.73 (0.11)	1.18 (0.10)	1.62 (0.13)
pasteurized × skim-free × lactose × carton	0.36 (0.11)	0.85 (0.08)	0.24 (0.06)	-0.37 (0.06)	1.05 (0.04)	0.44 (0.03)	-0.17 (0.06)
ultra pasteurized × whole × lactose × carton	1.56 (0.09)	1.28 (0.05)	1.18 (0.06)	1.09 (0.07)	1.40 (0.04)	1.30 (0.03)	1.21 (0.04)
ultra pasteurized × whole × lactose-free × carton	1.54 (0.07)	0.78 (0.17)	-0.06 (0.11)	-0.90 (0.12)	1.05 (0.08)	0.21 (0.08)	-0.63 (0.16)
ultra pasteurized × semi-skimmed × lactose × carton	2.66 (0.10)	2.18 (0.07)	2.33 (0.07)	2.47 (0.07)	2.17 (0.06)	2.31 (0.05)	2.46 (0.05)
ultra pasteurized × semi-skimmed × lactose-free × carton	2.75 (0.09)	3.11 (0.08)	3.23 (0.08)	3.35 (0.09)	2.89 (0.07)	3.01 (0.06)	3.14 (0.07)
ultra pasteurized × skim-free × lactose × carton	2.20 (0.09)	2.01 (0.06)	1.84 (0.06)	1.67 (0.07)	2.19 (0.03)	2.02 (0.03)	1.85 (0.04)
ultra pasteurized × skim-free × lactose-free × carton	2.83 (0.11)	2.25 (0.08)	2.72 (0.07)	3.19 (0.11)	1.47 (0.11)	1.94 (0.08)	2.41 (0.08)

Notes: Robust standard errors clustered at the market level computed via bootstrap shown in parenthesis. The table shows the raw coefficients from Equation 4 for the product characteristics parameters. Each row represents a different combination of product characteristics. Combinations of product characteristics that are not shown are not observed in our data.

Table OA-8: Demand Estimation Results – Private Market Quality Fixed Effects (10 Biggest Producers)

	Ineligible households	No children			With children		
		Socioeconomic status					
		Low	Med	High	Low	Med	High
(1)	(2)	(3)	(4)	(5)	(6)	(7)	
<i>Excluded group: Private 1</i>							
Government	-	-1.07 (0.08)	-1.32 (0.09)	-1.57 (0.11)	-2.60 (0.07)	-2.85 (0.07)	-3.10 (0.09)
Private 2	-0.13 (0.04)	-0.55 (0.03)	-0.36 (0.03)	-0.18 (0.04)	-0.19 (0.03)	0.00 (0.01)	0.18 (0.02)
Private 3	-0.89 (0.05)	-0.03 (0.04)	-0.39 (0.03)	-0.74 (0.05)	0.09 (0.03)	-0.26 (0.01)	-0.62 (0.04)
Private 4	-0.47 (0.09)	0.61 (0.06)	0.71 (0.04)	0.80 (0.04)	-0.58 (0.03)	-0.49 (0.02)	-0.39 (0.04)
Private 5	-2.56 (0.08)	-1.46 (0.09)	-1.56 (0.09)	-1.66 (0.12)	-1.97 (0.05)	-2.07 (0.03)	-2.17 (0.06)
Private 6	-0.87 (0.10)	-0.05 (0.05)	-0.29 (0.06)	-0.53 (0.10)	-1.07 (0.06)	-1.31 (0.02)	-1.56 (0.05)
Private 7	-1.19 (0.10)	-1.37 (0.09)	-1.24 (0.09)	-1.10 (0.10)	-1.06 (0.05)	-0.93 (0.02)	-0.80 (0.05)
Private 8	-1.78 (0.06)	-2.17 (0.13)	-2.54 (0.09)	-2.91 (0.12)	-2.69 (0.07)	-3.06 (0.05)	-3.43 (0.13)
Private 9	1.17 (0.08)	1.49 (0.05)	1.29 (0.04)	1.08 (0.06)	0.76 (0.04)	0.56 (0.02)	0.35 (0.04)
Private10	1.74 (0.03)	1.50 (0.06)	1.87 (0.05)	2.25 (0.06)	1.72 (0.04)	2.09 (0.02)	2.47 (0.03)

Notes: Robust standard errors clustered at the market level computed via bootstrap shown in parenthesis. The table shows the raw coefficients from Equation 4 for selected producer fixed effects. Each row represents a different producer. Private j is the j -th biggest firm in terms of the national retail revenues.

Table OA-9: Change in Producer Margins (Constant MXN) Under Counterfactual Arrangements

	Relative Profits (constant MXN)	
Direct provision	-	
Cash transfers	0.19	(0.00)
Universal vouchers	0.32	(0.00)
Vouchers for top 5 among the poor	0.37	(0.00)
Vouchers for poor households	0.31	(0.00)

Notes: This table shows results for selected counterfactuals. We report the simulation details in [Appendix H](#). Producer profits are computed as the weighted average margin (prices minus costs) from all markets. Margins are shown relative to the direct-provision system.

A.2 Additional Figures

Figure OA-1: Example of Milk, Stores and Ration Cards for Government Eligibility

(a) Government Milk



(b) Ration Card



(c) Example of Store

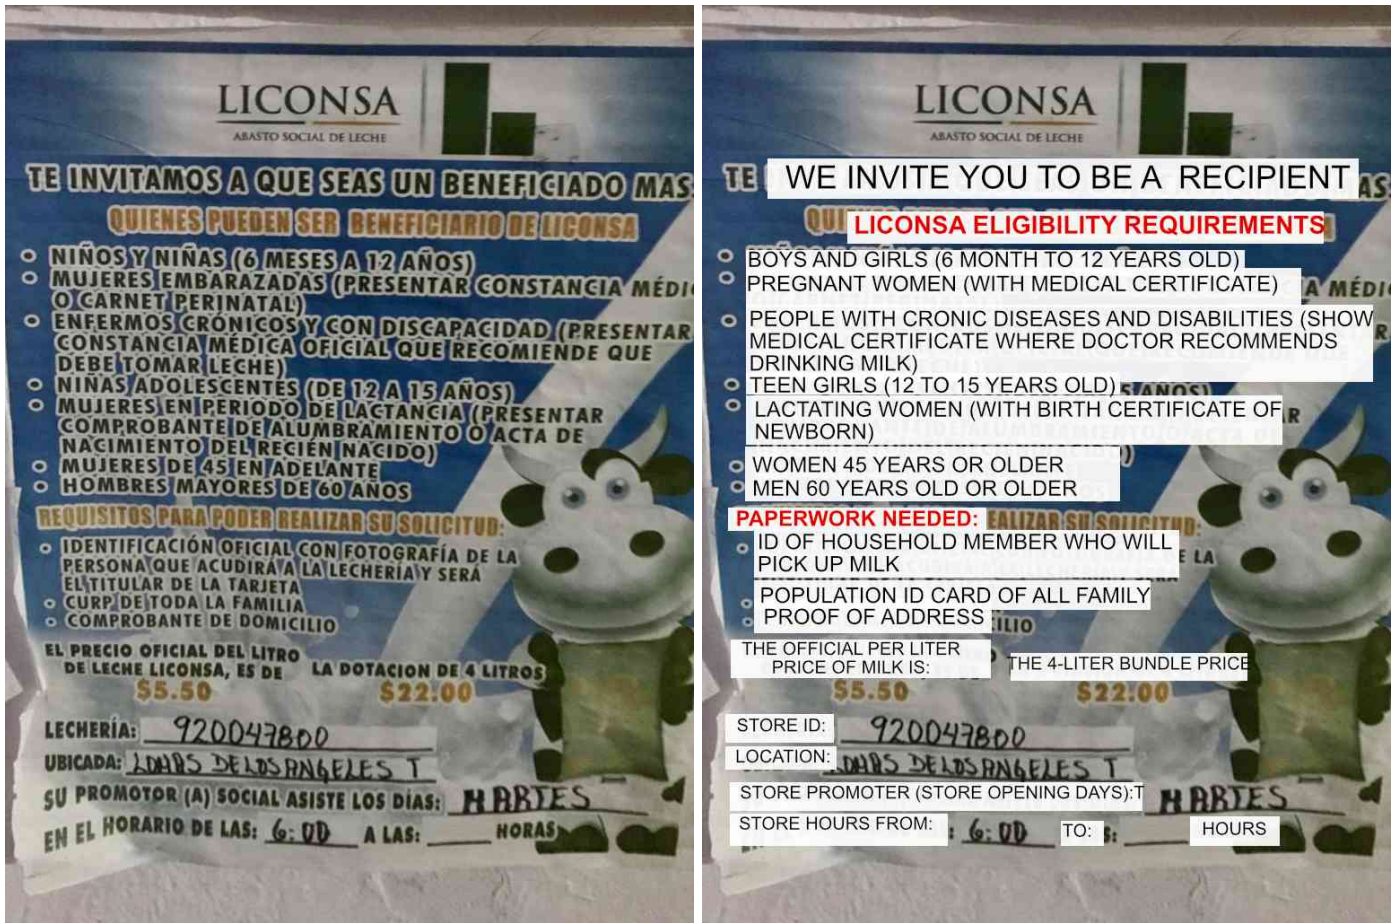


Notes: The panels in this figure show examples of (a) the government milk, (b) a ration card that households get to purchase milk, and (c) a government ration store with the store identification in the right top corner (c). Data identifying individuals in the card are censored for confidentiality purposes.

Figure OA-2: Posted eligibility rules at government ration stores

(a) Original


(b) Translated




Notes: The panels in this figure show examples of the posters the government sets up in all stores in the country. Panel (a) shows the original, which we translate in Panel (b) for convenience.

Figure OA-3: Proportion of Households Who Purchase from Government (January 2010)


(a) By Distance to Closest Ration Store Quintiles



(b) By Number of Children

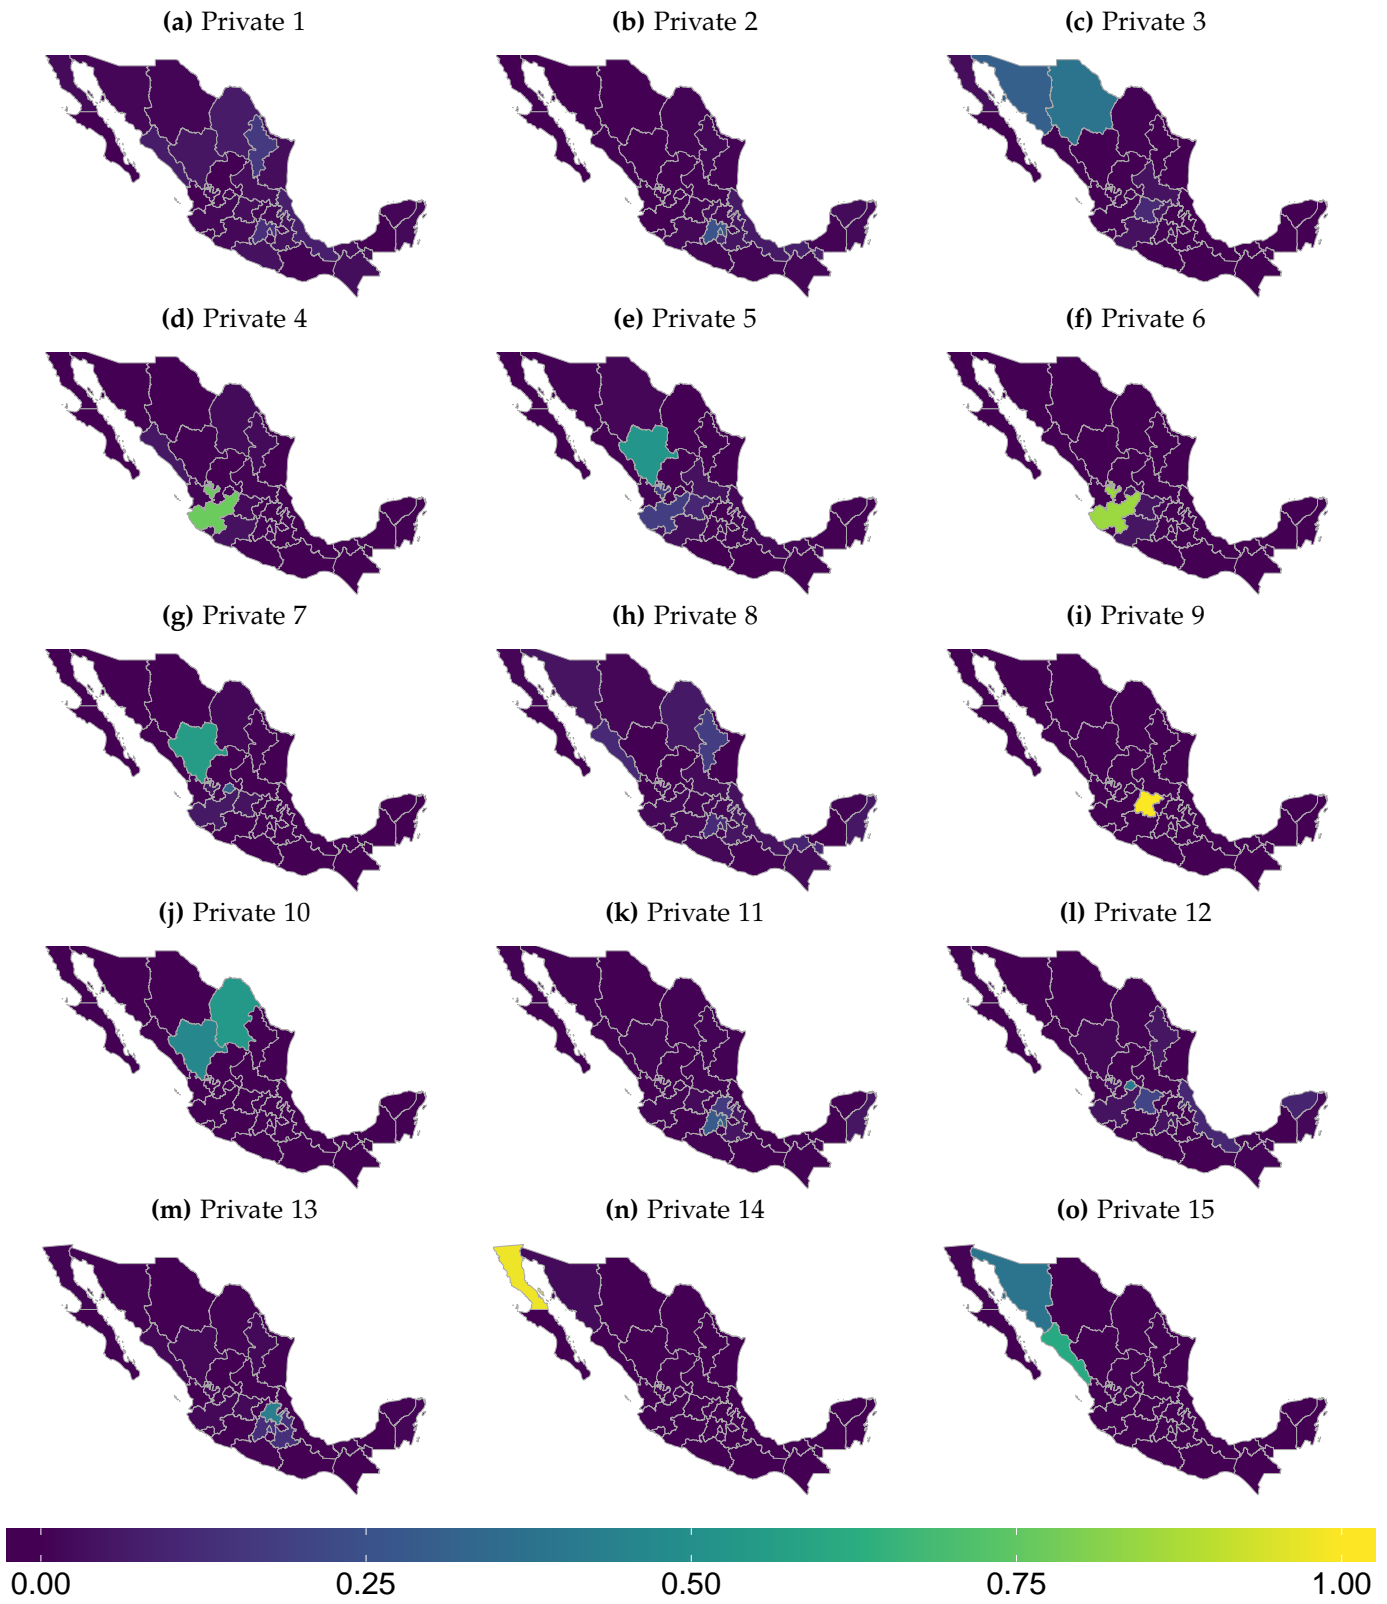


(c) By Socioeconomic Status




Notes: The sample are households in January 2010 in our data (described in Section 3). We plot the proportion of households who purchase at least one liter of government milk with 95 percent confidence intervals. Panel individual below 18 years old. Panel (c) separates household by their socioeconomic status.

