

Public Procurement and Private Product Distribution: Evidence from WIC and the Infant Formula Market

James Okun
MIT

Yuci Zhou
Princeton*
(Job Market Paper)

September 20, 2026

[\[Click here for the latest version\]](#)

Abstract

Governments can lower the cost of providing in-kind benefits by restricting beneficiaries to a single supplier's products, which induces manufacturers to compete through discounts. When private retailers distribute these benefits, their responses to procurement rules can affect the choices available to private-market consumers. We study this tradeoff in infant formula procurement by the Special Supplemental Nutrition Program for Women, Infants, and Children (WIC), where three manufacturers bid for contracts that require retailers to keep the contracted brand in stock. We find evidence that retailers' stocking decisions are a driver of the increase in the contracted brand's private-market share. We rationalize this with a model in which consumers substitute across brands during stockouts, so keeping the contracted brand reliably in stock lowers retailers' return to keeping competing brands in stock. Counterfactual simulations show that, among designs that keep formula free for beneficiaries, contracting with two brands and requiring retailers to keep both in stock raises annual private-market consumer surplus by \$340 million (12% of current private-market formula spending), while leaving net program spending essentially unchanged. This design broadens product availability for private-market consumers through the stocking requirements, while preserving competition for contracts, since manufacturers still face the threat of exclusion.

* Okun: jokun@mit.edu; Zhou: yucizhou@princeton.edu. We are grateful for Nick Buchholz, Janet Currie, Adam Kapor, Jakub Kastl, Nikhil Agarwal, Amy Finkelstein, Nancy Rose, and especially Kate Ho for their invaluable guidance and support throughout this project. This project has benefited from comments from seminar participants at Princeton Graduate IO Workshop, Princeton Industrial Relations Section Graduate Workshop, Center for Health and Well-being Brown Bag Lunch Seminars, the MIT Second Year Paper class, the MIT Industrial Organization Lunch, and the MIT Public Finance Lunch. Okun gratefully acknowledges financial support from the National Institute on Aging under grant T32-AG000186, the Alfred P. Sloan Foundation under grant G-2024-22690, the Jerry A. Hausman Graduate Dissertation Fellowship, and the George Obie Shultz Fund at MIT which funded the purchase of data for this project. Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

1 Introduction

This paper studies procurement designs for public programs that purchase goods from manufacturers and distribute them to beneficiaries through private retailers, focusing on the case of nutritional assistance programs.¹ Retailers play a key role in determining consumer welfare: they decide which products are available to buy and thereby affect how upstream manufacturers compete and set prices. Existing studies provide evidence on private-market spillovers of public programs, including how retailers respond (e.g., Duggan and Scott Morton, 2006; Jaravel, 2018; Handbury and Moshary, 2021; Meckel, 2020), but typically do not model the stocking incentives behind them. We model retailers’ stocking responses to procurement rules and ask how the design of these rules can balance program spending and private-market consumer surplus in such programs.

Our setting is the procurement of infant formula by the Special Supplemental Nutrition Program for Women, Infants, and Children (WIC). Infant formula provides complete nutrition for infants who are not breastfed, and WIC provides it as a means-tested, in-kind benefit to ensure access among low-income families. Stringent safety regulations and requirements, similar to those for pharmaceuticals, make manufacturer entry costly and help sustain a concentrated market. With three national brands dominating the market, WIC contracts with one of these brands in each state in an effort to minimize fiscal costs; manufacturers bid for contracts by offering discounts to the program. The program distributes benefits through the same retailers consumers ordinarily use for groceries and household items; these retailers also serve consumers paying for infant formula out of pocket. To implement the sole-source contract, program participants can redeem benefits only for eligible products of the contracted “WIC brand,” which retailers are required to keep in stock to ensure their access, in the form of a minimum stocking requirement.

We begin by demonstrating that WIC procurement design has large effects on the private market for formula via retailers’ stocking incentives. This analysis combines NielsenIQ’s retail and household purchase data with information on changes in WIC contract brands across states and time between 2006 and 2020. Consistent with earlier studies (e.g. Huang and Perloff, 2014; Rojas and Wei, 2019; An et al., 2023), winning a WIC contract substantially increases a brand’s market share among private-market consumers.

Among several potential explanations for this *private-market spillover*, we find evidence that

¹Examples of food assistance delivered through private retailers include programs in the United Kingdom (NHS Business Services Authority, n.d.), Indonesia (Mahardika, 2022), Chile (Ministerio de Economía, Fomento y Turismo, 2023), Nigeria (World Food Programme, n.d.), and of course the Supplemental Nutrition Assistance Program (SNAP) in the U.S., among many others.

retailers’ stocking decisions are an important driver. Grocery stores and pharmacies (“small-box stores”) are compared with mass merchandisers (“big-box stores”), which typically have larger square footage, hold more inventory, and can restock more easily to satisfy all consumers.² The spillover effect in small-box stores is nearly three times as large as in big-box stores. This difference points to the role of retailers in the spillover, consistent with Abito et al. (2026), who show that the spillover is concentrated in physical retail. To investigate product availability for private-market consumers, we use months with zero sales as a proxy for product unavailability. Losing the WIC contract also raises a brand’s probability of monthly zero sales in small-box stores, while the rate in big-box stores is largely unaffected. Tracking private-market consumers who previously purchased the old WIC brand, we find that small-box store shoppers shift over a quarter of their purchases to the new WIC brand, even though consumers are generally loyal to one brand. In contrast, the purchasing behavior of big-box shoppers changed little. Together, these findings suggest that small-box retailers use the WIC brand’s inventory to serve private-market consumers rather than keeping competing brands consistently available.

To understand how procurement affects retailers’ stocking decisions, pricing by manufacturers, and ultimately consumer surplus, we develop and estimate a model of households’ purchasing, retailers’ stocking decisions, and manufacturers’ pricing. The model rationalizes the observed shift in private-market purchases toward the WIC brand as an outcome of retailers’ optimal stocking responses to procurement rules. The model is used to simulate the effects of alternative procurement designs on program spending and private-market consumer surplus in equilibrium. Retailers choose inventories of substitutable products jointly to maximize their total profits from all products. Inventory held to serve program participants can also be used to satisfy private-market demand, changing retailers’ incentives to keep substitute brands available.

At the core of the model, we extend the standard single-product (s, S) dynamic inventory model of Arrow et al. (1951) to a joint inventory problem across substitutable brands. Empirical models of inventory typically abstract from cross-product inventory effects or capture them through estimated cost parameters. We model these effects directly, so that stocking one brand changes the return to replenishing its substitutes.³ This framework can apply more broadly to policies or other supply-side factors that change retailers’ incentives to stock particular products.

²Big-box stores typically occupy 100,000–200,000 square feet, roughly 3–5 times the size of an average supermarket and about 10 times that of a drugstore (Institute for Local Self-Reliance, 2008; Kaufman et al., 2000). Section 2.1 describes the differences in store size and distribution arrangements across these retail formats.

³For example, Sanders (2024) estimates a model of dynamic pricing and inventory decisions assuming independent demand between products, and Aguirregabiria and Guiton (2023) capture cross-product inventory effects through estimated cost parameters. Mahajan and van Ryzin (2001) develop a theoretical single-period inventory model with stochastic demand and consumer substitution during stockouts.

In the model, consumers substitute among available options, possibly purchasing an inferior alternative when the preferred option is out of stock (as in Conlon and Mortimer, 2013). Retailers decide when and how much of each brand to stock before observing that day’s demand. They trade off revenue from serving WIC and private-market consumers against wholesale and fixed ordering costs.

To capture how WIC’s minimum stocking requirement affects stocking choices, the model incorporates a penalty for stocking out of the contracted brand, which summarizes the expected cost of violating the requirement. This penalty incentivizes retailers to hold buffer stocks of the WIC brand, but it reduces the return to keeping non-WIC brands in stock because stores still earn revenue when consumers switch to the WIC brand during stockouts of other brands. The model is completed by incorporating manufacturers, who set wholesale prices and bid for WIC contracts anticipating retailers’ ordering decisions. The resulting equilibrium links procurement design to consumers’ choice sets and upstream competition through retailers’ stocking incentives.

A central challenge in estimating the model is separating consumer demand from product availability, both of which determine observed sales. This challenge is addressed by combining complementary datasets. Preferences are estimated by assuming complete product availability in big-box stores, consistent with empirical findings about their stocking behavior; this avoids bias from attributing the effects of stockouts to consumer preferences. Similarities in the observed demographics of consumers shopping at different store types motivates treating preference parameters are treated as common across store types, while price sensitivity is allowed to vary with household income.

Counterfactual simulations require stocking cost estimates for stores across the country, but NielsenIQ’s retail sales data do not record inventory orders. We therefore estimate the inventory model using daily orders and sales from more than 500 grocery stores during 2020–2024. Assuming that stocking costs vary with store characteristics in the same way in both datasets, we map these estimates to NielsenIQ stores and, together with the preference estimates, recover the primitives that determine stocking for that broader sample. We then use observed bids and model-implied manufacturer profits to estimate additional primitives governing manufacturers’ bidding behavior for WIC contracts.

The estimates imply that retailers place a high value on satisfying WIC’s stocking requirement and therefore have a strong incentive to keep the WIC brand in stock. Most manufacturers’ bids for WIC contracts imply net revenue per unit of formula sold to WIC below estimated production cost. These bids are consistent with manufacturers accepting losses on WIC purchases in exchange for the private-market spillovers. Competition for contracts thus allows WIC to

recapture part of the rent that the spillover generates for manufacturers.