Figure OA-4: Share of total retail value by state in Mexico (2010-2014)



Notes: Authors own calculation using household panel data. The total retail value for a given firm is sum (across products and all months in the 2010-2014 period) of the price times the quantity sold by the given company. The revenue share is the proportion of the total revenue that comes from the specific state selected.

Figure OA-5: Concentration worldwide

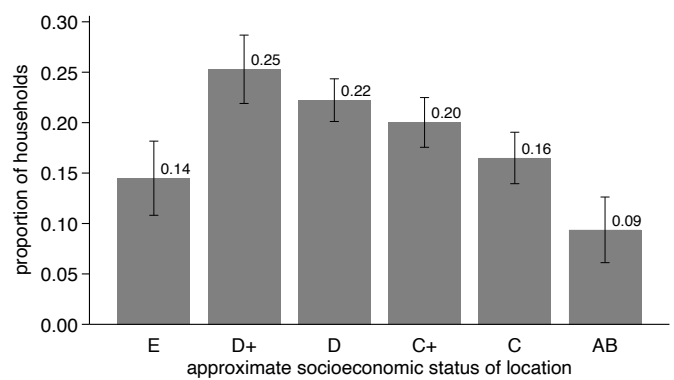



Notes: Authors own calculation with data from Euromonitor International. An observation is a country in 2018. The market retail value is the sum of the retail value of all products sold in a country. The HHI is computed at the parent company per country level.

Figure OA-6: Household consumption of government good by socioeconomic status


(a) National Health Survey, 2012

(b) Household panel, restricted to 2012



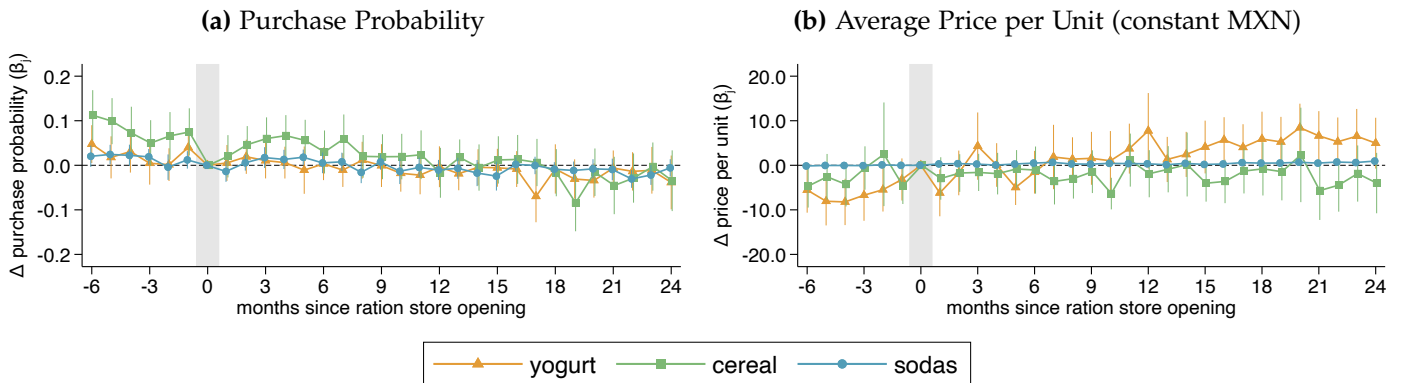
Notes: The data source for the figure on the left is the 2012 National Health and Nutrition Survey (ENSANUT). The data source for the figure on the right is the consumer panel described in Section 3, restricted to 2012. An observation in each source is a household. Panel (a) shows the proportion of households who answered yes to 'Is there currently at least one member of the household who receives Liconsa milk?'. The x-axis is the socioeconomic status of the household. Panel (b) shows the proportion of households that consume Liconsa milk at least twice in the year 2012 in this paper's household panel. The x-axis is constructed using the average socioeconomic index categorization of the municipality of residence of households.

Figure OA-7: Event Studies with Ration Store Openings
 (Additional Results for Low Socioeconomic Status Households)




Notes: These figures plot the estimates of β_j in Equation 1 for different values of j in the x-axis. An observation in this graph is a hexagon-month. Confidence intervals at the 90 percent statistical level are based on robust standard errors twoway clustered at the hexagon and month levels. Each line is a different regression. All regressions include hexagon and month fixed effects.

Figure OA-8: Placebos: Event Studies with Ration Store Openings in Other Markets



Notes: These figures plot the estimates of β_j in Equation 1 for different values of j in the x-axis. An observation in this graph is a hexagon-month. Confidence intervals at the 90 percent statistical level are based on robust standard errors twoway clustered at the hexagon and month levels. Each line is a different regression. All regressions include hexagon and month fixed effects.

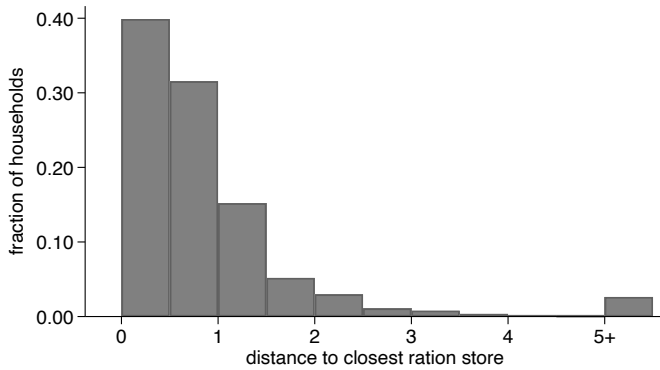
Figure OA-9: Event Studies with Ration Store Openings (Producer and Product Exit)




Notes: These figures plot the estimates of β_j in Equation 1 for different values of j in the x-axis. An observation in this graph is a hexagon-month. Confidence intervals at the 90 percent statistical level are based on robust standard errors twoway clustered at the hexagon and month levels. Each line is a different regression. All regressions include hexagon and month fixed effects. The dependent variable is the count of unique products / producers in the given hexagon month.

Figure OA-10: Variation in the Distance to the Closest Ration Store

(a) Cross-sectional variation (January 2010)




(b) Variation Over Time (Avg. Distance)



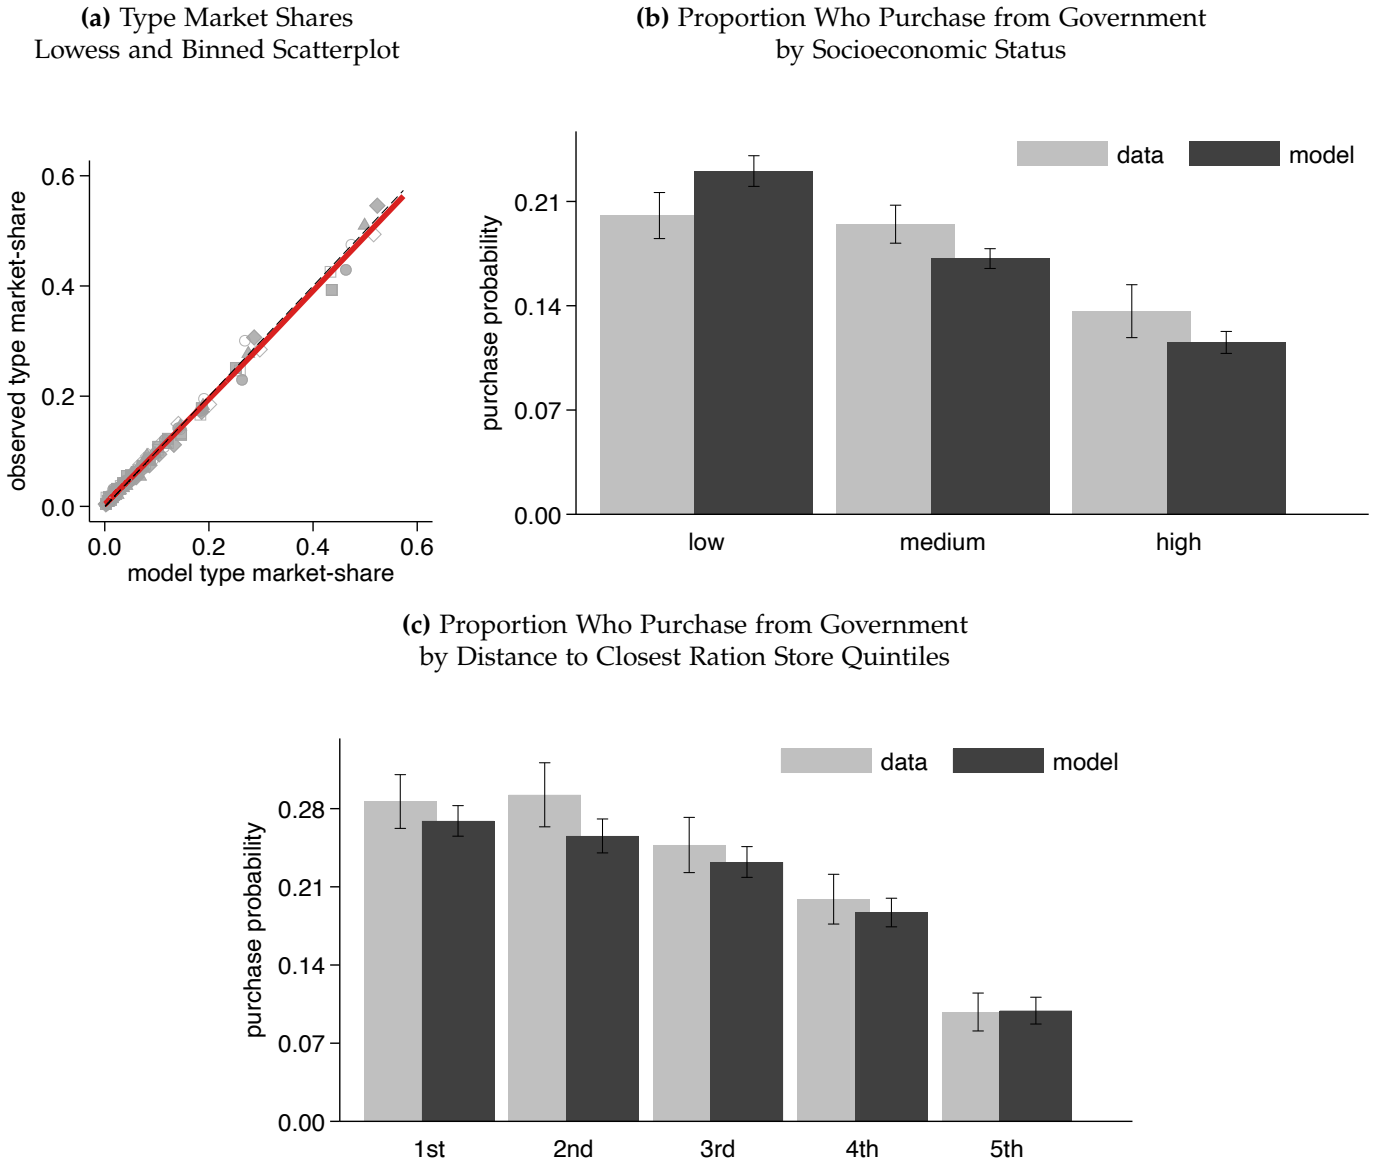
Notes: This figures plots the variation in the distance to the closest ration store in our data. Panel (a) shows the cross-sectional variation in the distance to the closest ration store for the first month of our sample period. We top-code the distance to the ration store to 5 kilometers. Panel (b) shows the average distance to the closest ration store for households that get a change in the distance. We focus on the all households in the rotating household panel, regardless on whether we observe them in all months. Before computing the average distance, we top code the variable to the 95th percentile to avoid outliers to drive the changes in the distance. The panel separates households by whether the distance in December 2014 is lower or higher than the distance in January 2010.

Figure OA-11: Price per liter (constant MXN) across consumer-months (Jan 2012 - Dec 2014)



Notes: Authors' own calculation with household panel data using population weights. An observation is a consumer-month. Each dot represents the average price per liter for the selected firm. Error bars denote the standard error of the estimation.

Figure OA-12: Model Fit



Notes: This figure plots the model fit. Panel (a) shows a binned scatterplot where an observation is a product - market - type. The red line is a lowess regression between the model shares in the x-axis, against the data shares in the y-axis. The dashed line is a 45-degree line. Panels (b) and (c) use as an observation a household month. The light gray bars represent the share of households who purchase from the government in the data. The dark gray bars represent the model implied average probability that a household purchases from the government.

Appendix B Additional Information of the Government Intervention

Procurement of government milk The Mexican government milk supply comes from three sources: own production, local purchases, and imports.⁴⁸ Figure OA-13 shows the collection centers at which the govern-

⁴⁸ There is no official disaggregation of these three sources.

ment buys milk from local producers and the storage facilities at which they store both liquid and powder milk. There are 43 collection centers each with a 988,000 liter per day capacity, and 10 storage facilities locations. In 2014, the government imported 44 thousand tons of powder milk. In 2015, the government imported 35 thousand tons of milk. According to government officials, the the government is shifting its milk input sources to be more local-oriented.

Figure OA-13: Collection, Storage, and Processing Locations in Mexico

(a) Storage and Processing Facilities



(b) Collection Centers



Notes: Source is [Liconsa \(2017\)](#). Panel (a) shows all of the storage and processing facilities used by the government, and the type of milks processed in each of them. Note that they are concentrated in the center of the country. Panel (b) shows all locations where the government buys milk from.

Appendix C Pilot Study for the Government Good's Quality

The survey's main purpose was to elicit individuals' subjective perceptions concerning different milk brands in the market. The sample consisted of 338 survey respondents (all of them students). The survey took place out at ITAM, Mexico City, on five different days (15, 19, 20, 21, 22) of November 2019. It was carried out in Spanish, took between 5 and 10 minutes to complete, and involved 22 survey collectors.

Alternatives Considered We considered three different milk brands: the government milk (leche Conasupo entera), whole milk by Private 1 (LALA Entera), and whole milk by Private 2 (Alpura Clásica Entera). We purchased 4 liters of government milk using one eligible individual. Each of these products is the flagship product from the corresponding firm.

Survey Protocol We implemented the survey as a double-blind tasting study where neither the surveyors nor the survey respondents knew the true identities of milk brands. We divided surveyors into pairs and installed two separate stands at ITAM's main quad, assigning one surveyor to each of the stands as shown in Figure OA-14. In the first stand, the surveyor 1 was in charge of pouring milk into three glasses marked with a letter A, B, or C. To avoid inertia, we randomized the milk that went to each glass at the survey level using a pre-specified list of length 50. In the second stand, survey respondents tasted each of these glasses monitored by surveyor 2. Surveyors were not allowed to mix between stands. Respondents who answered the survey were allowed to choose between a chocolate and an apple as a participation token. To incentivize respondents to pay attention, the final question involved labelling the brands. The three brands (Alpura, LALA, Liconsa) were then disclosed to respondents. Respondents who labelled the 3 glass of milk - brand pairs correctly entered a raffle to win \$1,000 MXN (~\$50 USD). Unfortunately, we could not incentivize willingness to pay.

Survey Questions The survey consisted of 8 questions. For the purpose of the survey, respondents and surveyors referred to the different alternatives as A, B, and C. Questions always followed the same order.

1. In terms of taste, which milk did you like the most? (rank, 1 = most liked, 3 = least)
2. In terms of sweetness, which milk did you like the most? (rank, 1 to 3)
3. In terms of consistency and viscosity, which milk did you like the most? (rank, 1 to 3)
4. In terms of color, which milk did you like the most? (rank, 1 to 3)
5. In nutrition terms, which milk do you think is the most nutritious? (rank, 1 = most nutritious, 3 = least)
6. In terms of fat content, which milk do you think contains more fat? (rank, 1 = least fat, 3 = most)
7. A liter of milk in the supermarket costs areound \$20, though there is variation across brands. What is the maximum amount that you would be willing to pay for each of the products? (positive number)
8. The three brands you tasted are Alpura, LALA, and Liconsa. Which letter do you think is which brand?

Survey Results Figure OA-15 shows under no uncertain terms that survey respondents rank the government milk last against the other alternatives, even if they do not know their identity. In all of our panels, we can reject the null hypothesis that respondents answered choosing one milk at random. In terms of subjective perceptions about flavor, Panels (a) through (e) show that respondents consider government milk to be worse in terms of taste, sweetness, consistency, color, perceived nutritional value, and fat content. The ordering suggests that Private 2's milk is slightly preferred to Private 1's milk — which is consistent with our demand estimates. In terms of prices, Panel (g) shows that respondents are willing to pay less for the government milk. Interestingly, the ordering of prices (not the magnitudes) is consistent with the true ordering of prices in our data. The large difference in magnitudes with respect to our sample prices is expected since we asked about prices in 2019 pesos and all numbers in our data are in 2010 pesos. Finally, Panel (h) shows respondents are more likely to identify the government milk. In fact, households that (blindly) rank the government milk third are more likely to identify it correctly.

Figure OA-14: Example Booths

(a) Serving Booth

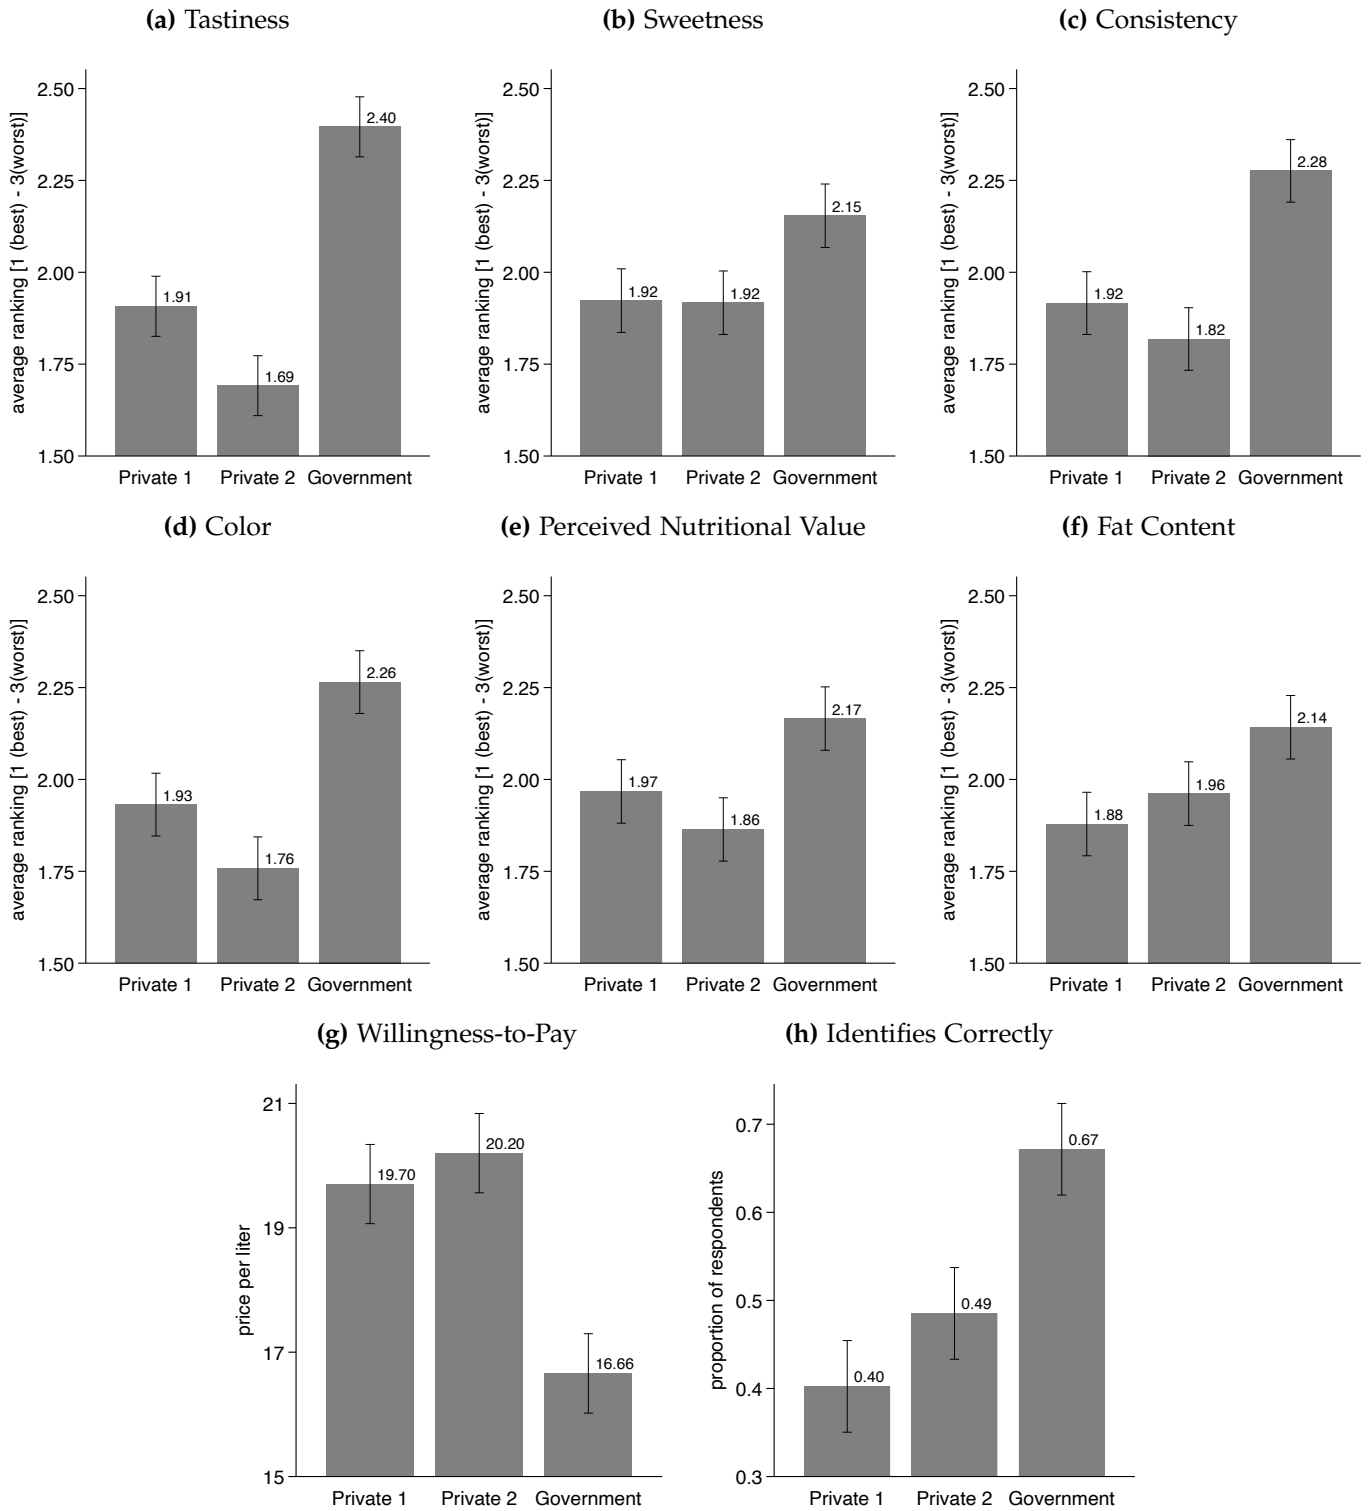
(b) Separator Between Booths

(c) Tasting Booth



Notes: This figure shows the booths used for the blind tasting experiment. Panel (a) shows an example of the serving booth from the point of view of the surveyor. Survey respondents received the questionnaire in the serving booth, and then went to the tasting booth if they decided to participate. After receiving consent, the first surveyor poured milk in glasses marked A, B, and C using the previously specified randomization. The box in Panel (a) and the wall in Panel (b) made it reasonably impossible for anyone to see which glass contained which milk. Panel (c) shows an example of the tasting booth, with a survey respondent in the middle of sipping one of their milks.

Figure OA-15: Pilot Blind Tasting Results



Notes: This figure shows the blind tasting experiment results. Panels (a) through (f) correspond to questions 1 through 6, where survey respondents had to rank (without indifference) the milk in terms of taste. Panel (g) corresponds to question 7 where survey respondents provided a numerical value corresponding to pesos. Panel (h) computes the proportion of survey respondents that identify correctly each brand.

Appendix D Dataset construction

This section provides a detailed explanation of the dataset construction.

D.1 Main Dataset Construction

Product-wise, we restrict our sample on three fronts. First, we restrict attention to one- and two-liter brands of milk. These sizes account for 91% of all unit sales in the sample and 91% of the total liters in the dataset. Second, we drop all transactions in our data where the product purchased had some type of promotion (such as buy 1 get 2), including those where products had a price of zero. These instances represent a relatively small fraction (0.03%) of our household-months. Third, we drop bulk-bought milk. Buying bulk-milk is a relatively uncommon occurrence (4.6% of transactions), and we do not observe the full set of product characteristics for them as we do with the other brands. Table OA-10 summarizes our sample restriction.

Table OA-10: Sample Restriction and Number of Dropped Observations

	Transactions	Households	Household Months	Products	Retail Value	Volume Value
	(1)	(2)	(3)	(4)	(5)	(6)
Initial dataset (all observations)	4,089,176	15,652	405,590	670	100%	100%
Sample restrictions:						
Dropping bulk purchases	187,303	2,718	21,769	1	2.3	3.0
Dropping price zero purchases	19,881	3,632	10,555	432	0.0	0.5
Dropping other milk sizes	336,958	9,123	82,108	232	9.2	9.2
Dropping glass containers	41	6	12	2	0.0	0.0
Intermediate dataset (all geographic areas)	3,652,986	15,552	390,065	559	90.0	89.4
Final dataset (only big metropolitan areas)	2,696,196	11,608	285,617	539	68.3	70.3

Notes: This table shows the number of transactions / households / household-months dropped from the sample due to various restrictions. The starting sample is all liquid milk transactions in our packaged food data. We define products as a combination of a brand name, a pasteurization process, lactose content, package type, and cream content (regardless of size).

Geography-wise, we focus on 36 of the 45 metropolitan areas in our consumer panel sample, dropping the smallest metropolitan areas to ensure we have at least 50 consumers per market. We also drop a few “urban” municipalities that are not part of a specific metropolitan area. Such sample restrictions are useful to estimate market shares for the model in Section 5, and we impose such restrictions in the descriptive evidence in Section 4 to keep the sample fixed, but all results in Section 4 hold without them. The metropolitan areas included in our sample are shown in Table OA-11. These metropolitan areas account for 56 million individuals, or 51% of the total population in Mexico. These regions also show a fair amount of cross-sectional variation in the average price per liter for private brands and on the concentration in the market.

Table OA-11: Metropolitan Areas in Sample

	Population (millions) (1)	Avg. Price per Liter (2)	HHI (0-10,000) (3)	Avg. Municipality Size (km^2) (4)
Valle de México	20.117	11.59	3,959	103
Guadalajara	4.435	9.33	4,441	338
Monterrey	4.106	11.31	8,354	518
Puebla-Tlaxcala	2.729	10.7	3,656	61
Toluca	1.936	11.49	5,153	146
Tijuana	1.751	10.71	4,261	1,494
León	1.61	8.83	2,386	872
Juárez	1.332	10.63	5,414	3,578
La Laguna	1.216	11.49	5,379	1,257
Querétaro	1.097	10.75	3,058	509
San Luis Potosí	1.04	10.06	4,109	885
Mérida	0.973	11.64	3,164	303
Mexicali	0.937	10.7	3,434	15,813
Aguascalientes	0.932	8.95	2,653	601
Cuernavaca	0.925	11.98	5,794	148
Acapulco	0.863	11.63	8,979	1,772
Tampico	0.859	9.94	2,884	1,045
Chihuahua	0.853	11.48	6,281	6,021
Morelia	0.83	10.13	3,568	587
Saltillo	0.823	10.51	6,025	4,625
Villahermosa	0.755	10.9	4,291	1,124
Tuxtla Gutierrez	0.684	11.56	6,574	507
Cancún	0.677	10.87	5,644	1,513
Celaya	0.602	9.36	2,581	387
Pachuca	0.512	11.13	4,266	170
Tlaxcala-Apizaco	0.5	11.15	5,736	37
Cuautla	0.434	12.61	7,982	163
Orizaba	0.427	10.62	6,191	51
Nuevo Laredo	0.384	11.2	8,383	1,216
Minatitlán	0.356	11.73	6,868	488
Coatzacoalcos	0.347	11.13	7,258	165
Colima-Villa de Álvarez	0.334	9.88	2,440	455
La Piedad - Pénjamo	0.25	9.87	1,737	916
Tula	0.206	11.18	4,040	117
Guaymas	0.203	11.34	2,291	4,253
Piedras Negras	0.181	10.67	6,860	689
All sample	56.218	10.89	–	478

Notes: This table shows additional information on the metropolitan areas in our final data. Population (Column 1) is the sum of all municipalities within the metropolitan area, based on the 2010 census. The average price per liter (Column 2) is the total expenditures in milk (across all selected transactions) divided by the total liters of milk. The HHI (Column 3) uses, as market shares, the proportion of expenditures that go to a given firm. The average municipality size (Column 4) is the metropolitan area's total area divided by the number of municipalities.

D.2 Hexagon Dataset Construction


We use Uber’s Hexagonal Hierarchical Spatial Index (also known as h3) to aggregate our household-level data onto hexagons. As we mention in the paper’s main body, the rotating panel nature of the household data requires us to aggregate our information at a coarser level and consider multiple households in a hexagon to capture within location fixed effects. Constrained by Uber’s resolution levels, we use the resolution level 5 (5.31 miles of edge length). It is the smallest resolution level for which we consistently observe multiple households per hexagon. Hexagons of this resolution level have an average hexagon area of 253 squared kilometers — 53% of the size of the average municipality in our data, as shown in Table OA-11. Figure OA-17 shows a comparison between municipalities and hexagons for the eight most populated metropolitan areas.