Counterfactual simulations reveal a tradeoff between broadening private-market consumers' choice sets and preserving the fiscal cost-savings from manufacturer competition for WIC contracts. For WIC contracts covering 2016, we compare designs that contract with one, two, or all three brands, keeping formula free for WIC beneficiaries and requiring retailers to stock every contracted brand, accounting for how retailers adjust their stocking under each design. Under the current design, each state awards its contract through a first-price auction. When some brands are excluded, contracts are awarded through an auction that selects the brands that maximize manufacturers' combined profit from WIC and private-market sales. Each winner pays the government the additional profit, including from the private-market spillover, that its rivals would have earned without it. Under this auction, contracting with two brands raises private-market consumer surplus by \$340 million annually relative to the current design, equivalent to 12 percent of private-market formula spending. Government spending on WIC remains essentially unchanged. Contracting with one brand but switching to this auction instead reduces net program spending by \$279 million annually, with little change in private-market consumer surplus.

Relative to two brands, contracting with all three brands achieves an additional gain in consumer surplus at the expense of \$5.84 in additional program spending per dollar of consumer surplus. The consumer surplus gains from contracting with two or all three brands are mainly driven by the expansion of private consumers' choice sets. The advantage of contracting with two brands is that one brand remains excluded, so manufacturers are still willing to pay the government to serve WIC beneficiaries and capture the associated private-market sales. Contracting with all three brands eliminates that competition and the associated manufacturer payments. This cost of eliminating exclusion is consistent with findings from An et al. (2023), who highlight the role of competition for contracts in containing program costs.

Additional counterfactual designs examine whether changes to other procurement rules can improve private-market welfare while controlling program spending and maintaining the quantity of formula provided through WIC. None of these designs delivers welfare gains comparable to contracting with two brands without increasing program spending or reducing WIC redemption. These alternatives either weaken the stocking incentives that support product availability or the manufacturer competition that finances the program.

Related literature Prior research shows that public programs can affect private markets through manufacturers' pricing (Duggan and Scott Morton, 2006) and through retailers' price responses (Hastings and Washington, 2010; Goldin et al., 2022; Leung and Seo, 2023), changes prod-

uct variety and assortments (Jaravel, 2018; Handbury and Moshary, 2021; Andreyeva et al., 2011), and program participation decisions (Meckel, 2020; Meckel et al., 2023; Byrne et al., 2025). These studies show how retailers respond to public programs but typically do not model the stocking incentives behind these responses. We show that because retailers serve program participants and private-market consumers from the same inventory, sole-source contracts can also affect how often private consumers face stockouts of competing brands. We then evaluate alternative procurement designs accounting for retailers’ endogenous stocking responses. More broadly, spillovers through retailers’ shared inventory can arise whenever policies or other supply-side factors change retailers’ incentives to stock particular products. Dean et al. (2026), for example, show that when commercial insurers restrict which drugs their enrollees can use, utilization of the covered drugs increases among Medicare patients, whose coverage is not restricted.

Two contemporaneous papers studying WIC’s infant formula contracts are closest to ours. Abito et al. (2026) also document the role of physical retail in the spillover, consistent with our descriptive evidence, and show that it persists among retailers with observed purchases of all three national brands. We provide evidence of intermittent stockouts in small-box stores as a channel and explain why this availability pattern can arise from optimal stocking even when retailers carry substitutes. An et al. (2023) evaluate alternative procurement rules using a model of manufacturer pricing and reduced-form bid functions and, as in our counterfactuals, find that competition for contracts contains program costs. In our model, private-market spillovers respond to the number of contracted brands and retailers’ stocking requirements, allowing us to quantify the resulting tradeoff between private-market consumer surplus and program spending.

The rest of the paper proceeds as follows. We describe the setting and the data in Section 2, and document the private-market effects of WIC procurement and present evidence on retailers’ stocking responses in Section 3. We present the model of retailers, consumers, and manufacturers in Section 4, and the estimation strategy in Section 5. We report the estimates in Section 6 and evaluate the alternative procurement designs in Section 7.

2 Setting and Data

2.1 The Infant Formula Industry

Infant formula, an alternative to breastmilk, provides complete nutrition for infants who are not breastfed. Infants receive it for roughly the first twelve months of life, after which whole

cow's milk becomes suitable (American Academy of Pediatrics Committee on Nutrition, 1992; Grummer-Strawn et al., 2008). Four firms manufactured infant formula in the United States during 2006–2020. We focus on milk-based powder products of the three national brands: Abbott Nutrition's Similac, Mead Johnson's Enfamil, and Nestlé's Gerber Good Start. A fourth manufacturer mostly produces private label products for retailers. Our analysis focuses on milk-based powder products from the three national brands.⁴

Stringent regulatory requirements imposed by the Food and Drug Administration (FDA) make entry to the infant formula market costly. Manufacturers must demonstrate, generally through clinical growth-monitoring studies, that their formula supports normal physical growth, and their production facilities must meet stringent manufacturing and sanitation requirements and are subject to FDA inspections.⁵

Consumers purchase formula through retailers with different stocking capacities. We distinguish *big-box stores* (mass merchandisers) from *small-box stores* (grocery and drug stores). Big-box stores generally have more shelf and storage space, and their overall scale and distribution arrangements facilitate maintaining product availability. Small-box stores have less space to hold large inventories, and restocking frequently is costly.⁶ These differences motivate our subsequent analyses of the two store types in Section 3.

2.2 The WIC Program and Formula Procurement

The Special Supplemental Nutrition Program for Women, Infants, and Children (WIC) is a means-tested nutrition assistance program for low-income pregnant and postpartum women, infants, and young children at nutritional risk, administered by the United States Department of Agriculture (USDA). WIC provides vouchers for specified nutritious foods, including infant formula. Its restriction of benefits to prescribed foods reflects a paternalistic rationale: directing consumption toward nutritional needs identified by the program (Currie and Gahvari, 2008). Formula benefits are available to infants who are partially or fully formula-fed. Participants redeem these benefits for specified quantities of eligible formula in retail stores at zero monetary

⁴Formula can use other protein sources, such as soy, and is also available as liquid concentrate or ready-to-use liquid. We exclude these alternatives because they represent a small share of sales—in NielsenIQ's Retail Scanner data for 2006–2020, milk-based products account for 92.2 percent of total prepared-ounce volume and powder products for 90.2 percent. The three national brands account for 98.4 percent of volume within milk-based powder products.

⁵Manufacturers must also register their production facilities and submit information about new formula at least 90 days before marketing it. See Code of Federal Regulations (n.d.a, §§106.96(b)–(c), 106.110, and 106.120).

⁶Drugstores generally occupy 11,000–15,000 square feet and the average supermarket approximately 40,000 square feet (Institute for Local Self-Reliance, 2008); supercenters occupy 100,000–200,000 square feet (Kaufman et al., 2000).

cost.⁷ The program is a large buyer of infant formula—in 2016, there are about 1.6 million infants who are registered as fully or partially formula-fed in the WIC program in an average month, around 40 percent of all infants born in the United States.⁸ WIC redemption accounts for at least a third of sales volume for infant formula.⁹

WIC contains its spending on formula by leveraging *sole-source procurement contracts*: each state signs a contract with one manufacturer whose formula beneficiaries redeem with WIC vouchers at eligible retailers. States award these contracts by auction every two to five years, either individually or jointly through purchasing alliances (at least two WIC state agencies that purchase together). Manufacturers offer a rebate to the program for each unit redeemed through the program’s vouchers. This rebate is specified by each manufacturer’s bid in the auction. The contract is awarded to the manufacturer that has the lowest net wholesale price, which is given by the wholesale list price net of the rebate.¹⁰

Retail stores approved by the WIC program dispense formula to beneficiaries and private-market consumers. For products redeemed by WIC beneficiaries, the program reimburses stores at the retail price of the formula, subject to state limits on allowable reimbursement.¹¹ The program also imposes a minimum stocking requirement—an approved store must keep the WIC-eligible products of the contracted brand in stock (Pelletier et al., 2017).

WIC restricts formula redemptions to eligible sizes and varieties of the contracted brand, meaning that some products are never redeemable with vouchers. Sales of these products allow us to study private-market outcomes separately from voucher redemptions.

2.3 Data

We use four main datasets. Additional data sources and construction details are described in Appendix Section A.

⁷Except Vermont and Mississippi, all states distribute WIC benefits through private retailers before 2020; we exclude these two states from the analysis sample.

⁸WIC’s coverage is calculated using USDA administrative data on program participation and CDC’s Natality data on births.

⁹Calculated among the sampled retailers and milk-powder products. Using a similar retail dataset, USDA estimates that WIC purchases account for over 50% of the retail value of all formula sold; the difference likely reflects the undersampling of stores with large WIC redemption volumes (Oliveira et al., 2010).

¹⁰If the manufacturer’s wholesale list price changes during the contract, the rebate changes by the same amount, keeping the net wholesale price fixed; see Code of Federal Regulations (2011, §246.16a(c)(6)(iv)).

¹¹WIC vendors may not charge the program more for a product than they charge other customers; see Oliveira and Prell (2004).

WIC Contracts We use administrative data from the USDA on all sole-source contracts during 2006–2020. The data record the winner, national wholesale price¹², the winning rebate and resulting net wholesale price, the contract duration, and the characteristics of the products WIC beneficiaries can redeem. The set of these *WIC products* helps us classify purchases by WIC participants. Davis (2016) supplements these data with losing bids for contracts that began before 2016, which we use to estimate the model for these auctions. We use the alliance membership recorded in these sources to identify states that auction their contracts jointly.

Scanner Data Store-level prices and quantities come from NielsenIQ’s Retail Scanner data from 2006 to 2020 (our study period), which we later refer to as the *Scanner Data*. The unit of observation in the raw data is a Universal Product Code (UPC) by store and week. For the descriptive analysis, we aggregate these observations to the brand–store–month level, sometimes restricting to UPCs not in the WIC product set. We observe about 34,000 stores that sold infant formula, each with its county, 3-digit ZIP code, and retail channel.