We observe 289 unique hexagons and 17,340 hexagon-months in our consumer panel data. We focus on the 230 (80%) hexagons and 13,920 (80%) hexagon-months that receive at least one store opening. The average hexagon-month contains 32 households (Figure OA-16(a)), and 19 stores (Figure OA-16(b)). The main advantage is that 68% (56%) of the hexagons appear in 50+ (all 60) months, allowing us to forget about household churn within the data and focus on hexagons as the unit of analysis. To generate event studies relative to a reference month within hexagon h , we define $\tau(h)$ as the date with the largest month-to-month change in the number of stores within the hexagon:

$$\tau(h) = \arg \max_t \{ \#stores_{h,t} - \#stores_{h,t-1} \} \quad (20)$$

and if more than two months satisfy this criterion, we define $\tau(h)$ by selecting one at random. The average increase in month $\tau(h)$ is 3.7 stores. Finally, to compute outcomes at the hexagon-month level, we compute outcomes at the household-month level and then average across households within the hexagon-month.


Figure OA-16: Counts of Observations per Hexagon-Month




Notes: These figures plot histograms of the count of households (Panel a) and government stores (Panel b) within a hexagon-month. For visibility purposes, we restrict histograms at 100 households and 50 government stores, respectively.

**Figure OA-17: Hexagon Examples, Households and Ration Store Location
(Top 8 Most Populated Metropolitan Areas)**


**(a) Valle de México
Municipalities**




**(b) Valle de México
Hexagons**




**(c) Monterrey
Municipalities**




**(d) Monterrey
Hexagons**




**(e) Guadalajara
Municipalities**




**(f) Guadalajara
Hexagons**




**(g) Puebla-Tlaxcala
Municipalities**




**(h) Puebla-Tlaxcala
Hexagons**




**(i) Tijuana
Municipalities**




**(j) Tijuana
Hexagons**




**(k) Mérida
Municipalities**




**(l) Mérida
Hexagons**




**(m) León
Municipalities**




**(n) León
Hexagons**



**(o) Juárez
Municipalities**



**(p) Juárez
Hexagons**



■ Households in consumer panel ● Government stores available in Jan/10

Notes: This figure exemplifies how we divided metropolitan areas into hexagons. Figures (a), (c), (e), (g), (i), (k), (m), and (o) show the standard division of the metropolitan areas onto municipalities. Figures (b), (d), (f), (h), (j), (l), (n), and (p) show the division of metropolitan areas based on Uber's h3 with a resolution level of 5. Turquoise squares denote the location of households within the metro area. Red dots show the location of ration stores that were open by January 2010.

Appendix E The Crusade Against Hunger and Distance to Ration Stores

E.1 Municipality selection criteria

This subsection explains the selection criteria through which the government chose municipalities to open ration stores. We first explain the criteria for the first 400 municipalities selected to participate in the Crusade in 2013. We then introduce the criteria used for the remaining 612 municipalities selected to participate in 2014. After describing the selection process in detail, we discuss some summary statistics at the municipal level that shed light on the differences between the Crusade municipalities and those that did not participate.

E.1.1 2013 Municipality Selection

To select the first 400 municipalities, the federal government used information on (i) the number of individuals that suffered from extreme poverty and (ii) that lacked access to sufficient food for a nutritious diet. The government classified households using the 2010 census as in extreme poverty (or not), and with lack of proper access to food (or not).⁴⁹ The government then used a two-step process to select the 400 municipalities.

For the first step, Table OA-12 shows the criteria used for the first 381 municipalities. Note that the criteria 1 and 2, by construction, focus on the *biggest* municipalities as these ones tend to have the highest count of individuals who live in extreme poverty or without access to food. Similarly, criteria 2 and 4 focus on the *poorest* municipalities. Figure OA-18 paints a similar picture to the one described above, where the government selected the north-most municipalities in the scatter plot (representing the poorest municipalities and those without access to food) both in terms of population count in Panel (a) and as a percent of the total population in Panel (b) for the first round of the Crusade.

Table OA-12: Selection criteria, Crusade Against Hunger (2013 municipalities)

Selection criteria (sorting method, from highest to lowest)	Municipalities selected
1 Population (count of individuals) living in extreme poverty	167
2 Population (as percent of total) living in extreme poverty	184
3 Population (count of individuals) living in extreme poverty and without access to food	140
4 Population (as percent of total) living in extreme poverty and without access to food	150
Total municipalities (eliminating duplicates)	381
Municipalities in metropolitan areas	69

Notes: An observation is a municipality. This table shows the selection criteria for the first 381 municipalities selected by the government to participate in the Crusade in 2013. For each criteria, the federal government sorted municipalities from worst (e.g., those municipalities where the count of poor individual was the highest) to best (e.g., those with the lowest count). Then, the government chose an arbitrary threshold to select the worst X municipalities given this selection criteria.

⁴⁹The first variable is computed as the proportion of households whose income is below a certain threshold. The second variable is computed using a questionnaire similar to the one used by FAO (2019) to measure food insecurity. For more information see the Coneval website (in Spanish) [[link](#)].

For the second step, the remaining 19 (out of the 400) municipalities were selected arbitrarily, with the understanding that small municipalities (with a high share of poor, but not among the poorest) had a lower chance to be selected based on the above-defined criteria. For example, the government selected 1 municipality per state (4 in total) to make sure that all states were represented in the final sample. Then the government selected 9 additional municipalities to ensure that at least 50 percent of the eligible population (i.e., those in extreme poverty) were covered in these first round of the program. The additional 6 municipalities were chosen to participate based on the 6 “worst” municipalities that were not selected before — a completely arbitrary measure with no precise definition of “worst”. For more details, see [SEDESOL \(2013\)](#).

E.1.2 2014 Municipality Selection

In March 2014, the government selected 612 additional municipalities. They selected these additional municipalities with the main goal of reaching at least 75 percent of the total population living in extreme poverty. Similar to the 2013 municipalities, the government implemented a two step process to select these additional municipalities. Table [OA-13](#) shows the criteria used for the first 503 municipalities. Note that the selection criteria is similar to the one shown on Table [OA-12](#), but with more municipalities. To this end, one additional criteria to be selected in 2014 was that the municipality was not part of the first selection round.

Table OA-13: Selection criteria, Crusade Against Hunger (2014 municipalities)

Selection criteria (sorting method, from highest to lowest)	Municipalities selected
1 Population (count of individuals) living in extreme poverty	455
2 Population (as percent of total) living in extreme poverty	455
3 Population (count of individuals) living in extreme poverty and without access to food	455
4 Population (as percent of total) living in extreme poverty and without access to food	455
Total additional municipalities (eliminating duplicates and eliminating those selected in 2013)	503
Municipalities in metropolitan areas	82


Notes: An observation is a municipality. This table shows the selection criteria for the first 503 municipalities selected by the government to participate in the Crusade in 2014. For each criteria, the federal government sorted municipalities from worst (i.e., those municipalities where the count of poor individual was the highest) to best (e.g., those with the lowest count). Then, the government chose an arbitrary threshold to select the worst X municipalities given this selection criteria.

Similar to 2013, the remaining 109 municipalities were selected arbitrarily. The government selected 1 municipality per state (2 in total) to make sure that all states were in the final sample. Additionally, the government selected 45 municipalities in the Northern and Southern border (34 in the Northern border and 11 in the Southern border). The remaining municipalities were selected from the poorest states (Guerrero and Oaxaca), and states heavily affected by the drug-related violence (Michoacan, Sinaloa and Coahuila).

E.1.3 Comparison of Crusade and non-Crusade Metropolitan Area Municipalities

For this subsection, we focus on municipalities within our metropolitan area sample. We compare the 151 selected municipalities to the remaining metro-area municipalities that did not participate in the Crusade.

Figure OA-18: Municipality Selection Criteria (2013 and 2014)



Notes: An observation is a municipality. This figure plots the four variables that the federal government used to select municipalities for the Crusade. Panel (a) shows the count of individuals within a municipality living in extreme poverty (y-axis) and with lack of access to food (x-axis). Panel (b) shows the analogous numbers when divided by the total population within a municipality. Selected municipalities in the first and second rounds are shown in yellow triangles and blue squares, respectively.

One key concern for our results' external validity is that Crusade municipalities may differ substantially from non-crusade ones. If that is the case, our estimated "treatment effects" would be local and the potential "treatment effects" would be different elsewhere. In Table OA-14, we show that it is not necessarily the case in our metropolitan area sample. This table shows that Crusade municipalities are similar to non-crusade ones (before the program was implemented) in terms of different milk market variables. The only variable that is somewhat different is the number of ration stores per 10,000 households, where non-selected municipalities had more stores.

E.2 Within-municipality selection of ration stores location

In this section, we show that the exact location where ration stores opened depended on the proportion of eligible households within the census tract. In conversation with government officials, we were told that one often used way to decide where to open ration stores is by using the 2010 Mexican census data and selecting locations based on the share of the population that would be eligible to purchase from the government within a given location. With this spirit, we try to replicate such strategy using census tracts — the finest-level of disaggregation available. We focus on the within-municipality choice of where to open ration stores.

We first construct our measures of the eligible population and the change in the number of stores at the census tract level. We measure the change in the number of ration stores as the number of stores by December 2014 minus the stock of stores by January 2013. To count the share of eligible individuals within the census

Table OA-14: Crusade municipalities comparison in January 2013 (mean and standard errors)

	Non Selected (1)	Round 1 (2)	Round 2 (3)	p-value (4)
Weighted avg. price per liter (private brands)	12.06 (0.14)	12.12 (0.10)	12.17 (0.12)	0.97
Total liters purchased by month	9.76 (0.81)	10.37 (0.71)	9.63 (0.80)	0.94
Households that purchase from government	0.14 (0.03)	0.19 (0.03)	0.17 (0.03)	0.72
Government stores per 10,000 eligible households	3.90 (0.36)	2.59 (0.16)	2.83 (0.26)	0.01
Average distance to closest ration store (kms)	2354.96 (519.31)	737.97 (51.18)	1511.98 (274.61)	0.10

Notes: An observation is a municipality. Column (1) shows selected means for municipalities not selected for the government program. Columns (2) and (3) show the corresponding means for municipalities selected in the first and second rounds, respectively. Column (4) shows the p-value of the joint null hypothesis that all means (Columns 1 through 3) are the same.

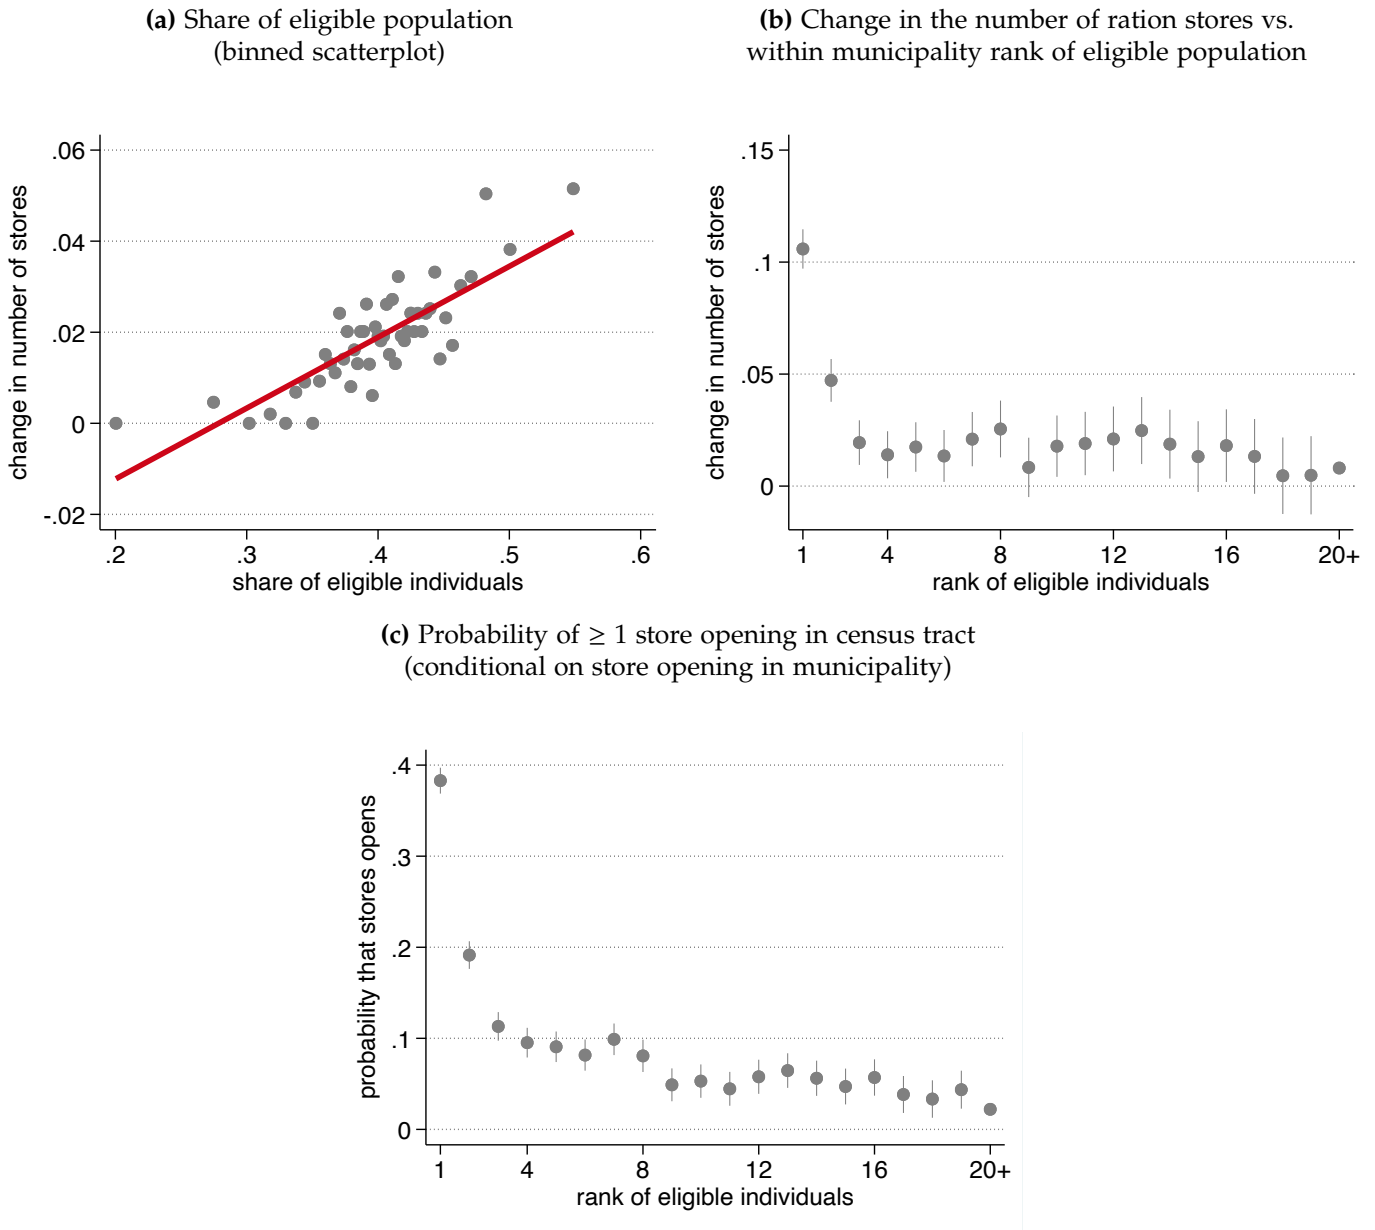
track, we use publicly-available data files and sum the counts of (i) girls between 0 and 14 years old, (ii) women above 50 years old, (iii) boys between 0 and 14 years old, and (iv) men above 65 years old. Next, we divide this sum by the total population within the census tract. This measure is the closest approximation to that of the government using the non-restricted use access data.⁵⁰

Our data show a significant correlation between the change in the number of stores and the share of eligible individuals within the census tract. In Figure OA-19(a), we show there is a positive correlation between these two variables. Next, we ordered all census tracts with respect to the share of households that are eligible (among all census tracts in the municipality). In Figure OA-19(b) we regress the change in the number of ration stores against census tract rank dummies and plot the coefficients of these dummies. The census tract with the highest share of eligible households has an average change of 0.11 ration stores. For comparison, the unconditional store change across census tracts is 0.016 — one order of magnitude smaller. Note that this store change probability rapidly decays, with the second-ranked (and third) census tracts having half (and one-tenth) of a change in the number of stores, relative to the change in the first ranked tract. Figure OA-19(c) repeats the analysis among those municipalities in which a ration store opened, and shows that the probability that the ration store opens in the census tract with the highest number of eligible individuals is twice that of the second census tract.

In Table OA-15, we show that the above correlation persists even after multiple fixed effects and a host of alternative explanatory variables. Column (1) shows the correlation without any additional controls. Column (2) adds municipality fixed effects and shows there is still a persistent correlation between the change in the number of stores and the share of eligible individuals. In Column (3), we add several potential explanatory variables and show the relationship of the share of eligible population and the change in the number of ration stores does not change. These variables include the share of the population who cannot read or write,

⁵⁰We were supposed to get access to the restricted use files of the data in the Summer of 2020. Unfortunately, COVID-19 happened.

Figure OA-19: Change in the number of ration stores versus share of eligible population



Notes: An observation in each panel is a census tract in Crusade municipalities. Panel (a) plots a binned scatterplot of the change in the number of ration stores (measured as the change in the stock of stores between Dec/14 and Jan/13) against the share of eligible households within the tract. Panel (b) ranks all census tracts by their share of eligible individuals within the census tract, and predicts the change in the number of stores using rank dummies. Each dot represents a regression coefficient and confidence intervals are shown using error bars. The right-most dot contains all census tracts ranked 20 or above.

the share of population without middle school or above, the share of households with no access to laundry machines, the share of households without access to running water, and the share of households without a refrigerator.

Table OA-15: Change in the number of ration stores at the census tract level versus census tract demographics

	dependent variable		
	change in number of stores		
	(1)	(2)	(3)
share of eligible population	0.067 (0.008) ***	0.033 (0.009) ***	0.036 (0.016) **
municipality FE		✓	✓
census tract controls			✓
observations	53,538	53,154	49,610
r-squared	0.00	0.07	0.08

Notes: An observation is a census tract. Each column is a separate regression. The dependent variable in each of these regressions is the change in the number of ration stores, measured as the number of stores in December 2014 minus the number of ration stores by January 2013. The second equation (Column 2) controls for fixed effects. The third equation (Column 3) controls for the share of the population who cannot read or write, the share of population without middle school or above, the share of households with no access to laundry machines, the share of households without access to running water, and the share of households without a refrigerator.

Appendix F Stylized Model Description and Intuition

This section introduces a simplified version of the model in Section 5 with the main objective of providing intuition of the main channels through which direct provision can be beneficial. The model is extremely stylized, and reduces the set of parameters to the minimum. In the first section of this appendix, we restate supply and demand in this oversimplified market and show that for a given set of parameters, direct provision can be more beneficial than equivalent cash transfers. We then tweak parameters one by one to show how changes in each of the parameters may flip the result. In summary, whether direct provision or cash transfers generate larger consumer surplus gains, depends on the relative efficiency (i.e., costs) of the government and private producers, the quality of the private and public goods, and the price elasticity of consumers.

F.1 Model Set-Up

Market Primitives Consider a market with two firms (government and private) and two household types. Households are indexed by i and have heterogeneous preferences over the two goods in the market. Household types are indexed by $t(i) \in \{1, 2\}$ (e.g., rich and poor) and a proportion ω_t are of type t . On the supply side, the private (profit-maximizing) firm sells a single representative product, henceforth called the “private good”. The public (not-profit maximizing) firm sells a government brand of milk, which we will call the “government good”. We index firms by $j \in \{f, g\}$, respectively.

Household Demand Utility of household i from choosing good j is given by $u_{ij} = \gamma_{t(i),j} - \alpha_{t(i)} \cdot p_j + \varepsilon_{ij}$ where ε_{ij} is an iid EV1 unobservable component of utility. We normalize the utility of not consuming anything to

zero. With this structure, demand for good j at prices \mathbf{p} is given by

$$\sigma_j(\mathbf{p}) = \sum_t \omega_t \left(\frac{\exp(\gamma_{tj} - \alpha_t \cdot p_j)}{1 + \sum_k \exp(\gamma_{tk} - \alpha_t \cdot p_k)} \right) \quad (21)$$

with weights that sum up to one.

Private Supply The private firm chooses the price for the private good (p_f), taking the government price as given (p_g^*), to maximize profits (e.g., $p_f(p_g^*) = \arg \max(p_f - c_f) \times \sigma_f(p_f, p_g^*)$ where c_f is the marginal cost per unit sold). The first order condition gives the price best response of the private firm with respect to the the price of the government (conditional on the model parameters). The first order condition is

$$\sigma_f(p_f, p_g^*) + (p_f - c_f) \cdot \frac{\partial \sigma_f(p_f, p_g^*)}{\partial p_f} = 0 \quad (22)$$

Government Supply The government has a fixed budget B that they can use to expend on households, either via cash transfers or by selling a government good for which they can choose the quality (γ_{1g} and γ_{2g} in Equation 21) and the price (p_g). For simplicity, we assume that the government is forced to sell to all households, and, if they decide to use cash transfers, to transfer the same amount of money to everyone. The budget constraint in this case for the government is

$$(c_g - p_g) \cdot \sigma_g(\mathbf{p}) + \tau = B \quad (23)$$

where c_g is the government good unit cost ($c_g - p_g$ is the per unit subsidy) and τ is the per household transfer.

Equilibrium An equilibrium is a tuple $\langle \mathbf{p}, \sigma, \tau \rangle$ such that Equations 21, 22 and 23 hold.

Consumer surplus Based on the model described above, we follow the literature and define the consumer surplus in the market as the expected utility divided by the marginal utility of income. Such formula is

$$CS(\mathbf{p}, \tau) = \sum_t \omega_t \cdot CS_t(\mathbf{p}, \tau) = \sum_t \omega_t \cdot \left\{ \frac{1}{\alpha_t} \log \left[1 + \sum_k \exp(\gamma_{tk} - \alpha_t \cdot p_k) \right] + \tau \right\} \quad (24)$$

F.2 Parametrization and Initial Results

To show comparative statistics by moving a single parameter, we first fix all parameters at specific values given by Table OA-16. The parameters are chosen to mimic specific features of our market. First, poorer households are more reactive to prices than richer ones (α). Second, conditional on type, the government good is always seen as inferior than the private good (γ_f compared to γ_g). Poorer households find the

government good to be better targeted at them (γ_g across types). Third, the government good — even with its lower utility — is more expensive to produce than the private good (c_f and c_g).

Table OA-16: Parameter Values for Toy Model

	products		disutility on price	proportion of households
	private	public		
<i>(a) – Demand side parameters</i>	γ_f	γ_g	α	ω
low socioeconomic status households	1	0.5	2.5	0.5
high socioeconomic status households	1	0.25	1.75	0.5
<i>(b) – Supply side parameters</i>	c_f	c_g		
cost per liter	0.5	0.6		

Comparison of government selling at $p_g = 0$ and no government + cash transfers The following table summarizes outcomes for two budget neutral scenarios detailed above. The amount of cash to transfer was selected so that it is the same as the deficit of the government firm.

Table OA-17: Comparison of counterfactuals

	Scenario 1	Scenario 2
	Government provision	Equivalent cash transfers
Price of private product	1.08	1.14
Price of government product	0	∞
Market share of private product	0.11	0.20
Market share of government good	0.53	0
Consumer surplus	0.49	0.44
Consumer surplus rich	0.57	0.50
Consumer surplus poor	0.42	0.38


F.3 Comparative Statics

In all the subsequent figures we plot in the left panel the consumer surplus under both scenarios from Table OA-17, and in the right panel the price of the private firm under the Scenario 1 (government sells goods) and Scenario 2 (monopoly with cash transfers).

Changes in cost of government good If the cost of the government good is higher, then the government needs more money to fund the government provision. Because we chose to keep the price of the government good fixed at 0, then changing the value of λ implicitly changes the amount of money households would


receive in the absence of the government good, which scales consumer surplus linearly. When the government is more inefficient than before to produce the good (e.g., twice the cost of the original parameters), direct provision is dominated by getting the same amount of money in cash.

Figure OA-20: Changes in the cost of government good ($c'_g = \lambda \cdot c_g$)



Changing quality of government good We first consider changes in the quality of the government good that apply to all households (rich and poor) in our data. We even consider negative values for the quality, as it is allowed by the model. Note that when the quality of the government good goes to minus infinity, then it is as if you introduce a good no one would buy, and prices would tend to the monopolist prices. In this extreme world of a very crappy government product, the equivalent cash amount you could transfer to households would be zero dollars. Therefore, we would not have any surplus gains from direct provision.

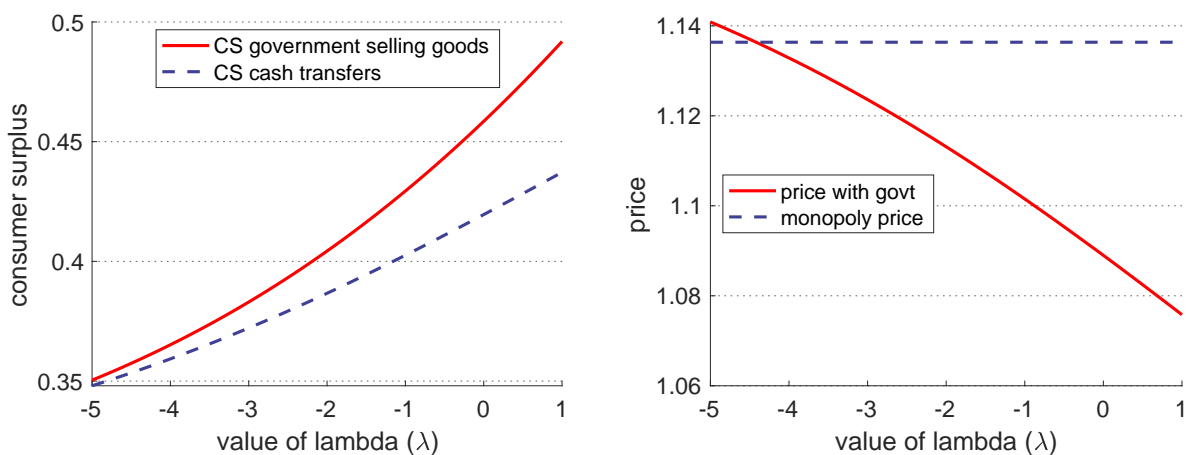
Figure OA-21: Changes in overall quality of government good ($\gamma'_{ig} = \lambda \cdot \gamma_{ig}$ for both $i = \{H, L\}$)



One often considered policy is to better target goods to poor households. We operationalize such policy


by decreasing the perceived quality by high socioeconomic status households while keeping constant the quality perceived by the low socioeconomic status households. Note that decreasing the perceived quality may actually induce higher prices for the private goods (as can be seen in the leftmost values of the right panel). The intuition is straightforward: poorer households would purchase the government good, while the richer ones are constrained to the private market (as the government good is too crappy for them). However, by removing the poor households from the private market you remove the most price elastic consumers, decrease the slope of demand, and therefore firms may increase their price to extract more rents from rich households. In turn, there would be no surplus gains as some households (the rich) may be hurt by these higher prices in the private market.

Figure OA-22: Changes in perceived quality by rich households ($\gamma'_{Hg} = \lambda \cdot \gamma_{Hg}$)



Changes in the price elasticity of households We consider changes to the price sensitivity of households (α_i) that apply both to low and high socioeconomic status households. The takeaway from both graphs is that the more price sensitive the households are (higher price sensitivities are to the right of the graph), the smaller the markups in the monopoly case, and therefore, you get smaller price decreases from the government extra competitor. Therefore, the consumer surplus gain from the added competitor decreasing prices is null — and thus the surplus is smaller than just getting the cash.

Figure OA-23: Changes in price elasticity of all households ($\alpha'_i = \lambda \cdot \alpha_i$ for both $i = \{H, L\}$)



Appendix G Model Estimation Details

In this Appendix we lay out our estimation procedure of the supply and demand model. We discuss each stage in turn and describe the variation in the data that allows us to identify the relevant parameters. We proceed in two steps, first describing demand, and later supply. Unless otherwise specified, we keep the same notation as in Section 5.

G.1 Implementation details

We defined markets as a combination of metropolitan area and a month.

Model in per Liter Terms To focus on the choice among the available alternatives, in the simplest version of our model we collapse all package sizes onto one and denominate products without taking into consideration its size (of either 1 or 2 liters). This assumption helps us with parsimony on two fronts. First, it allows us to reduce the product space considerably. Second, it also helps us to interpret all magnitudes in our model in per liter terms. In previous versions of the paper, we used package size as one of our product characteristics (with the caveat that the instrument is harder to interpret if not in per liter terms) and obtained qualitatively the same results. For this reason, we consider removing product sizes as a very useful simplifying assumption.

Selection of Most-Consumed Milk per Month To focus on their most consumed alternative, we make the simplifying assumption of only allowing households to purchase a single good per month. One concern of this assumption is that households may be complementing their private-milk purchases with government milk. Such behavior seems, if anything, not as likely. For 86% of the household-months with government purchases, the government good is the product with the highest number of liters. Conditional on purchasing from the government in a given month, households source 85% of their total consumption from the government on average.

Proxying Eligibility Unfortunately, we do not directly observe whether a household is eligible or not to buy milk from the government. Nonetheless, we use the observed demographics in our household panel to proxy for eligibility. In particular, we observe the age and gender of the head of household and their partner, and the age of the youngest household member. We use these demographics from our household panel to proxy for eligibility, while acknowledging that we may induce small inclusion and exclusion errors.⁵¹ We define a household as eligible if there is a women above 45 years old; or a man above 60; or the youngest household member is below 15 years old. We also define as eligible those households that we observe purchase government milk in the data, whose are not considered eligible through our initial definition. These households represent one percent of the total households in our data.

Household Types Table OA-18 shows the average household composition in a market. Only 5 percent of households are ineligible to purchase the government milk. Among the eligible, on average, 29, 47, and 19 percent of households fall in the low, medium and high socio economic status, respectively. In terms of number of children in the household, 74 percent of households have at least one child, with poorer households being more likely to have at least one.

Table OA-18: Average Household Composition in a Market

	ineligible	low socio		med socio		high socio	
		no kids	kids	no kids	kids	no kids	kids
percentage of households	5	6	23	10	37	5	14

Notes: This table shows the proportion of households of each type in an average market. To construct the Table, we compute the average (across markets) of each given type of households.