Inventory Data Store-level inventory decisions come from a proprietary dataset covering 544 grocery stores from four chains in seven states, which we call the *Inventory Data*. From April 2020 through June 2024, we observe wholesale prices, retail prices, quantities ordered from manufacturers or distributors, and quantities sold at the UPC level and daily frequency.¹³ We do not observe inventory levels.¹⁴ Based on our knowledge of these stores, we believe they would be classified as small-box stores by NielsenIQ. Store identities are masked in the Scanner Data, so the two datasets cannot be linked at the store level.

Consumer Panel Household-level formula purchases come from NielsenIQ’s Consumer Panel data, 2006–2020. The panel records each purchase’s UPC, quantity, store channel, and the household’s demographics. The demographic variables include household size, coarse income bins, and a self-reported WIC participation indicator. A missing indicator could mean either nonparticipation or survey nonresponse. We use these data to infer consumers’ income eligibility for WIC, follow individual private-market consumers across WIC contract changes, and

¹²The wholesale price is for the “contracted product,” the 90-prepared-ounce version of the brand’s non-specialty formulation.

¹³We exclude January 2022 through January 2023 to omit the disruption surrounding the 2022 infant formula shortage. For more details, see FDA (2023). We include observations during 2020–2021 (during the COVID-19 pandemic) because ordering and sales pattern are similar to those in the subsequent period. This allows us to use more observations from our limited inventory samples.

¹⁴We do not observe stores’ starting inventory, and the data does not record product spoilage, so we cannot reliably reconstruct inventory levels from orders and sales alone.

construct statistics for model estimation.

Data Linkages and WIC Sales Classification We link WIC contracts to UPCs by matching each contract’s product name, ingredient base, and size, supplemented by the states’ Approved Product Lists (APL), which we compile through public records requests.¹⁵ Within our milk-based powder sample, we define a product for the descriptive analysis as a brand and WIC eligibility combination. WIC-eligible products are the UPCs that have been a WIC product in any state at any time; private-market products are the complement. For the descriptive analysis, we use private-market products to isolate purchases paid for out of pocket. Because these products are never eligible for voucher redemption, this approach identifies a subset of private purchases without having to separate private market sales and WIC redemptions of WIC-eligible products.¹⁶

In model estimation, we also account for private-market purchases of WIC-eligible products. Using Consumer Panel households classified as private-market consumers, we estimate the private-market share of eligible-product purchases separately by brand and retail channel. We use these shares to derive private-market purchases from stores’ eligible-product sales of the local contract brand and assign the remaining volume to voucher purchases. Appendix Section A describes the household classification and sales allocation. We do not observe stores’ WIC authorization status and assume that all stores in our sample are WIC-authorized when estimating the model.¹⁷

3 Private-Market Effects of WIC Procurement

This section documents how WIC procurement affects the private market for infant formula and provides evidence that retailers’ stocking decisions are a driver of these effects. *Private-market spillover* denotes the effect of WIC contract status on a brand’s private-market outcomes. We first show that winning a WIC contract substantially increases a brand’s private-market share. We then assess the role of retailer stocking by comparing big-box and small-box stores and examine product availability. Finally, tracking consumers across contract changes to shows how their brand choices are affected by the stores they shop at.

¹⁵See Appendix Section A for a detailed description of the APL data.

¹⁶Since we do not have data on WIC redemptions, this is challenging to do without strong assumptions.

¹⁷Store surveys find that over 80% of larger grocery retailers are WIC authorized (Andreyeva et al., 2011; Gosliner et al., 2018).

3.1 WIC Contracts and Private-Market Shares

We quantify the effect of the WIC program on private-market outcomes using an event-study design. WIC contract auctions generate geographic and temporal variation in which brand holds the WIC contract.¹⁸ Taking the contract change as the event, we estimate the regression

$$y_{jst} = \sum_{k=-12}^{12} \delta_k \mathbb{1}\{t - t_{js}^* = k\} + \alpha_{js} + \tau_t + \varepsilon_{jst}, \quad (1)$$

where y_{jst} is an outcome for brand j in state s and month t , and t_{js}^* is the month brand j 's WIC status changes in state s . Depending on the outcome, y_{jst} uses all of the brand's products or only its private-market products (Section 2.3). The indicators $\mathbb{1}\{t - t_{js}^* = k\}$ are relative-month dummies; we normalize $\delta_{-1} = 0$ and bin event times beyond ± 12 months into the endpoints. The regression includes brand-by-state fixed effects α_{js} and calendar-month fixed effects τ_t . Standard errors are clustered by state. Brand-state pairs whose WIC status never changes during the sample period form the control group, and pairs that both gain and lose a contract are excluded (Appendix Figure A.1).¹⁹ We estimate (1) separately for the winning and losing brands to assess symmetry.


Figure 1 reports the estimates, where the outcome is the brand's share of prepared ounces sold. Panel (a) uses all products, including both WIC and non-WIC products. The winner's share rises by about 46 percentage points within three months of winning the contract, and the loser's falls by about the same amount. This partly reflects a mechanical increase due to the winner's WIC voucher redemptions. Panel (b) examines private-market spillovers using products never eligible for WIC vouchers, as described in Section 2.3. In the month before the contract change, the incoming and outgoing WIC brands account for 26 and 61 percent of private-market products' sales, respectively. The winner's private-market share then rises more gradually, reaching about 18 percentage points at 12 months; the loser's falls by a similar amount. Small pre-event trends are consistent with stores responding to contract changes that are announced before taking effect, usually three to six months ahead. The estimated spillovers, valued at observed wholesale prices, generate roughly three times as much manufacturer revenue as WIC voucher sales net of rebates,²⁰ suggesting that the private market is a principal source of contract value and helps explain the steep discounts firms offer the government.

¹⁸During 2006–2020, 65 of 170 contract awards (38 percent) changed the contracted brand. Appendix Figure A.1 shows the timing and geographic distribution of the events used in the event studies.

¹⁹As a robustness exercise, we also estimate all event studies with difference-in-differences estimators that are robust to the negative-weighting problem under heterogeneous treatment effects; see Appendix Figure A.2.

²⁰Per 90 prepared ounces, the median winning bid in 2006–2015 nets the manufacturer \$1.18, and the spillover adds about \$3.50 in private-market revenue.

Figure 1: WIC’s Effect on the WIC Brand’s Market Share



Notes: Event-study estimates of the WIC brand’s share of milk-based powder infant formula sales, by month relative to a state’s WIC contract change (vertical dashed line). Panel (a) uses all products; panel (b) restricts to private-market (never WIC-eligible) products, which are purchased only by private-market consumers. The *winning* series follows the brand that gains the state’s WIC contract; the *losing* series follows the brand that loses it. Shaded bands are 95% confidence intervals. Sample size: winning—42 states, 180 months, and 15,120 brand–state–month observations; losing—40 states, 180 months, and 14,760 brand–state–month observations.


3.2 Retailer Stocking and Consumer Choice

Private-market spillovers could arise from demand-side factors, such as consumers’ perception of the WIC brand as the government’s endorsement, or supply-side factors, such as retailer stocking or manufacturer scale economies. We argue that retailer stocking is a central driver of the spillover effects we observe.

If retailer stocking drives the spillover, the effect should be larger in stores where maintaining broad product availability is more difficult. This analysis compares the spillover effects observed in big-box and small-box stores. Compared with big-box stores, small-box stores have less inventory space and lack the scale and distribution arrangements that facilitate replenishment (Section 2.1). The private-market event study is re-estimated separately for the two store types. Figure 2 shows the estimates: 12 months after switching, the new WIC brand gains about 23 percentage points of market share in small-box stores, compared to about 8 percentage points in big-box stores, roughly a third of the effect in small-box stores. Big-box stores sell about a quarter of milk-based powder formula, measured in prepared ounces, and small-box

stores sell the remainder.²¹

Figure 2: Effect of WIC Status on Private-Market Shares, by Store Type



Notes: As in Figure 1, estimated separately for big-box (mass merchandiser) and small-box (drug and grocery) stores. The *winning* series follows the brand that gains the state's WIC contract; the *losing* series follows the brand that loses it. Shaded bands are 95% confidence intervals.

To examine whether this difference reflects stockouts of private-market products, we estimate a store-level event study around the loss of the WIC contract. We re-estimate (1) for the losing brand with stores in place of states. The outcome is an indicator for zero sales in a month, which is used as a proxy for product unavailability.²² Figure 3 shows that the outgoing WIC brand's monthly zero sales rate barely increases in big-box stores, but sharply increases in small-box stores by about six percentage points after 12 months.²³


Zero sales events are generally temporary and resolve quickly; 70.7% resolve within two months and 79.2% within three months. This pattern is consistent with intermittent stockouts rather than permanent withdrawal of the product from the store. The outcome variable is a proxy for stockouts lasting for a whole calendar month and does not capture more short-lived stockout events. The evidence indicates that stores keep less inventory of the outgoing brand after it loses the contract. Both results use private-market products, which the stocking require-

²¹Authors' calculation from the Scanner Data for 2016, the year we evaluate counterfactual designs nationally. Each store is reweighted by the inverse of NielsenIQ's coverage of its 3-digit ZIP code and channel. The calculated share is lower in earlier years, which could reflect changes in the consumption trends as well as composition of sampled stores.

²²The sales data do not reveal inventory or the duration of an individual stockout. Zero sales can also reflect low demand, so we interpret this proxy alongside the evidence on subsequent sales and household substitution.

²³For the brand that wins the contract, zero sales rate falls in both store types, and the difference between them is smaller than around losses (Appendix Figure A.6). A decrease in zero sales rate could reflect adding the brand to the assortment or fewer stockouts. We focus on contract losses, where the brand was already carried before the change, to rule out assortment changes.

Figure 3: Loss of the WIC Contract and Private-Market Product Availability



Notes: Event-study estimates of the probability that a private-market product of the losing brand has zero sales in a store-month. Estimates are shown separately for big-box (mass merchandiser) and small-box (drug and grocery) stores, by month relative to the WIC contract change (vertical dashed line). Shaded bands are 95% confidence intervals.

ment does not cover. The contract therefore changes how stores stock the brand as a whole, and not only the products that WIC beneficiaries redeem.

A primary concern with results is that the difference between store types could instead reflect selection—consumers who prefer the WIC brand may systematically sort into small-box stores. To rule this out, we examine individual consumers’ responses to changes in the WIC brand using NielsenIQ’s Consumer Panel data. It is reasonable to hypothesize that a consumer’s formula preference does not change over a purchasing spell, defined as a sequence of purchases with no more than 100 days between purchases, spanning less than a year, which likely represents a series of transactions for the same child.²⁴ Appendix Figure A.9 shows that households are largely loyal to a single brand, consistent with parents’ reluctance to switch their infant’s formula.²⁵ If the contract change nonetheless shifts their choices, it likely operates through what is available at their preferred retailer. We therefore examine private-market consumers²⁶

²⁴Our definition is motivated by the fact that formula-fed infants consume formula through roughly the first year of life — see Section 2.1 for more details.

²⁵Switching formula is often medically disruptive for the infant. Smilowitz et al. (2024) show that infants developed digestive problems when parents were forced to switch formula during the 2022 U.S. shortage.

²⁶We use reported WIC participation, household income, and household size to identify private-market households for the consumer-switching analysis. See Appendix Section A for the classification rule.


whose brand before the contract change was the outgoing WIC brand that lost the auction, comparing those who shop in big-box and small-box stores after the contract change. Consumers are classified by where they shop after the contract change to measure the effect of their retail environment. Figure 4 shows the result. Compared to the pre-period share of 98% purchasing the outgoing WIC brand, consumers who shop in small-box stores in the post period reallocated 28% of their purchases to the incoming WIC brand, while those who shop in big-box stores are mostly unaffected, reallocating only 5% of purchases to the new WIC brand. Appendix Table A.2 reports household-level regressions of brand and store choice around contract changes. In Appendix Table A.1, NielsenIQ’s Consumer Panel data is used to show that the consumers who shop in small-box and big-box stores are balanced on their observable characteristics, as additional supportive evidence that consumer selection does not drive the difference in spillovers between store types.

3.3 Alternative Explanations

We assess a variety of alternative explanations in this section. Ultimately, the most compelling result on mechanisms is retailer stocking so we develop our model with this mechanism in mind. We allow for other demand-side preferences for the WIC brand to be present in the model as well.

We also consider two channels that could generate private-market spillovers without changes in retailer stocking. First, it is possible that WIC benefits do not cover the full demand of WIC consumers, who may then purchase more of the WIC brand out of pocket. We find that WIC households in the Consumer Panel likely stay within the limits of their benefit packages, so spillovers of this magnitude are unlikely to be due to WIC beneficiaries supplementing their benefits. Second, winning a WIC contract also raises a manufacturer’s production volume, which could lower its cost of serving the private market if there are economies of scale. However, there are little response of private-market retail prices to contract changes, and a firm’s bid does not respond to plausibly exogenous growth in its volume in other markets. Finally, Abito et al. (2026) find that spillovers persist among retailers selling all three national brands. In our Scanner Data, 97% of small-box stores sell all three brands at some point during a contract period, so we argue that their finding is consistent with intermittent stockouts in small-box stores. Appendix Section A.2 reports these analyses in detail.

Figure 4: Consumers' Brand Choice After WIC Contract Change, by Store Type



Notes: Formula-purchasing households whose dominant brand before the contract change was the outgoing WIC brand (at least 75% of pre-period quantities), split by the store type where they predominantly shop after the contract change. Because the classification requires dominance rather than exclusivity, the group's pre-period share on the outgoing brand is 98% rather than 100%. Each pair of bars gives the quantity-weighted composition of the group's formula purchases before and after the change: the share staying on the outgoing brand, moving to the incoming WIC brand, and (big-box only) moving to a third national brand. A purchasing spell is a sequence of purchases with no more than 100 days between successive purchases, spanning less than 365 days. Weights are pre-period total quantities. Quantities are in prepared ounces. Sample sizes: $N_{\text{small-box}} = 37$ and $N_{\text{big-box}} = 48$ household purchasing spells. Source: NielsenIQ Consumer Panel, 2006–2020.

4 Modeling the Infant Formula Industry

We develop a vertical model of the infant formula industry, connecting consumers, retailers, and manufacturers. The model answers two questions raised by the descriptive evidence: how sole-source procurement affects private market outcomes through retailer stocking, and how the value of the contract determines manufacturers' bids and wholesale prices in equilibrium.

The model consists of three components: consumer demand, retailer stocking, and manufacturers' wholesale pricing and bidding. The first part, based on Conlon and Mortimer (2013), is a model for calculating sales to consumers given inventory constraints, assuming consumers choose the best product available to them. This model pins down retailers' revenue as a function of inventory. Second, at the core of the model, consumers' choice sets result from each store's profit-maximizing stocking decisions. We develop a dynamic model of each retailer's inventory problem, wherein the retailer chooses inventory to maximize its expected profit across all the brands it carries, extending the standard (s, S) inventory model (Arrow et al., 1951) to a multi-product setting. The retailer's equilibrium inventory policy captures how cross-product cannibalization shapes stocking in the presence of the WIC program. Third, retailers' stocking decisions transmit to manufacturers upstream. Stocking policies determine the value of each WIC contract through private-market spillovers, which determines how much of a discount manufacturers give the WIC program on voucher purchases to win each contract. The resulting allocation affects how retailers' orders respond to wholesale prices, thereby affecting wholesale price competition.

4.1 Sales under Inventory Constraints

Following Conlon and Mortimer (2013), sales is modeled as consumers choosing sequentially from choice sets that shrink as products stock out. At a given store, a random number of private-market consumers arrive each period. We suppress the period index t and write the realized arrivals as N , an i.i.d. draw across periods. Consumers choose in the order they arrive, indexed $i = 1, 2, \dots, N$. Each chooses one available brand from a set of J brands or the outside option, picking brand j from a choice set a with probability $d_j(p, \xi, a)$,²⁷ a function of prices p , demand shocks ξ , and the choice set a . Write $y_j^i(p, \xi, a) \in \{0, 1\}$ for the indicator that consumer i buys brand j at prices p and demand shock ξ when facing a , so that $\Pr(y_j^i(p, \xi, a) = 1) = d_j(p, \xi, a)$. For simplification, each purchase is assumed to be a transaction of $\bar{y}$ ounces. Let Q_j denote the store's inventory of brand j available to private-market consumers. In the retailer model

²⁷ s_j would have been the more standard notation in the literature, but d is chosen to avoid notation conflict with the s component of the inventory policy in the retailer model.

(Section 4.2), WIC voucher consumers buy from the same pool of inventory before private-market consumers arrive. A brand is available to i until the consumers before her purchase the last unit,

$$j \in a_i \iff \bar{y} \sum_{i'=1}^{i-1} y_j^{i'}(p, \xi, a_{i'}) < Q_j, \quad (2)$$

so the first consumer faces all stocked brands, $a_1 = \{j : Q_j > 0\}$. The store's profit from private-market consumers depends on realized sales,

$$Y_j^{\text{priv}} = \bar{y} \sum_{i=1}^N y_j^i(p, \xi, a_i). \quad (3)$$

The subscript $-j$ denotes all brands other than j ; thus $Q_{-j} = (Q_k)_{k \neq j}$.

4.2 Retailers

We present the inventory problem for small-box stores, where stockouts could arise from profit-maximizing stocking decisions when serving both WIC voucher and private-market consumers and satisfying the program's stocking requirement.

We model inventory at the brand level. Although the stocking requirement covers only the WIC-eligible products of the contracted brand (Section 2.2), there are spillover and availability effects in the brand's other products (Section 3), indicating that stores respond to the contract at the brand level. We therefore model the store's choice of how much of each brand to stock.

Small-box stores do not compete with each other over price or inventory.²⁸ Retail prices are fixed within each inventory problem. Store s prices brand j at a fixed multiple of its national wholesale price,

$$p_{js} = m_{js} w_j, \quad (4)$$

which allows manufacturers' pricing decision to reach retailers and consumers.

Setup Time is discrete, and a period is one day. At the start of each period, the store's state $q = (q_1, \dots, q_J)^{29}$ is the leftover inventory from the prior period — for simplicity, the base model

²⁸There is a large incidence of uniform retail prices within chains in our data, consistent with DellaVigna and Gentzkow (2019); this motivates treating prices as given, a standard assumption in inventory models.

²⁹Suppressing the store subscript s for ease of exposition, since the inventory problem for a small-box store is a single-agent problem.

omits the cost shocks the store observes before ordering, which are added back as state variables in the empirical model (Section 5.2). The store chooses $Q_j \geq q_j$, how much inventory to make available for consumers, which we refer to as the *post-order inventory* hereafter. The store pays an inventory cost for each brand,

$$C_j(q_j, Q_j) \equiv w_j(Q_j - q_j) + K \cdot 1\{Q_j > q_j\}. \quad (5)$$

This cost has two components: (1) the wholesale price w_j per unit, which the retailer pays to the manufacturer for any order; (2) the ordering cost K , incurred each time the store places an order for a brand.³⁰ These costs are known to the retailer.