G.2 Algorithm Description

We follow the literature (Berry, 1994; Berry, Levinsohn, and Pakes, 1995) and break down consumer utility into two separate arguments. The decomposition is useful to explain—loosely speaking—which moments in the data are the most informative about which model parameters. Formally, for consumer i , product j , in market m , the arguments are defined as

$$\delta_{jm} = \mathbf{x}_j \cdot \boldsymbol{\beta} - \alpha \cdot p_{jm} + \xi_{jm} \quad (25)$$

$$\mu_{ijm} = \mathbf{x}_j \cdot \tilde{\boldsymbol{\beta}}_{t(i)} - \tilde{\alpha}_{t(i)} \cdot p_{jm} - \tilde{\gamma}_{t(i)} \cdot d_{im} \cdot \mathbb{1}_{(j=g)} \quad (26)$$

⁵¹On exclusion, we miss households whose eligibility arises from pregnancy, disability, or chronic illnesses (they purchase less than 4 percent of the total milk per government records). On inclusion, we incorrectly mark as eligible a small number of households that are in reality ineligible to purchase, but the disaggregation in our data is not sufficient to identify them (e.g., those with boys between 14 and 15 years old).

where tilded parameters make it easier to point out idiosyncratic-type taste shifters.⁵² The first argument (Equation 25) governs the mean valuations common to all consumers in a market and includes the unobservables at the product-market level. The second argument (Equation 26) includes household-specific deviations from the mean valuations. To ease notation, we denote by $\tilde{\theta}$ the set of tilded parameters across all household types.⁵³ With the help of these two arguments, we can write the probability that household i purchases product j in market m

$$\sigma_{ijm}(\delta_m, \tilde{\theta}) = \frac{\exp(\delta_{jm} + \mu_{ijm}(\tilde{\theta}))}{1 + \sum_{k \in \mathcal{J}(t(i), m)} \exp(\delta_{km} + \mu_{ikm}(\tilde{\theta}))} \quad (27)$$

which is a function of δ_m and $\tilde{\theta}$.

Note that if we were to observe values for the vector δ_m , then it would be feasible to estimate α and β using Equation 25 and an instrument for prices. With this intuition, we devise a two-step estimation approach. In the “inner-step”, we approximate the values for δ_m as a function of a specific value for the type taste shifters and the observed aggregate market shares ($\delta_m(\tilde{\theta}, s_m)$), and recover the values of $\alpha(\tilde{\theta}, s_m)$ and $\beta(\tilde{\theta}, s_m)$ using an excluded instrument. In the “outer-step”, we iterate over $\tilde{\theta}$ until we minimize a GMM objective function based on additional moment conditions for the tilded parameters. In Section G.2.1, we focus in the inner step moments. We describe the additional outer-step moments in Section G.2.2. In Section G.2.3, we introduce the full GMM estimation procedure.

G.2.1 Inner step: Estimation of market mean valuations using excluded supply cost shifters

We first focus on estimating the mean valuation parameters (α and β in Equation 25). Conditional on a given value of $\tilde{\theta}$, we define $\delta_m(\tilde{\theta}, s_m) = (\delta_{1m}, \dots, \delta_{jm})$ as the vector of mean valuations that equate the implied model aggregate shares with the observed shares (s_m) in the data.

$$s_m = \sigma_m(\delta_m, \tilde{\theta}) \iff \delta_m(\tilde{\theta}, s_m) = \sigma_m^{-1}(\tilde{\theta}, s_m)$$

Berry (1994) and Berry et al. (1995) show that, conditional on the value of the tilded parameters and the observed market shares, there is a unique value for such vector, $\delta_m(\tilde{\theta}, s_m)$.⁵⁴ Recovering the vector of $\delta_m(\tilde{\theta}, s_m)$, often known as “demand inversion”, allows us to estimate the implied value of $\alpha(\tilde{\theta}, s_m)$ and $\beta(\tilde{\theta}, s_m)$ using an instrumental variables approach

$$\delta_{jm}(\tilde{\theta}, s_m) = x_j \cdot \beta - \alpha \cdot p_{jm} + \xi_{jm} \quad (28)$$

$$p_{jm} = x_j \cdot \pi - \kappa \cdot z_{jm} + \eta_{jm} \quad (29)$$

⁵²With these parameters, we can restate utility as $u_{ijm} = x_j \cdot (\beta + \tilde{\beta}_{t(i)}) - (\alpha + \tilde{\alpha}_{t(i)}) \cdot p_{jm} - \tilde{\gamma}_{t(i)} \cdot d_{im} \cdot \mathbb{1}_{(j=g)} + \xi_{jm} + \varepsilon_{ijm}$.

⁵³These parameters are often known as the non linear parameters in the discrete choice literature.

⁵⁴Formally, define the value of $\delta_m(\tilde{\theta}, s_m)$ as the value that satisfies: $s_m = \sigma_m(\delta_m, \tilde{\theta}) \iff \delta_m(\tilde{\theta}, s_m) = \sigma_m^{-1}(\tilde{\theta}, s_m)$.

with an appropriate instrument, z_{jm} that satisfies the standard IV assumptions of exclusion and relevance.⁵⁵ We introduce our proposed instrument below.

Moment definition for mean market valuations As mentioned above, Equations 28 and 13 jointly determine a set of moment conditions that estimate α and β for a given value of the tilded parameters, $\tilde{\theta}$. To recapitulate, we first fix the type-taste shifters (Equation 26) at a specific value using $\tilde{\theta}$. Then, we recover the vector of mean market valuations $\delta_m(\tilde{\theta}, s_m)$ as the unique values that match the model market shares to the observed ones using the demand inversion procedure. Next, we define the first set of moment conditions as

$$g_{1,jm}(\Theta) = \begin{pmatrix} x'_{jm} \\ z_{jm} \end{pmatrix} \cdot \left(\delta_{jm}(\tilde{\theta}, s_m) - x_{jm} \cdot \beta + \alpha \cdot p_{jm} \right) \quad (30)$$

which, when its sample expectation equals to zero, returns a unique vector (α, β) , conditional on the non-linear parameters, $\tilde{\theta}$. In this equation, we denote by Θ the set of all parameters in our model. Also, we write x_{jm} to emphasize that, in addition to the product characteristics (x_j), we control for metropolitan area, producer, calendar month, and calendar year fixed effects. Note that our estimation procedure requires additional moment conditions to identify $\tilde{\theta}$. We detail such moments below.

G.2.2 Outer step: Household-specific type deviations using micro moments

We augment our estimation procedure using additional micro moments (Berry et al., 2004; Petrin, 2002) derived from the data. Loosely speaking, the main role of these moments allow is to capture the heterogeneity in choices across households, and especially across household-observable types. Such data is of particular importance, as it provides within-market variation of choices across consumers holding fixed the product-market-specific unobservables (Berry and Haile, 2020). We introduce these moments below.

Moment definition for type-level market shares We include the market shares at the type-level as moments that we want our model to match. Adding these moments implies that the model fit implicitly chooses the value of $\tilde{\theta}$ that best-matches the heterogeneity in choices across household types within a market. For households of type t , we define the moment condition as

$$g_{2,jm}^{(t)}(\tilde{\theta}) = \frac{\partial \sigma_{jm}^{(t)}}{\partial \tilde{\theta}} \times \left(\sigma_{jm}^{(t)} \left(\delta_{jm}(\tilde{\theta}, s_m), \mu_{jm}(\tilde{\theta}) \right) - s_{jm}^{(t)} \right) \quad (31)$$

where $s_{jm}^{(t)}$ and $\sigma_{jm}^{(t)}$ are the data and model market shares of product j in market m among households of type t , respectively. Similarly, $\partial \sigma_{jm}^{(t)} / \partial \tilde{\theta}$ is the vector of derivatives of $\sigma_{jm}^{(t)}$ with respect to the tilded parameters. These moments are the equivalent GMM moment conditions to using non-linear least squares over the type

⁵⁵ At the true value of the tilded parameters, the IV approach yields consistent estimates when both the number of households within a market and the number of markets grows. We incorporate a set of additional moment restrictions for $\tilde{\theta}$ in the next subsection.

market shares. Implicitly, the parameters α and β enter this moment condition through δ_m . Note that we have one moment per product-market-type, and thus, useful variation to identify the type coefficients, $\tilde{\theta}$.

Moment definition for distance-to-government store regressions We complement the moments defined above with a set of moments whose main goal—loosely speaking—is to identify the distance coefficients, $\tilde{\gamma}_t$. Through our distance regressions (Equation 2), we showed that, in the data, households that are closer to ration stores are more likely to purchase from the government. In the model, households (of the same type) that are closer to ration stores can have a higher purchase probability, relative to households in the same market that live farther away, through the distance parameter $\tilde{\gamma}_t$. Since the distance to the closest ration store varies across consumers within a market, and within consumers across time, these individual choice probabilities provide useful variation to identify these $\tilde{\gamma}_t$ coefficients.

We implement the analogous regressions to Equation 2 and the corresponding Table 2. Instead of using the data as the dependent variable, we replace the dependent variables with the simulated choice probabilities. We define the additional moment conditions as

$$g_3(\tilde{\theta} | s, x, p, d, z) = \begin{pmatrix} \frac{\partial \lambda_1(\tilde{\theta})}{\partial \theta} \times (\lambda_1(\tilde{\theta}) - (-0.014)) \\ \frac{\partial \lambda_2(\tilde{\theta})}{\partial \theta} \times (\lambda_2(\tilde{\theta}) - (-0.050)) \\ \frac{\partial \lambda_3(\tilde{\theta})}{\partial \theta} \times (\lambda_3(\tilde{\theta}) - (+0.004)) \\ \frac{\partial \lambda_4(\tilde{\theta})}{\partial \theta} \times (\lambda_4(\tilde{\theta}) - (+0.050)) \end{pmatrix} \begin{matrix} \text{distance coefficient, col 1 of Table 2} \\ \text{distance coefficient, col 1 of Table OA-5} \\ \text{distance} \times \text{children, col 1 of Table OA-5} \\ \text{distance} \times \text{socio status, col 1 of Table OA-5} \end{matrix} \quad (32)$$

where we match our model distance coefficients to the ones from the regressions using the data. Specifically, λ_1 represents the OLS distance regression coefficient when using, as dependent variable, the share of households that purchase from the government in the model. We match the regression to the point estimate from Column (1) of Table 2. In a similar fashion, $\lambda_2, \lambda_3, \lambda_4$ use the same dependent variable but different regressors. We match that regression's coefficient estimates to Column (1) of Table OA-5.

G.2.3 Generalized method of moments (GMM) estimation procedure

The estimation of $(\alpha, \beta, \tilde{\theta})$ is done through two-step GMM following [Berry, Levinsohn, and Pakes \(2004\)](#) and the subsequent literature using a nested fixed point algorithm (BLP with micro moments). Note that g_2 and g_3 provide a set of overidentifying restrictions and thus require us to find the optimal weights for our moment conditions. We weight all moments equally in the first step, and compute the optimal weighting matrix in the second step. In the outer-loop, we minimize the GMM objective function based off on the micro-moments (Equations 31 and 32) choosing over the non-linear parameters, $\tilde{\theta}$. In the inner-loop, we invert the market share equations to find the implied linear parameters, α and β , through the mean-market moments (Equation 30). After the demand estimation, we recover costs using our demand estimates, observed market shares, and data on retail prices by inverting the private firms' first order conditions. We provide additional details in [Appendix G](#) and show, through simulations, that the estimation algorithm recovers the

true parameters consistently as the number of markets grows.

G.3 Demand Estimation

We estimate demand using GMM searching over the linear parameters $\tilde{\theta}$. The identification assumption is given by $\mathbb{E}[G_{jm}(\Theta)] = 0$, with $G_{jm}(\Theta) = (g_{1,jm}(\Theta), g_{2,jm}(\Theta)^\top, g_3(\Theta))^\top$ and g_2 is a $T \times 1$ vector, with every type as an element. Given our observed data on choices, prices, product characteristics, and our distance to metropolitan area instrument (s, p, x, z) , the corresponding objective function and estimand is given by

$$\hat{\Theta}(\Omega) = \arg \min_{\Theta} \sum_{jm} G_{jm}(\Theta)^\top \Omega G_{jm}(\Theta) \quad (33)$$

for a given weighting matrix Ω . We follow Hansen (1982) to estimate our parameters through two-step feasible GMM. In the first step, we estimate parameters using the identity matrix as Ω . In the second step, using the optimal weighting matrix based off on the first step. To solve our model, we employ a quasi-Newton algorithm with a tolerance level of 10^{-7} . We supply the gradient to the problem and explain our analytic derivatives below.

We follow the standard inner-outer loop using the nested fixed point algorithm similar to Berry et al. (2004). Note that, because we have as many linear parameters as restrictions in g_1 , we are “just-identified” for the linear parameters. Thus, we can eliminate g_1 from the objective function as the sample mean of those moments is numerically equal to zero. In consequence, we can simplify the problem to the standard inner-outer loop algorithm. In the outer loop we search over the non-linear parameters, $\tilde{\theta}$. In the inner loop we “invert” demand to recover the linear parameters, θ_L . Without loss of generality, the problem can then be simplified to:

$$\hat{\theta}_{NL}(\Omega) = \arg \min_{\tilde{\theta}} \sum_{jm} \left[g_2(\tilde{\theta})^\top \Omega_2 g_2(\tilde{\theta}) \right] + g_3(\tilde{\theta})^\top \Omega_3 g_3(\tilde{\theta}) \quad (34)$$

such that $\delta = \delta(\tilde{\theta})$

for the appropriate weighting matrices Ω_2 and Ω_3 . We then recover the linear parameters through $\mathbb{E}(g_1) = 0$. Implicit in this simplification is that the two moment sources are independent. Our standard errors are computed using the consistent estimates.

Computation of the distance to government ration store regressions The computation of our $g_3(\tilde{\theta})$ moments require further explanation. Recall that our two moments are regression coefficient estimates based on the distance to the closest ration store regressions (Equation 2). These regressions include as controls (i) municipality-month fixed effects, (ii) household demographic controls, and (iii) household census tract controls. To reduce the computational burden in our numerical estimation, we implement the fixed effects using the Frisch-Waugh-Lovell Theorem and restrict the household census tract controls to the proportion of

eligible households. Below we explain how to compute the $\lambda_1(\tilde{\theta})$. A similar approach can be used for the other moments, $\lambda_2(\tilde{\theta})$, $\lambda_3(\tilde{\theta})$, $\lambda_4(\tilde{\theta})$, by simply changing the regressors.

Denote by \mathcal{X} the matrix of controls listed above; by $P_{\mathcal{X}}$ the projection matrix for these controls; by $M_{\mathcal{X}}$ the annihilator matrix for these controls; by $Y_{\text{gov}}(\tilde{\theta})$ the vector of the model probability that household i in month m purchases from the government; and by D the vector of distance to the closest ration store for household i in month m . For a given level of the non-linear parameters, $\tilde{\theta}$, we follow three steps to compute $\lambda_1(\tilde{\theta})$:

1. Regress D on the controls \mathcal{X} , and recover the residuals $\hat{\varepsilon}_D = M_{\mathcal{X}} \cdot D$.
2. Regress $Y_{\text{gov}}(\tilde{\theta})$ on the controls \mathcal{X} , and recover the residuals $\hat{\varepsilon}_Y(\tilde{\theta}) = M_{\mathcal{X}} \cdot Y_{\text{gov}}(\tilde{\theta})$.
3. Recover $\lambda_1(\tilde{\theta})$ as the coefficient of a regression of $\hat{\varepsilon}_Y(\tilde{\theta})$ on $\hat{\varepsilon}_D$.

The advantage of this approach is a three-fold. First, we use an iterative method to reduce the computational burden of Steps 1 and 2, instead of inverting a very sparse matrix for the fixed effects and controls. Second, note that Step 1 can be preallocated, so there is no need to compute it more than once. Third, Step 3 uses a single regressor and can be computed easily.