The retailer also serves WIC voucher consumers, which extends the sales model of Section 4.1. Voucher sales are assumed to be random and price-inelastic, since WIC consumers do not pay any monetary costs on top of the voucher. Voucher sales are denoted by Y_j^V , so total sales are $Y_j = Y_j^V + Y_j^{\text{priv}}$.

As in any inventory model (e.g. Arrow et al., 1951), demand is stochastic, but the store chooses inventory before demand realizes. First, voucher demand ξ^V realizes; for the WIC brand, voucher sales are the minimum of demand and inventory, $Y_j^V = \min(\xi^V, Q_j)$, and are zero for non-WIC brands. Then, the private market size N and the demand shocks $\xi_1, \dots, \xi_J$ realize and, jointly with the inventory left after voucher sales, determine $Y_j^{\text{priv}}(p, Q - Y^V, \xi)$. $\xi = (\xi^V, \xi_1, \dots, \xi_J)$ denotes the full vector of demand shocks. The law of motion governs the next period's inventory,

$$q'_j = Q_j - Y_j^V - Y_j^{\text{priv}}. \quad (6)$$

The Value Function The store sets inventory by maximizing its present discounted value of profits from all products that it stocks, using a discount factor β . As analyzed in Arrow et al. (1951), in a single-product setting the store's strategy takes an (s, S) form: order up to S when leftover inventory q falls to or below s , otherwise do not order. When the retailer optimizes the stocking of J products, its profit is the sum over all products. Therefore, it must consider the fact that raising one product's inventory steals business from another, so the total revenue effect of stocking more j depends on the inventory of $-j$ and substitution patterns between j and $-j$.

To incorporate this dependency across products yet keeping the computation of inventory policies tractable, we restrict attention to policies that are separable by products: product j 's

³⁰The standard (s, S) inventory model (Arrow et al., 1951) also includes a per-unit carrying cost cQ_j ; it is omitted for simplicity.

inventory choice depends only on q_j and on the stationary marginal distribution of other brands' post-order inventory. To achieve this, the store's problem is formulated as follows:

1. Each product sets its own inventory policy to maximize the discounted value of the store's total profit given its belief about other products' stationary inventory distribution; and
2. In the optimum, the belief of each product coincides with the true stationary distribution of post-order inventory.

Then, the joint inventory optimization problem reduces to J value functions V_j with a one-dimensional state variable q_j ; these J value functions are internally consistent such that together they represent the store's joint inventory optimization problem. This problem is convenient to compute via a fixed-point algorithm, yet captures the business stealing effect of a product's high level of stock in the long run, which is central to our application.

Formally, let Φ_{-j} be the stationary distribution of substitute products' inventory Q_k for $k \neq j$. Given the store's own Q_j and a realized substitute inventory Q_{-j} and demand shock draw ξ , the expected per-period profit that can be influenced by brand j 's stocking policy is the store's *total revenue* net of j 's own inventory cost,

$$\pi_j(Q_j, Q_{-j}, \xi; q_j) \equiv \sum_{k=1}^J p_k \mathbb{E}[Y_k(p, Q_j, Q_{-j}, \xi)] - C_j(q_j, Q_j), \quad (7)$$

where the inner expectation integrates over arrivals N and consumer taste shocks. The revenue sums over all brands because stocking more of j steals business from its substitutes; only j 's own cost enters, since substitutes' costs do not respond to how much j orders.

As discussed in Section 2.2, the WIC program requires the store to keep the contracted brand's WIC-eligible products in stock. This requirement is modeled as a perceived penalty on the risk that the WIC brand runs out,

$$\Lambda_j(Q_j; \Phi_{-j}) \equiv 100 \lambda_{\text{WIC}} 1\{j = \text{WIC}\} \mathbb{P}(Y_j = Q_j). \quad (8)$$

The coefficient λ_{WIC} measures the perceived cost in dollars per percentage point of stockout probability. The penalty enters each period's payoff when choosing Q_j , before the randomness of demand and consumer arrivals realizes.

Inventory is modeled at the brand level, so the penalty is also modeled as applying to stockouts of the whole brand, not just the products that are part of WIC's minimum stocking requirement. A brand's varieties are close substitutes, such as different container sizes of the same formulation. Stockouts of varieties not eligible for WIC redemption would redirect private-market

consumers to WIC-eligible products, increasing their stockout risk. Therefore the store has an incentive to hold more inventory of both sets of products. We interpret λ_{WIC} as summarizing the incentive behind product-level stocking choices that stores make to satisfy the program's requirement.

The recursive formulation of the store's problem can then be written as a value function

$$V_j(q_j; \Phi_{-j}) = \max_{Q_j \geq q_j} \left\{ \mathbb{E} \left[\pi_j(Q_j, Q_{-j}, \xi; q_j) + \beta V_j(q'_j; \Phi_{-j}) \right] - \Lambda_j(Q_j; \Phi_{-j}) \right\}, \quad (9)$$

where β is the store's discount factor and the expectation is over the belief $Q_{-j} \sim \Phi_{-j}$ and the period's transitory shocks (N , ξ , and the consumer taste shocks that determine Y_j).

Characterizing Optimal Inventory Policy The policy functions can then be represented by the (s_j, S_j) form, such that the post-order inventory is

$$Q_j = \begin{cases} S_j(\Phi_{-j}) & \text{if } q_j \leq s_j(\Phi_{-j}), \\ q_j & \text{otherwise.} \end{cases} \quad (10)$$

We characterize them to show how voucher demand and cross-brand business stealing enter the retailer's marginal stocking incentives.

The store chooses S_j by comparing the wholesale cost of additional inventory with the gains in expected revenue and future profits. The revenue calculation includes sales lost by substitute brands. For the WIC brand, additional inventory also reduces the stockout penalty. At the target, neither stocking more nor stocking less would increase expected profits.

When the store orders, its value is the value of starting with inventory S_j , less the fixed and variable costs of the order. The reorder point s_j is the highest inventory level at which ordering is optimal. At s_j ,

$$V_j(S_j(\Phi_{-j}); \Phi_{-j}) - V_j(s_j(\Phi_{-j}); \Phi_{-j}) = K + w_j(S_j - s_j). \quad (11)$$

The WIC brand's inventory also serves voucher consumers, and the stocking requirement makes the store hold more buffer stock of it than it otherwise would. Let j be a non-WIC brand; through Φ_{-j} entering V_j , Q_{WIC} crowds out inventory of the non-WIC brands, because the store makes similar revenue either way, but stocking the non-WIC brands means paying for more ordering and wholesale costs today.

Big-Box Stores Motivated by the empirical evidence in Section 3, big-box stores are treated as passive retailers with complete product availability. Their consumers choose among all J brands in all big-box stores in their 3-digit ZIP code, according to the choice probabilities $d_j(p, \xi, a)$ from Section 4.1. We assume consumers do not substitute across store types and discuss this assumption in Section 5.

4.3 Manufacturers

Manufacturers first set national wholesale prices given existing WIC contracts. Subsequent auctions take these prices as given and allocate new contracts. Each decision is modeled as a static best response: wholesale pricing responds to the existing allocation, and bidding responds to posted prices and the downstream profits associated with each possible winner.

Wholesale Pricing Let m index a state or purchasing alliance³¹ of WIC in a given contracting period, and let k_m denote its current WIC contract holder. Manufacturer j has marginal production cost c_j per prepared ounce and no fixed production cost. Write $Y_{jm|WIC=k_m}^{\text{priv}}(w)$ for its expected private-market sales under wholesale prices w and allocation k_m , and Y_m^V for expected voucher sales. Expected orders equal expected sales in a stationary inventory equilibrium, so these quantities also determine manufacturers' expected revenue.

Each firm revises its national wholesale price w_j at regular intervals. We assume that firms set prices simultaneously, taking existing WIC contract allocations as given, and maximize their expected long-run profits. When manufacturers change wholesale prices, we assume that retail prices change proportionally in both store types. Existing contracts adjust rebates one for one with posted wholesale prices, so that manufacturers' revenue from WIC voucher redemptions is fixed to $b_{jm} = w_j - r_{jm}$, where r_{jm} is the contracted rebate.³² Holding expected voucher volume fixed in this pricing calculation, firm j solves

$$\max_{w_j} \sum_m (w_j - c_j) Y_{jm|WIC=k_m}^{\text{priv}}(w) + \sum_{m:k_m=j} (b_{jm} - c_j) Y_m^V. \quad (12)$$

The margin from voucher units is fixed. The first-order condition with respect to w_j thus cap-

³¹As explained in Section 2.2, some states run their WIC auctions together through a purchasing alliance.

³²See Code of Federal Regulations (2011, §246.16a(c)(6)(iv)).

tures the effect of a small price change on profit from the private market

$$\sum_m \left[Y_{jm|WIC=k_m}^{\text{priv}}(w) + (w_j - c_j) \frac{\partial Y_{jm|WIC=k_m}^{\text{priv}}(w)}{\partial w_j} \right] = 0. \quad (13)$$

Procurement design affects wholesale competition by changing how stores order, which directly affects their responsiveness to wholesale prices.

The WIC Auctions The three manufacturers compete for a single sole-source contract in each market, again indexing states or WIC's purchasing alliance. Suppressing the market subscript m for simplicity, the manufacturer of brand j bids a rebate $r_j \in [0, \infty)$, taking its national wholesale price w_j as given, and the manufacturer with the lowest net wholesale price $b_j \equiv w_j - r_j$ wins. Wholesale prices are common knowledge. The auction's information structure and voucher volume are specified as follows.

Assumption 4.1 (Common knowledge on costs). Each firm's marginal cost of production c_j is common knowledge to all bidders.

Assumption 4.2 (Publicly known downstream environment). The map from the wholesale prices w and the winner's identity to each firm's expected sales in the market is common knowledge.

Assumption 4.3 (Winner-independent voucher volume). Each market m 's expected voucher volume Y_m^V does not depend on the winner's identity.

This is an auction with identity-dependent externalities (Jehiel and Moldovanu, 2000; Kuehn, 2019), generated by the retailers' stocking response to WIC program's procurement rules. With Assumption 4.2, the externalities each winner imposes on losers are common knowledge among manufacturers.

Firm j 's profit, as implied by the downstream model, depends on whether it wins and, when it loses, on *which* rival holds the contract, since the downstream stocking equilibrium depends on the WIC brand:

$$\Pi_{jm}^{\text{win}}(r_j) = (w_j - c_j) Y_{jm}^1 - r_j Y_m^V, \quad (14)$$

$$\Pi_{jm}^{\text{lose}}(k) = (w_j - c_j) Y_{jm|WIC=k}^{\text{priv}}, \quad (15)$$

where $Y_{jm}^1 = Y_m^V + Y_{jm|WIC=j}^{\text{priv}}$ is j 's total sales when it holds the contract and $Y_{jm|WIC=k}^{\text{priv}}$ is its private market sales when rival k does.

Modeling the auctions accurately requires cost heterogeneity within a firm across markets. Observed bids and allocations vary across markets in ways that the modeled operating profits alone do not account for. Our wholesale pricing first-order condition does not allow for such heterogeneity, since each manufacturer sets one national price given a constant marginal cost c_j .³³ We therefore introduce a firm- and market-specific cost shock η_{jm} to the per-unit cost of supplying WIC vouchers, part of which is private information. η_{jm} is referred to as *the bidder's type*.

We use the following assumption to formally introduce the type parameter.

Assumption 4.4 (Bidding heterogeneity and private information). For the auction in market m , firm j has type η_{jm} which shifts its willingness to bid beyond production costs and modeled private market spillovers:

$$\eta_{jm} = \mu_j^\eta + \tilde{\eta}_{jm}, \quad \tilde{\eta}_{jm} \sim F_{\tilde{\eta}}, \quad (16)$$

μ_j^η is common knowledge among bidders. $\tilde{\eta}_{jm}$ is private information of bidder j , and its distribution $F_{\tilde{\eta}}$ is common knowledge among bidders.

μ_j^η can be interpreted as the *average type* that bidders know about each other. $\tilde{\eta}_{jm}$ is the only private information in the model.

For bidding, define the *type-adjusted WIC unit cost* as

$$\tilde{c}_{jm}^{\text{WIC}} = c_j + \eta_{jm}. \quad (17)$$

The adjustment can represent an omitted economic cost or the bidder's perception about the cost of supplying the contract. The type-adjusted profit from winning separates WIC and private-market business:

$$\Pi_{jm}^{\text{win}}(r_j) - \eta_{jm} Y_m^V = (w_j - r_j - \tilde{c}_{jm}^{\text{WIC}}) Y_m^V + (w_j - c_j) Y_{jm|WIC=j}^{\text{priv}}. \quad (18)$$

The losing payoff is unchanged by the type.

We consider a Bayesian Nash equilibrium in which net bids are weakly increasing in types. Appendix Section B.1 characterizes the equilibrium strategies. Their interior optimality conditions are described as follows.

³³We maintain the assumption of constant marginal cost of production because one or two factories supply formula to the entire country, manufacturers generally do not pay distribution costs, and we find little evidence of economies of scale in production (Section 3).

Let G_k denote rival k 's equilibrium net-bid distribution, conditional on the public market environment, with density g_k on its continuous part. At an interior bid b , firm j wins with probability $\bar{G}_j(b) \equiv \prod_{k \neq j} (1 - G_k(b))$. Write $\bar{g}_j(b) \equiv -\bar{G}_j'(b)$; we then define the *externality weights* as

$$\lambda_{k|j}(b) \equiv \frac{h_k(b)}{\sum_{n \neq j} h_n(b)}, \quad h_k \equiv \frac{g_k}{1 - G_k}. \quad (19)$$

These weights describe which rival wins in the events where an arbitrarily small increase in j 's bid changes j from winner to loser. They condition on the lowest bid of the competitors being at b . The relevant private market spillover for the bidding first-order condition is therefore an externality-weighted average of the potentially different payoffs for losing,

$$\Delta \bar{Y}_{jm}^{\text{priv}}(b) \equiv Y_{jm|WIC=j}^{\text{priv}} - \sum_{k \neq j} \lambda_{k|j}(b) Y_{jm|WIC=k}^{\text{priv}}. \quad (20)$$

For ease of exposition, define the *pseudo-cost* for supplying each voucher unit as

$$\tilde{c}_{jm}(b) = \tilde{c}_{jm}^{\text{WIC}} - (w_j - c_j) \frac{\Delta \bar{Y}_{jm}^{\text{priv}}(b)}{Y_m^V}. \quad (21)$$

At a net bid equal to the pseudo-cost, the type-adjusted profit from winning equals the externality-weighted profit from losing.

At an interior bid ($b_{jm} < w_j$), the equilibrium bids satisfy

$$b_{jm} - \tilde{c}_{jm}(b_{jm}) = \frac{\bar{G}_j(b_{jm})}{\bar{g}_j(b_{jm})}. \quad (22)$$

The nonnegative-rebate constraint implies that bidders could submit $b_{jm} = w_j$. Such a bid is optimal only if no feasible positive rebate yields a higher expected payoff. Section 5.3 explains how we use interior and boundary bids in estimation.

4.4 Industry Equilibrium

The timing of the model is as follows:

1. Manufacturers set wholesale prices given the existing WIC contracts.
2. Retailers optimize inventory given WIC contracts, procurement rules, and wholesale prices.

3. Consumers buy from retailers given the products retailers stock.
4. In the next WIC auction, manufacturers take their own and competitors' wholesale prices as given, and bid for contracts according to their types.

The equilibrium conditions for each stage are:

1. **Retailer subgame.** Given w and the WIC contract holder k , each small-box store's stocking policy is the optimal (s_j, S_j) rule from the rational-expectations dynamic programming problem of Section 4.2.
2. **Wholesale pricing.** Given the existing WIC contracts, each manufacturer's wholesale price w_j solves the profit-maximization problem (12).
3. **Contract allocation.** Given the resulting wholesale prices w , the WIC contract in each market is allocated by the auction of Section 4.3. Firms bid in a Bayesian Nash equilibrium given the publicly known externalities, and the contract goes to the bidder who submits the lowest net wholesale price.

4.5 Discussion: How Procurement Design Shapes the Private Market

The equilibrium conditions establish the relationship between WIC's procurement design and private market outcomes through three channels.

Consumers' choice sets Retailers order jointly for WIC and private-market consumers, and their ordering determines private-market consumers' choice sets. The stocking mandate raises the retailer's target inventory for the WIC brand. Because brands are substitutes and replenishing each brand entails an ordering cost, this can reduce the retailer's incentive to replenish non-WIC brands, shrinking the variety available to private-market consumers. At the same time, the WIC brand's buffer stock is also available to private-market consumers, so they are less likely to find no brand in stock.

Rebates and the value of the contract Sole-source procurement makes manufacturers compete for access to voucher demand and the associated private-market spillovers. (22) implies that the rebate per voucher unit rises with $(w_j - c_j)\Delta\bar{Y}_{jm}^{\text{priv}}/Y_m^V$, the private-market spillover profit per unit. The amount captured by the government through a procurement mechanism depends on competition and bidding behavior.

Wholesale competition Procurement design changes retailers' stocking policies and therefore the response of retailer orders to wholesale prices. These changes pass upstream to manufacturers' equilibrium wholesale prices.

The tradeoff between fiscal cost and private-market consumer surplus in WIC's procurement design is determined by the importance of these three channels in equilibrium. Section 7 compares the current design and a wide range of reasonable alternative counterfactuals that vary them.

5 Empirical Model

In this section, we specify the empirical model for consumer demand, retailer stocking, manufacturer pricing, and bidding in WIC auctions. The estimation procedure is described along with model specifications.

This section proceeds in three steps. Step 1 specifies consumer demand and estimates it from big-box stores in the Scanner Data, where all brands are available. Step 2 specifies the retailer's empirical inventory model, holds the Step 1 preferences fixed, and estimates the remaining demand parameters and retailer costs from small-box orders, wholesale prices, and sales in the Inventory Data. It then re-estimates the store intercept for each small-box store in the Scanner Data. Step 3 holds the downstream model fixed and recovers manufacturers' production costs from wholesale pricing and the distribution of bidder types from bids.

5.1 Step 1: Demand

A random-coefficient logit model (Berry et al., 1995) is used to generate the choice probabilities for brand j under demand shock draw ξ under choice set a , $d_j(p, \xi, a)$ (first appeared in Section 4.1). A private market consumer i visiting store s in period t has indirect utility

$$u_{ijst} = -\alpha_i p_{jst} + \beta_{\text{WIC}} \text{WIC}_{jst} + \delta_s + \bar{\xi}_j + \xi_{jst} + \varepsilon_{ijst} \quad (23)$$

from brand j , and chooses the available brand $j \in a_i$ that yields the highest utility, or the outside option $j = 0$.³⁴ Here p_{jst} is the retail price, WIC_{jst} indicates j being the WIC brand, δ_s is a store-level intercept common to all inside goods, $\bar{\xi}_j$ is a brand fixed effect, ξ_{jst} is a store-, brand-, and period-specific demand shock, and ε_{ijst} is an i.i.d. Type-I extreme-value

³⁴Voucher (WIC) consumers do not make this discrete choice. They redeem their voucher for the contracted brand, so their demand is price-inelastic and enters directly as voucher sales Y_j^V in Section 4.2.