G.3.1 Analytic derivatives of the GMM function

The computation of our optimal solution heavily relies on the gradient. To show our derivatives with simplified notation, we make the following notation change. The change, shown below, makes it easier to show all derivatives in a simplified way.

$$\mathbf{w}_{ijm} = \begin{pmatrix} \mathbf{x}_{jm} & -p_{jm} & d_{ijm} \end{pmatrix} \quad \boldsymbol{\varphi}_{t(i)} = \begin{pmatrix} \tilde{\boldsymbol{\beta}}_{t(i)}^{\top} & \tilde{\alpha}_{t(i)} & \gamma_{t(i)} \end{pmatrix}^{\top}$$

With this notation change, we denote (suppressing the function arguments) the choice probability for good j from household i of type $t(i)$ in market m as:

$$\sigma_{ijm} = \frac{\exp(\delta_{jm} + \mathbf{w}_{ijm} \cdot \boldsymbol{\varphi}_{t(i)})}{1 + \sum_{k \in \mathcal{J}(t(i), m)} \exp(\delta_{km} + \mathbf{w}_{ikm} \cdot \boldsymbol{\varphi}_{t(i)})} \quad (35)$$

and rewrite the model type-level shares $(\sigma_{jm}^{(t)})$, and the model aggregate market shares (σ_{jm}) with such notation change accordingly.

Equation 35 is useful to show the derivatives of the choice probability with respect to the non-linear parameters. In particular, the derivative of the choice probability with respect to the l -th non linear parameter is given by:

$$\frac{d\sigma_{ijm}}{d\varphi_{t(i)}^l} = \frac{\partial\sigma_{ijm}}{\partial\varphi_{t(i)}^l} + \sum_{k \in \mathcal{J}(t(i), m)} \frac{\partial\sigma_{ijm}}{\partial\delta_{km}} \times \frac{d\delta_{km}}{d\varphi_{t(i)}^l} \quad (36)$$

where the last term makes explicit that changing any of the non-linear parameters changes the demand

inversion, and thus δ_m . The first two elements of Equation 36 are:

$$\frac{\partial \sigma_{ijm}}{\partial \varphi_{t(i)}^l} = \sigma_{ijm} \left(w_{ijm}^l - \sum_{k \in \mathcal{J}(t(i), m)} \sigma_{ikm} \cdot w_{ikm}^l \right) \quad \frac{\partial \sigma_{ijm}}{\partial \delta_{km}} = \sigma_{ijm} (\mathbb{1}_{(j=k)} - \sigma_{ikm})$$

We use the implicit function theorem to compute the last element of Equation 36. Note that, conditional on the value for the non linear parameters,

$$\sigma_{jm}(\delta_m; \tilde{\theta}) - s_{jm} = 0 \text{ for all } j = 1, \dots, \mathcal{J}(m)$$

is a system with $\mathcal{J}(m)$ equations and $\mathcal{J}(m)$ unknowns. Then, by the implicit function theorem:

$$\frac{d\delta_m}{d\varphi_{t(i)}^l} = -D(\sigma_m)^{-1} \cdot \frac{\partial \sigma_m}{\partial \varphi_{t(i)}^l} \quad (37)$$

is the total derivative of the linear parameters with respect to the non linear ones of a given type. In Equation 37, both $d\delta_m/d\varphi_{t(i)}^l$ and $\partial \sigma_m / \partial \varphi_{t(i)}^l$ are $\mathcal{J}(m) \times 1$ vectors, and $D(\sigma_m)$ is a $\mathcal{J}(m) \times \mathcal{J}(m)$ matrix of derivatives. The (j, k) element of the $D(\sigma_m)$ matrix is given by $\partial \sigma_{jm} / \partial \delta_{km}$. We use these equations to generate our gradient.

Gradient for type-level market shares moments Based on our notation change, the non-linear parameters are the collection of the φ parameters, $\tilde{\theta} = (\varphi_t)_t$. Thus, we can explain the derivatives of our moments using the same notation. The derivative of $g_2^{(t)}$ with respect to the l -th parameter in φ_t is given by:

$$\frac{dg_{2,jm}^{(t)}(\tilde{\theta})}{d\varphi_t^l} = \frac{d\sigma_{jm}^{(t)}}{d\varphi_t^l} = \int_{i \in \mathcal{J}(t, m)} \left(\frac{d\sigma_{ijm}}{d\varphi_{t(i)}^l} \right) \quad (38)$$

with the inner element of the integral given by Equation 36 above. Somewhat obvious is that, if the type (subscript t) of the moment condition (Equation 31) and the parameter φ_t do not coincide, then the derivative is equal to zero.

Gradient for distance-to-government store regressions The advantage of using the annihilator matrix M_X when estimating these moments is that we can differentiate every element of the vector $Y_{\text{gov}}(\tilde{\theta})$ with respect to a given parameter φ_t^l and propagate the derivative by premultiplying it by M_X to compute the gradient.⁵⁶ To compute the gradient, we slightly modify the pseudo-algorithm steps described to compute the distance to ration store regressions:

⁵⁶Loosely speaking, the coefficients of a regression of Y on X are just weighted averages of Y (and thus linear in the elements of Y) where the weights depend on the values of X . If Y is a function of an ancillary parameter θ , then the derivative of the regression coefficients with respect to θ can be obtained by the element-wise derivative of Y with respect to θ premultiplied by the projection matrix.

1. Regress D on the controls \mathcal{X} , and recover the residuals $\widehat{\varepsilon}_d = M_{\mathcal{X}} \cdot D$.
2. Perform the element-wise derivative of $Y_{\text{gov}}(\widetilde{\theta})$ with respect to φ_t^l . Denote that as $dY_{\text{gov}}(\widetilde{\theta})/d\varphi_t^l$.
3. Regress $dY_{\text{gov}}(\widetilde{\theta})/d\varphi_t^l$ on the controls \mathcal{X} , and recover the residuals $\widehat{\varepsilon}_{dY}(\widetilde{\theta}) = M_{\mathcal{X}} \cdot dY_{\text{gov}}(\widetilde{\theta})/d\varphi_t^l$.
4. Recover $d\lambda_1(\widetilde{\theta})/d\varphi_t^l$ as the coefficient of a regression of $\widehat{\varepsilon}_{dY}(\widetilde{\theta})$ on $\widehat{\varepsilon}_D$.

Note that Step 1 is still the same, so we do not need to recompute it—just as we did for the moment estimation. With all of this information, the derivative of g_3 with respect to the l -th parameter in φ_t is given by:

$$\frac{dg_3(\widetilde{\theta})}{d\varphi_t^l} = \begin{pmatrix} \frac{d\lambda_1(\widetilde{\theta})}{d\varphi_t^l} & \frac{d\lambda_2(\widetilde{\theta})}{d\varphi_t^l} & \frac{d\lambda_3(\widetilde{\theta})}{d\varphi_t^l} & \frac{d\lambda_4(\widetilde{\theta})}{d\varphi_t^l} \end{pmatrix}^{\top} \quad (39)$$

where, for completeness, an arbitrary element of $dY_{\text{gov}}(\widetilde{\theta})/d\varphi_t^l$ is the derivative of the model-computed government choice probability (e.g., $d\sigma_{igm}/d\varphi_{t(i)}^l$ from Equation 36 where $j = g$ denotes the government good).

Hessian of moment parameters to compute standard errors In addition to the first derivatives, the second order derivatives are necessary in order to compute the standard errors. We differentiate Equation 36 once again (but now with respect to the n -th parameter) to show that:

$$\frac{d^2\sigma_{ijm}}{d\varphi_{t(i)}^l d\varphi_{t(i)}^n} = \frac{d}{d\varphi_{t(i)}^n} \left(\frac{\partial\sigma_{ijm}}{\partial\varphi_{t(i)}^l} \right) + \sum_{k \in \mathcal{J}(t(i), m)} \left\{ \left[\frac{d}{d\varphi_{t(i)}^n} \left(\frac{\partial\sigma_{ijm}}{\partial\delta_{km}} \right) \right] \cdot \frac{d\delta_{km}}{d\varphi_{t(i)}^l} + \frac{\partial\sigma_{ijm}}{\partial\delta_{km}} \cdot \left[\frac{d}{d\varphi_{t(i)}^n} \left(\frac{d\delta_{km}}{d\varphi_{t(i)}^l} \right) \right] \right\} \quad (40)$$

where the new second derivatives are given by

$$\begin{aligned} \frac{d}{d\varphi_{t(i)}^n} \left(\frac{\partial\sigma_{ijm}}{\partial\varphi_{t(i)}^l} \right) &= \frac{d\sigma_{ijm}}{d\varphi_{t(i)}^n} \left(w_{ijm}^l - \sum_{k \in \mathcal{J}(t(i), m)} \sigma_{ikm} \cdot w_{ikm}^l \right) - \sigma_{ijm} \left(\sum_{k \in \mathcal{J}(t(i), m)} \frac{d\sigma_{ikm}}{d\varphi_{t(i)}^n} \cdot w_{ikm}^l \right) \\ \frac{d}{d\varphi_{t(i)}^n} \left(\frac{\partial\sigma_{ijm}}{\partial\delta_{km}} \right) &= \frac{d\sigma_{ijm}}{d\varphi_{t(i)}^n} (\mathbb{1}_{(j=k)} - \sigma_{ikm}) - \sigma_{ijm} \left(\frac{d\sigma_{ikm}}{d\varphi_{t(i)}^n} \right) \end{aligned}$$

and for the second order derivative of δ_m , we employ matrix differentiation that gives us

$$\frac{d}{d\varphi_{t(i)}^n} \left(\frac{d\delta_m}{d\varphi_{t(i)}^l} \right) = D(\sigma_m)^{-1} \cdot \left(\frac{\partial D(\sigma_m)}{\partial\varphi_{t(i)}^n} \right) \cdot D(\sigma_m)^{-1} \cdot \frac{\partial\sigma_m}{\partial\varphi_{t(i)}^l} - D(\sigma_m)^{-1} \cdot \left[\frac{d}{d\varphi_{t(i)}^n} \left(\frac{\partial\sigma_m}{\partial\varphi_{t(i)}^l} \right) \right]$$

where $\partial D(\sigma_m)/\partial\varphi_{t(i)}^n$ is the element-wise derivative of the $\mathcal{J}(m) \times \mathcal{J}(m)$ matrix $D(\sigma_m)$ with respect to $\varphi_{t(i)}^n$ and $d/d\varphi_{t(i)}^n(\partial\sigma_m/\partial\varphi_{t(i)}^l)$ is the $\mathcal{J}(m) \times 1$ vector of element-wise second derivatives of σ_m .

G.4 Supply Estimation

Denote by $\bar{\sigma}_m, \bar{p}_m, \bar{c}_m$ as the vector of market shares, prices and costs for all products in market m *excluding* the government product and the outside option. Denote by Λ_m the matrix where the (j, k) element is equal to zero if products j and k are not produced by the same firm, and equal to the derivative of the share demanded from product k with respect to price j otherwise. We can stack all of the elements of Equation 9 onto a single vector:

$$\bar{\sigma}_m(\Theta) + \Lambda_m(\Theta) \cdot (\bar{p}_m - \bar{c}_m) = \mathbf{0} \implies \bar{c}_m = \bar{p}_m + \Lambda_m^{-1}(\Theta) \cdot \bar{\sigma}_m(\Theta) \quad (41)$$

and recover the cost estimates for each product j in each market m using this equation.

We recover supply conditional on our demand estimation. Denote by $\widehat{\Theta}$ the demand model parameters. Given the estimates of consumer demand, data on the retail prices, and observed market shares, we recover the marginal cost estimates for privately-supplied products in market m using Equation 41 with Λ_m being the element-wise multiplication of an ownership matrix times the matrix of demand-price derivatives. The (j, k) -th element of matrix Λ_m is equal to $\partial\sigma_{jm}/\partial p_{km}$ if products j and k are produced by the same firm, and zero otherwise.

G.5 Influence of Moment Parameters

To study the influence of specific moments, we repeat our estimation exercise while perturbing their point estimate. We do so over a random sample of 50 percent of the markets to ease the computational cost of this exercise. We compute one-step GMM estimates starting from the paper's main estimation results and using the two-step GMM matrix to weight across moments. We start by re-estimating the model in the 50 percent random sample. Then, we modify each of the selected moments to produce the new estimates. Table OA-19 shows our results in parameter differences. For example, the price coefficient from Panel (a), Column (1), means that increasing the share of households that purchase the government good by 10% increases the government producer fixed effect parameter among poor households without kids by 0.081.

Panel (a) shows that increasing the share of households that purchase the government milk translates in the model into a lower value of the omitted product category (the product by Private 1) and into a higher government producer fixed effect. Similarly, product fixed effects from other privately-offered products decrease their point estimates. This is equivalent to saying that households substitute from the private market onto the government good. Panels (b) and (c) show that when we concentrate the increase in the government market share to a specific set of households (e.g., the poor) then the increase in the government producer fixed effect and the decrease in the constant concentrates among the same set of households. These increases the government market share induce smaller (to no-)parameter changes among richer households. This is mostly because a 10% (not percentage points) increase is smaller for richer household.

Similarly, Panel (d) shows that increasing the set of households that purchase milk increases the constant term (e.g., the value of Private 1 milk). Panels (e) and (f) repeat our exercise among the poor and rich, and

shows (similar to Panels (c) and (d)) that the moments most affected are those of households whose parameter is shocked.

Table OA-19: Influence of Moments in Parameter Estimates

	Ineligible households (1)	No children			With children		
		Socioeconomic status					
		Low (2)	Med (3)	High (4)	Low (5)	Med (6)	High (7)
<i>(a) - Increase the government market share by 10%</i>							
Price (α)	0.003	0.005	0.005	0.004	0.003	0.002	0.002
Distance to ration store (γ)	0.000	0.003	0.001	-0.000	0.004	0.003	0.001
Government producer FE	-	0.081	0.070	0.058	0.005	-0.006	-0.018
Constant	-0.031	-0.094	-0.089	-0.084	-0.016	-0.011	-0.006
<i>(b) - Increase the government market share among the poor by 10%</i>							
Price (α)	0.005	0.007	0.006	0.004	0.005	0.004	0.002
Distance to ration store (γ)	0.000	0.005	0.002	0.000	0.004	0.001	-0.001
Government producer FE	-	0.100	0.055	0.009	0.043	-0.002	-0.048
Constant	-0.051	-0.115	-0.086	-0.057	-0.057	-0.029	0.000
<i>(c) - Increase the government market share among the rich by 10%</i>							
Price (α)	-0.001	-0.001	-0.000	0.000	-0.001	-0.001	-0.000
Distance to ration store (γ)	0.000	0.004	0.002	-0.000	0.005	0.003	0.001
Government producer FE	-	-0.001	0.021	0.044	-0.028	-0.005	0.017
Constant	0.007	0.000	-0.018	-0.036	0.027	0.009	-0.009
<i>(d) - Increase the set of households that purchase milk by 10%</i>							
Price (α)	-0.002	-0.000	-0.001	-0.002	-0.001	-0.002	-0.002
Distance to ration store (γ)	0.000	-0.001	-0.000	0.000	-0.001	-0.000	0.000
Government producer FE	-	-0.009	-0.018	-0.026	0.003	-0.005	-0.013
Constant	0.079	0.030	0.079	0.127	-0.050	-0.001	0.048
<i>(e) - Increase the set of poor households that purchase milk by 10%</i>							
Price (α)	-0.002	0.000	-0.000	-0.001	-0.001	-0.002	-0.003
Distance to ration store (γ)	0.000	-0.000	-0.000	0.000	-0.001	-0.000	-0.000
Government producer FE	-	-0.000	-0.003	-0.005	-0.008	-0.011	-0.013
Constant	-0.015	0.044	0.011	-0.021	0.050	0.018	-0.015
<i>(f) - Increase the set of rich households that purchase milk by 10%</i>							
Price (α)	-0.002	-0.003	-0.003	-0.003	-0.002	-0.002	-0.002
Distance to ration store (γ)	0.000	-0.000	-0.000	-0.000	0.000	0.000	0.000
Government producer FE	-	-0.010	-0.017	-0.023	-0.003	-0.009	-0.015
Constant	0.007	-0.006	0.066	0.137	-0.050	0.021	0.093

Appendix H Counterfactual Estimation Procedure and Additional Results

H.1 Estimation Procedure

To estimate the equilibrium outcomes under alternative arrangements, we find a fixed point for demand (Equation 7) and supply (Equation 9) with the altered market structure that satisfies the government budget constraint

$$(c_g - p_g) \cdot \sigma_g(\mathbf{p}) + \tau = B \quad (42)$$

where τ is the per household net transfer (either through cash or vouchers) and B is the market budget.