(T1EV) taste shock. Price sensitivity varies across consumers only through household income, entering as $D_i = \log(\text{income}_i/10,000)$: $\alpha_i = \bar{\alpha} + \alpha_D D_i$ (Nevo, 2001). The store-level intercept δ_s measures the appeal of buying formula at store s relative to the outside option, which has utility normalized to zero.

The following assumptions are made about consumers and their choice sets in order to estimate demand using sales in big-box stores:

Assumption 5.1. A consumer i is either a big-box store consumer or a small-box store consumer.

Assumption 5.2. A big-box store consumer's choice set is all J brands in all big-box stores in the 3-digit ZIP code (ZIP3) where consumer i lives.

Under these assumptions we estimate the demand model on big-box stores in the Scanner Data at monthly frequency, so t indexes months in this step. Assumption 5.2 is what makes the big-box data informative about preferences: with the complete set of brands always available, our estimates do not confound a brand's low availability with low demand.³⁵

In estimation, income are drawn from household income bins supplied by the American Community Survey, aggregated to the ZIP3 level. Prices and WIC status potentially suffer from endogeneity concerns since they are set in equilibrium as a function of demand shocks. Retail prices are instrumented with the national milk futures price, a shifter of formula input cost, using its level and its interaction with the ZIP3's average log household income. Milk prices shift the wholesale cost the big-box retailer pays, and hence the retail price, but are uncorrelated with the focal market's demand shocks ξ_{jst} . The WIC status is not instrumented in the main specification.³⁶

From the estimates, the preference parameters we treat as common to all consumers are recovered: the price coefficients $\bar{\alpha}$ and α_D , the WIC coefficient β_{WIC} , and the brand fixed effects $\bar{\xi}_j$. For big-box stores, ξ_{jst} is recovered as a structural residual and δ_s as a store fixed effect. The store-level intercept δ_s and the distribution of demand shock ξ_{jst} are not transferred to small-box stores and instead are estimated jointly with the inventory model in Step 2.

³⁵The Scanner Data are weekly; we aggregate to monthly because weekly brand-level sales could be zero even if the brand is in-stock, simply because not enough consumers arrived to generate enough positive purchases. Assumption 5.1 additionally implies that big-box demand does not respond to small-box prices or availability, so we need not model substitution across store types when estimating it; such substitution is empirically rare, since NielsenIQ Consumer Panel households are mostly either big-box or small-box shoppers (Appendix Figure A.8 shows the bimodal distribution of households' big-box purchase shares).

³⁶In a plain logit specification, we instrumented for the WIC indicator with the same instruments, and the estimate of β_{WIC} does not differ substantially from the one reported in Section 6.

5.2 Step 2: Retailer Costs

This step contains the bulk of our estimation procedure. The procedure combines the preferences from Step 1 with daily orders, wholesale prices, and sales in the Inventory Data to estimate the small-box store’s daily inventory model for each store in the Inventory Data. We hold fixed the Step 1 price coefficients $\bar{\alpha}$ and α_D , the WIC coefficient β_{WIC} , and the brand fixed effects $\bar{\xi}_j$, and then estimate the remaining demand parameters jointly with retailer costs. The remaining demand parameters are the coefficients γ_0 of the store intercept δ_s , the scale of demand shocks σ_{ξ_j} , and the private-market arrival rate ρ . The cost parameters are the coefficients γ_K of the fixed ordering cost K_s , the brand markups κ_j , the dispersion $\sigma_{\bar{\omega}}$ of persistent wholesale cost shocks, the scale σ_v of ordering shocks, and the perceived WIC stocking penalty λ_{WIC} . For stores in the Scanner Data, we hold all of these fixed except δ_s , which are re-estimated store by store (Section 5.2.5).

5.2.1 Assumptions

The following assumptions are made demand and costs in the small-box stores:

Assumption 5.3. A small-box store consumer with income D_i has the same α_i , β_{WIC} , and $\bar{\xi}_j$ as a big-box store consumer with the same value of D_i .

Assumption 5.4. All small-box stores have the same inventory cost function of store observables, and their private-market demand shocks ξ_{jst} are drawn from a common distribution across stores.

Practicalities around the available data and our model structure motivate us to make such assumptions. In principle, one could estimate all the transferred demand parameters jointly with stocking costs using the Scanner and Inventory Data stores, leveraging the identification argument in Agarwal and Somaini (2022) that stocking data and retail channel classification provide exogenous shifters in choice sets. This multi-step procedure is chosen mainly for computational tractability—it allows us to separate the estimation of a random-coefficient logit model, which is standard, and the stocking model, which is computationally burdensome. We consider Assumption 5.3 plausible because big-box and small-box shoppers who live in the same ZIP code are similar in race, ethnicity, income, and employment in the Consumer Panel, and a joint test does not reject balance across these characteristics (Appendix Table A.1). Small-box shoppers may still differ in the appeal of the outside option and the scale of demand shocks, since δ_s and σ_{ξ_j} are estimated separately for small-box stores. Assumption 5.4 effectively assumes that conditional on observables, the differences in stores’ stocking policies are driven by underlying

demand heterogeneity rather than cost heterogeneity; it is considered plausible because we observe differences in inventory in stores that are otherwise similar during fieldwork visits to retailers.

Under these assumptions, we transfer the preference coefficients conditional on income and integrate over each store’s local income distribution. Below we discuss the role of additional demand-side parameters in small-box stores that are estimated jointly with the inventory model.

5.2.2 Remaining Demand Parameters

Demand Shocks and the Outside Option In the Inventory Data, we specify $\delta_s = \gamma'_0 X_s^0$, where X_s^0 contains ZIP-code demographics and chain indicators. The demand shocks are specified as $\xi_{jst} = \sigma_{\xi_j} z_{jst}$, with z_{jst} i.i.d. standard normal across brands, stores, and periods.

Consumer Arrival, Transaction Sizes, and Stockouts An inventory-model period is one day, indexed by t . After the store restocks, each of N_s potential private-market consumers independently arrives with probability ρ , then chooses from the available brands and the outside option.

Voucher demand and transaction sizes are calibrated outside the estimation. Voucher demand ξ^V is the purchases of N_s^V potential voucher consumers; each arrives with daily probability ρ^V and buys $\bar{y}^V$ ounces. Private-market consumers buy $\bar{y}$ ounces per transaction. ρ^V is calibrated from the frequency of WIC households’ formula purchases in the Consumer Panel, and $\bar{y}$ and $\bar{y}^V$ from average transaction sizes. N_s is set in proportion to the upper tail of each store’s daily private-market sales. N_s^V is calibrated so that expected daily voucher demand, $N_s^V \rho^V \bar{y}^V$, equals the store’s mean daily voucher sales, assuming that mean voucher sales equal mean voucher demand (Appendix Section C.2).³⁷

Calculating expected revenue over consumer arrivals, stockouts, and demand shocks is burdensome. We therefore make some simplifying approximations. First, we approximate the binomial arrival count by $N \sim \text{Poisson}(N_s \rho)$, where $N_s \rho$ is expected number of daily arrivals. Then, subsequent substitution is approximated by reallocating unsatisfied demand among the remaining brands and the outside option, again using a Poisson distribution. Appendix Section C.5 describes the algorithm in detail. In simulations, our approximation performs well for

³⁷Given the demand system, the brand shares of private-market sales provide information about which brands were stocked out. Voucher sales have no such structure, so we cannot tell how many voucher consumers a stockout turns away.

the range of N_s present in the data compared with the discrete stockout process used by Conlon and Mortimer (2013).

5.2.3 Cost Specifications

As modeled in (4), a store's retail price for brand j is $p_{js} = m_{js} w_j$. Each brand's retail price is fixed at its store-level average list price over the sample period, excluding promotional discounts, and set m_{js} to the ratio of this price to w_j . m_{js} is held fixed in counterfactuals.³⁸

Availability is modeled at the brand level, while observed average wholesale prices of a given day also reflect the varieties that the store restocked at that day. We therefore specify

$$\log w_{jst} = \log w_j + \log m_{js} - \kappa_j + \omega_{jst}, \quad \omega_{jst} = \tilde{\omega}_{jst} + e_{jst}, \quad (24)$$

where w_j is the national wholesale price.³⁹ The deterministic term $\log m_{js} - \kappa_j$ maps the posted price to the store's wholesale price. Here m_{js} summarizes the store's usual composition of UPCs within the brand, and κ_j is the difference between brand j 's log retail price and the store's log wholesale cost. Because this construction uses only the store's observed retail prices, Assumption 5.4 extends it to the Scanner Data stores, whose wholesale prices we never observe.

ω_{jst} is specified as containing persistent and transitory variation in a brand's average wholesale cost per ounce at a given day. The persistent component $\tilde{\omega}_{jst}$ is observed before ordering, so the store's state for brand j is $(q_j, \tilde{\omega}_{jst})$, and it pays w_{jst} per unit in place of w_j in C_j as written in Section 4. Each day, with probability ζ_w , it updates according to

$$\tilde{\omega}' = \rho_w \tilde{\omega} + \sqrt{1 - \rho_w^2} \sigma_{\tilde{\omega}} z^{\omega}, \quad z^{\omega} \sim N(0, 1), \quad (25)$$

and otherwise remains unchanged. The parameter ρ_w governs persistence across these updates.⁴⁰ The transitory component $e_{jst} \sim N(0, \sigma_e^2)$ captures variation in delivered costs around this persistent state. We calibrate ζ_w , ρ_w , and σ_e (Appendix Section C.3) and estimate $\sigma_{\tilde{\omega}}$ jointly with the other parameters.