Direct cash transfers To compute the equilibrium outcomes that would arise absent direct provision, we iterate market by market start from the market shares and prices observed in the data. Denote by $\sigma_m^{(0)}$ and $\mathbf{p}_m^{(0)}$ the initial values of the market shares and prices of all private goods (excluding the government) in a given market. For the n -th step, we start by updating prices using the first order condition for prices (Equation 9 and Equation 41) holding market shares fixed and our predicted cost estimates by

$$\mathbf{p}_m^{(n)} = (1 - \rho) \cdot \mathbf{p}_m^{(n-1)} + \rho \cdot \mathbf{p}_m^* \left(\sigma_m^{(n)} \right) \quad (43)$$

where ρ is the size of our tâtonnement step and $\mathbf{p}_m^* \left(s_m^{(n)} \right)$ is the vector of prices based on the first order conditions. We then update the market shares using our functional form of demand (Equation 7) holding prices fixed by

$$\sigma_m^{(n)} = (1 - \rho) \cdot \sigma_m^{(n-1)} + \rho \cdot \sigma_m^* \left(\mathbf{p}_m^{(n)} \right) \quad (44)$$

with $\sigma_m^* \left(\mathbf{p}_m^{(n)} \right)$ is the vector of market shares with the updated prices in the n -th step. We continue updating until both σ_m and \mathbf{p}_m satisfy a tolerance criteria that ensures we are in a fixed point (tol = 0.001). With a sufficiently small step, this process converges to a new equilibrium of market shares and prices.

Finally, to transfer the money back to households we implicit per household transfer for a given year t by

$$\tau_t = \left(\frac{1}{N_t} \right) \cdot \sum_{m \in \mathcal{Y}(m,t)} M_m \cdot (c_{gm} - p_{gm}) \cdot \sigma_{gm}$$

where N_t is the number of eligible households in the given year, $\mathcal{Y}(m,t)$ is the set of markets (recall that markets are combinations of metropolitan area and months) that correspond to the given year, M_m is the number of households in a given market, and σ_{gm} is the market share of the government good in the market. Because utility is quasilinear in income, changes in income do not generate changes in demand and thus we can first compute the equilibrium that would arise absent the government intervention, and then transfer the money back to households.

Vouchers To formalize the notation, let v_{ijm} be the voucher to household i in market m to be applied to product j , by $\mathbf{v}_{im} = (v_{i1m}, \dots, v_{ijm})$ the voucher structure for a given household and by $\mathbf{v}_m = (v_{1m}, \dots, v_{jm})$ the collection of all vouchers in a market. We use the above notation to accommodate all possible voucher structures, but in the paper focus in simpler designs (e.g., universal vouchers to all goods, or to a subset of products). Vouchers shift individual demand (Equation 5) for product j by $\mathbf{p}_m - \mathbf{v}_{im}$, and thus shift aggregate demand for product j by

$$\sigma_{jm}(\mathbf{p}_m, \mathbf{v}_m) = \mathbb{E}_i [\sigma_{ijm}(\mathbf{p}_m - \mathbf{v}_{im})] \quad (45)$$

We repeat our iterative process (Equations 43 and 44) using as starting values the cash transfers allocation. Using vouchers, some households would decide to select out of the market, which means that those consumers do not reap the benefits of the voucher. Because of the existence of such individuals, we need to compute one additional step to make sure we obtain a balanced budget. We consider the budget to be balanced if the difference between the total budget transferred to households and the government budget is less than $\epsilon 1$. The government expenditures in this program are given by:

$$B_t = \sum_{m \in \mathcal{Y}(m,t)} M_m \sum_{(i,j)} (p_{jm} - v_{ijm}) \cdot \sigma_{ijm} \quad (46)$$

where the interior sum denotes all individuals and all goods in the market.

Price Caps We consider changes in the maximum price per liter (in constant MXN) at which firms sell their products. Denote by κ the maximum allowed price. To implement caps, we proceed market by market in three steps. In the first step, we restrict the alternatives available to the set of products whose marginal cost per liter is lower than κ , that is

$$A_m(\kappa) = \{j \in \mathcal{J}(m) \mid c_{jm} \leq \kappa\} \quad (47)$$

whereas if $c_{jm} > \kappa$, we assume that those products are no longer offered. In the second step, we find the new “equilibrium” prices for the remaining products in the market. To do so, we find the set of prices $\mathbf{p}'_m(\kappa)$ that satisfy the (unconstrained) firms’ first order conditions (Equation 9) when the set of available products is $A_m(\kappa)$. To compute these prices, we repeat our iterative process (Equations 43 and 44) using as starting values the cash transfers allocation. Having computed these prices, we then define the capped prices as

$$\mathbf{p}^*_m(\kappa) = \min\{\mathbf{p}'_m(\kappa), \kappa\} \quad (48)$$

so that when firms would rather price above the cap, then we set prices exactly at the cap. In the third and final step, we compute market shares and surplus based on the capped prices. Since we remove the government intervention, we return the government budget to households via direct cash transfers.

Uniform Price Changes in Government Good To estimate what happens when the government good’s price changes, we proceed using a two-step analysis. We first change the price per liter for the government

good in nominal terms for a given change x

$$p'_g(x) = \max\{p_g + x, 0\} \quad (49)$$

(which is Equation 19 in the paper) and then we convert such change to constant MXN. Once we compute the new government prices, we estimate the new equilibrium market share and prices using our iterative process described above (Equations 43 and 44) using as starting values the original status quo shares. Denote by $p'_m(x)$ the new equilibrium prices and by $\sigma'_m(x)$ the new equilibrium shares for market m .

The final step involves balancing the budget (Equation 42). The new allocation requires a budget B'

$$B'(x) = (c_g - p'_g(x)) \cdot \sigma'_g(x)$$

per market, which we compare to the observed budget and use cash transfers to net out any additional surplus / deficit. When $B'(x)$ is smaller than B , that means the current allocation requires *less* money than the observed status quo. We use the additional revenue to fund a cash transfer system of $B - B'(x)$ divided by the total population. All of the eligible households within the same calendar year receive the same amount of transfer. Ineligible households receive nothing. To keep the problem symmetric, when $B'(x)$ is bigger than B , we raise the extra revenue needed by drawing money from the eligible households.

H.2 Robustness

As mentioned in the paper's main body, the model makes certain simplifying assumptions that we consider second order. This section discusses how our results would change if we were to change some of these assumptions. We focus on the entry of new competitors in the market and changes in the input prices (i.e., the marginal cost per liter of milk) below.

Entry and Exit We consider a stylized version of the entry of new competitors. To formally discuss entry, define a product χ_j as a combination of its observable product characteristics, unobservable product characteristics, and cost (x_j, ξ_j, c_j) . We simulate the entry of (randomly chosen) new products, repeating the process R times. For a given repetition r in market m , we proceed in two steps:

1. Draw $\chi_j^{(r)}$ from the estimated product distribution, and assume a new firm sells the new product.
2. Estimate counterfactual equilibrium prices and market shares based on existing products and $\chi_j^{(r)}$.

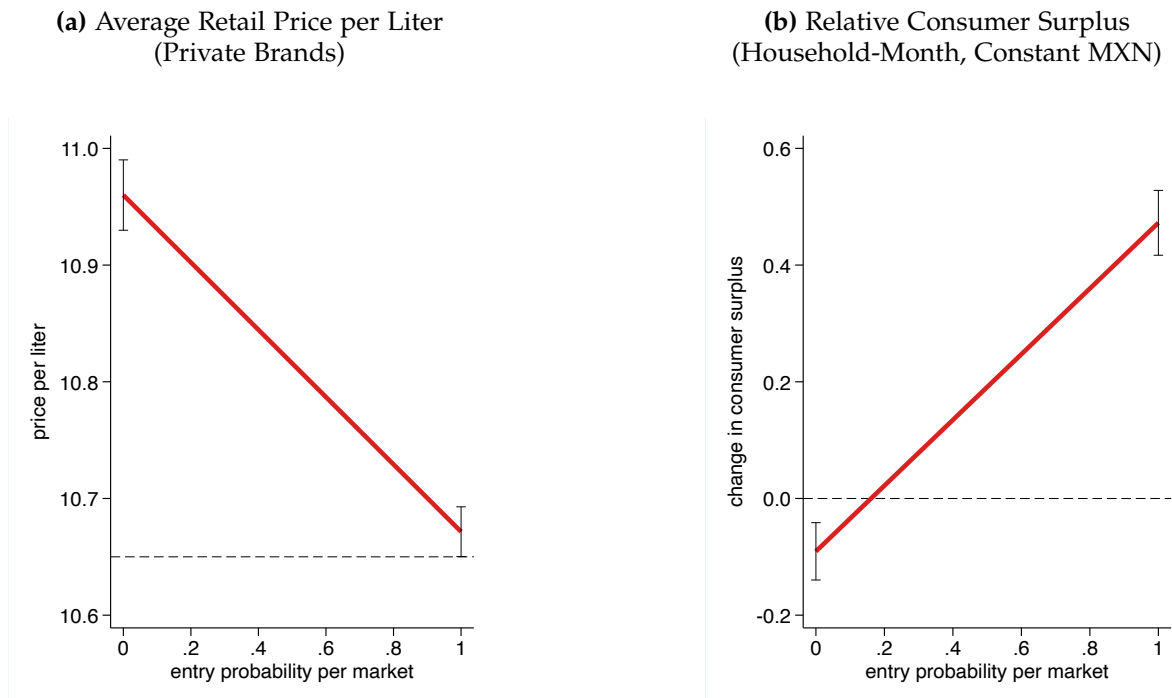
After simulating entry R times, we then compare the average across simulations, holding the transfers to households at the cash transfers level.

Figure OA-24 plots our simulation results. In the x-axis, we show the proportion of markets in which there is one entrant. On the left side of each graph, we have the cash transfer equilibrium outcomes with no entry. In the rightmost part of each graph, we have the market outcomes if an entry happens in all markets.

Panel (a) shows the average price per liter among the private brands. Panel (b) shows the consumer surplus relative to the direct provision world. Negative values in the y-axis indicate that surplus is higher with direct provision; positive values indicate that surplus is lower. The gray dashed lines are the average price per liter and the consumer surplus in the direct provision observed status quo.

As shown by the graph, we would need an entry in at least 20% of the markets to have the same effect in consumer surplus as the one with direct provision of the (low quality) government good. If there were an entry in more than 20% of the markets, then consumer surplus would be higher with cash than with direct provision. Since the government good is of a lower quality, we would not even need an entry in all markets to have a higher consumer surplus. Private firms' entry generates a higher surplus than direct provision, as private competitors have a higher quality than the government good. To have the same price effect as direct provision, we would need private firms to enter all markets.

Figure OA-24: Entry of New Retailers in the Market



Notes: This figure summarizes the entry simulation results.

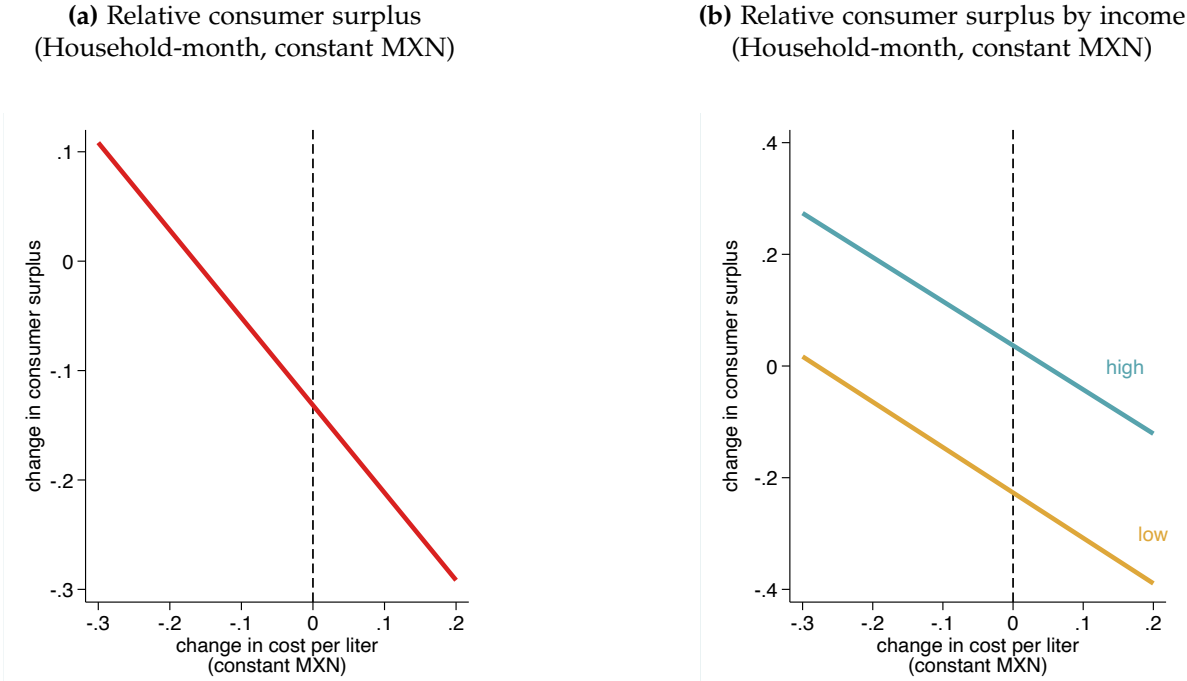
Cost Changes with Government Provision To estimate the robustness of our results to changes in the input costs, we compute what would happen if the marginal cost in the private market were to change. We consider uniform changes of the following form for product j in market m

$$c'_{jm}(x) = c_{jm} + x \tag{50}$$

and compute the new equilibrium outcomes using the new costs per liter in the private market, holding transfers to households at the level of the cash transfers.

Figure OA-25 summarizes our estimation results. The x-axis shows the different values of x in Equation 50 that we consider in the analysis. With $x = 0$, we show the market outcomes under the cash transfers counterfactual. Panel (a) shows the average consumer surplus with cash. Panel (b) disaggregates surplus across socioeconomic status levels. Negative values in the y-axis indicate that surplus is higher with direct provision; positive values indicate that surplus is lower. Costs would need to decrease by 19 cents—60% of the change in prices under the main counterfactual comparison or around 2.3% of the average cost per liter—for cash to dominate direct provision. Note that even for the largest changes (in absolute terms) in the cost per liter considered ($x = -0.4$), consumer surplus among the poor is larger with direct provision than with cash, highlighting the redistributive nature of direct provision.

Figure OA-25: Changes in the Cost per Liter of Milk (Constant MXN)



Notes: This figure summarizes the simulation results with new costs in the private market absent the government intervention.