The fixed ordering cost is specified as $K_s = \exp\{X_s^K \gamma_K\}$, where X_s^K includes a constant and log average daily formula volume in the observed data. Because we lack direct data on stores'

³⁸Daily variation in a brand's average retail price within a store mostly reflects which UPCs sell that day. UPC identity explains 99% of this variation. A model of the price variation would not enter the store's state, because the store cannot anticipate which UPCs consumers buy, so the structure of the stocking policy would not change.

³⁹We use the national wholesale price submitted in the WIC contract data and hold it fixed for the duration of the contract.

⁴⁰Scaling the innovation variance by $(1 - \rho_w^2)$ gives the level a stationary distribution $N(0, \sigma_{\tilde{\omega}}^2)$, so $\sigma_{\tilde{\omega}}$ measures the dispersion of the persistent component faced when ordering.

inventory infrastructure, which determines their fixed ordering costs, baseline sales volume is used as a proxy for shelf space. The estimated costs are held fixed in counterfactuals, assuming that the underlying structure remains unchanged. The perceived WIC stocking penalty λ_{WIC} in (8) is also estimated. In addition, independent T1EV payoff shocks to ordering and waiting are included, with scale σ_v , allowing variation in orders conditional on the state $(q_j, \tilde{\omega}_{jst})$. The ordering probability can then be written in closed form, which we specify in Appendix Section C.4.

We solve this empirical model and derive the stationary stocking policy for each store for the duration of the WIC contract. We interpret the inventory policy as a description of a store's aggregate restocking behavior of a brand, which is revised intermittently, as the state $\tilde{\omega}$ changes. It can arise from, for example, a store manager or an automated inventory management system establishing a stocking rule for the underlying varieties of a brand over an extended period. Our estimated inventory cost parameters help us recover the long-run distribution of consumers' choice sets. The empirical model inherits the base model's structure, in which stores endogenously respond to WIC's procurement rule, captured by the perceived stockout penalty for the WIC brand.

Lastly, the daily discount factor is calibrated as $\beta = 0.95^{1/30}$, which discounts by 5 percent per month.

5.2.4 Moments and Identifying Variation

The remaining demand and cost parameters are estimated jointly by method of simulated moments, matching each store's orders, wholesale prices, and sales. At each candidate parameter vector, we solve each store's inventory problem and find the best parameters that match four classes of moments in each store: (1) order frequencies and sizes of orders; (2) mean and variances of sales, as well as zero sales frequencies; (3) log wholesale prices, variances, and autocorrelations; (4) WIC brand's private market sales and its private market shares. Appendix Section C.4 lists these four classes of moments and explains their construction.

The sources of identifying variations rely on the comparative statics of the stocking model. Given price sensitivities and brand preferences estimated in Step 1, order frequencies and the average size of orders are jointly informative about inventory costs and the remaining demand parameters. Fixing demand and wholesale costs, infrequent large orders indicate a greater incentive to economize on larger ordering costs. Observed wholesale prices and their variances pin down the distribution of its persistent component $\tilde{\omega}$. The WIC brand's private-market sales and market share are informative about the program's stocking penalty once the the WIC coefficient β_{WIC} in demand and stocking costs are accounted for. The scale σ_v is informed by how

often stores order, or skip ordering, at inventory levels where the policy without shocks would do the opposite. Heterogeneity in moments by store pins down (1) the heterogeneity in δ_s and (2) the ordering cost function. In addition, variation in sales and frequency of zero sales are decomposed into randomness induced by consumer arrival and randomness induced by demand shocks. Since the private-market consumer arrival rate ρ are assumed to be the same across stores, stores with larger calibrated N_s would have less lumpy sales. The variation of N_s thus helps pin down ρ and variance of ξ jointly.

5.2.5 Demand and Costs for Scanner Data Stores

The manufacturer model and national counterfactuals require primitives for stores in the geographically broad Scanner Data sample, whose inventory orders are unobserved. Assumptions 5.3 and 5.4 let us apply the estimated cost functions and common demand parameters to these stores. The remaining unknown is δ_s , which determines consumer substitution to buying elsewhere or not buying in the event of a stockout. At a given cost, δ_s determines how much revenue the retailer can capture when some brands stock out.

Holding the transferred parameters fixed, we estimate each store's δ_s again using the method of simulated moments. We find the δ_s that best matches each store's mean sales, zero-sales frequencies, and the WIC brand's private market sales and its private-market shares. At each candidate δ_s , we solve its inventory problem and aggregate simulated daily sales to the weekly frequency of the Scanner Data. This yields a separate δ_s and model-implied stocking policy for each store. The moments are those in Appendix Section C.4, excluding the order and wholesale-price moments, and Appendix Section C.2.2 describes how we calibrate N_s from weekly sales.

5.3 Step 3: Manufacturer Costs and Bidding Behavior

Using the estimated downstream model, production costs are recovered from the manufacturers' wholesale-pricing condition (13). Given these costs and profits under each WIC allocation, we estimate the type distribution governing firms' bidding behavior.

Production Costs We use the contracted product's national wholesale price per prepared ounce, w_j , to approximate average manufacturer revenue per prepared ounce across the brand's products, before WIC rebates. We recover costs from the elasticity of aggregate private-market orders with respect to the national wholesale price, ϵ_j^Q .⁴¹ Following the proportional retail-price

⁴¹The Scanner Data contains a nonrandom sample of stores in the country. Appendix Section C.6 describes aggregation to auction markets and the sampling weights used to project Scanner Data stores to the retailer

response in Section 4.3, retail prices are varied in both store types when calculating this elasticity; small-box wholesale costs also change, affecting stocking. Because rebate adjustments keep revenue from voucher units fixed, only private-market quantities enter the pricing optimality condition. This gives the cost estimate $c_j = \theta_j w_j$, where $\theta_j = 1 + 1/\epsilon_j^Q$. The elasticity is estimated using the cross-section of contracts in effect in 2016 and hold the implied cost share θ_j fixed across auction years, allowing cost levels to move with posted wholesale prices.⁴²

Bidding Behavior Bids reflect profits under each WIC allocation and the bidder’s type η_{jm} in (16). Its mean is written as $\mu_{j,\ell}^\eta$ and the private component as $\tilde{\eta}_{jm} \sim N(0, \sigma_{\eta,j,\ell}^2)$, where ℓ denotes period 2006–2009 or 2010–2015. μ^η and σ_η are allowed vary by brand and period. The period definition is motivated by the observation that Nestlé stopped submitting many zero-rebate bids starting in 2010.

These parameters are estimated in two steps. First, we fit a lognormal distribution for net bids, censored at w_j , with a point mass at w_j corresponding to zero-rebate bids. Mean of log bid is specified as a function of firm j ’s gain from winning, the spread of its payoffs across rival winners, its posted wholesale price, and voucher volume ((C.8)), informed by the characterized equilibrium bid distribution in Appendix Section B.1; we specify the variance of log bid as brand- and period-specific. Then, using the fitted bid distributions and production costs, we invert (22) to recover η_{jm} . A zero-rebate bid instead reveals that the draw of η_{jm} must be high enough that no feasible positive rebate improves the bidder’s expected payoff relative to submitting a zero-rebate bid. This cutoff is calculated by doing a grid search on feasible deviations, accounting for which rival would otherwise win. We then estimate $(\mu_{j,\ell}^\eta, \sigma_{\eta,j,\ell})$ by maximum likelihood, using the recovered interior values and these bounds. Appendix Section C.6 contains details about the cutoff calculation and likelihoods.

6 Estimation Results

6.1 Demand

Our estimated demand parameters are presented in Table 1. The parameters are interpreted in the context of prices in 2016, the year that we simulate the counterfactuals for. The average

population.

⁴²Milk is an input to formula production, and we use CME Class III milk futures prices (CME Group, n.d.) as an input-cost shifter in demand estimation. Co-movement between input costs and posted wholesale prices motivates allowing production costs to vary over time; the constant cost share is an additionally maintained assumption.

retail price in big-box stores is \$0.21 per prepared ounce.⁴³ The estimated choice probabilities imply that, averaged across the income distribution, households behave as if they are willing to pay \$0.034 per prepared ounce more for the WIC contract brand, about 17 percent of price. Holding prices and availability fixed, β_{WIC} alone implies that a brand gaining the WIC contract from a rival raises its private-market share by about 8 percentage points, close to the private-market spillover estimated in big-box stores (Figure 2). Relative to Abbott, consumers are willing to pay \$0.080 per prepared ounce more for Mead Johnson (39 percent of price) and \$0.126 per prepared ounce less for Nestlé (61 percent of price).

6.2 Inventory Costs

Table 2 shows our estimates of inventory costs in small-box stores. The ordering cost for a brand is estimated to have a median of 6.25 dollars, with an interquartile range of 3.77–9.82 dollars within the Inventory Data. A brand order in our data includes 1.92 UPCs on average, so the median cost per UPC is about 3.3 dollars, comparable to the median ordering cost of about 2 dollars per product that Aguirregabiria and Guiton (2023) estimate for liquor stores in Ontario.

We estimate that the perceived penalty for 1 percentage point of *ex ante* daily stockout probability of the WIC brand is 40.13 dollars, so a day on which the WIC brand stocks out with certainty carries a perceived cost of about 4,000 dollars. For scale, the civil money penalty the WIC program can impose for a violation, which is about 3,200 dollars for the median store in our sample (interquartile range 1,900–5,100 dollars).⁴⁴ These estimates indicate that retailers place a high value on complying with the WIC program’s minimum stocking requirement.

To assess model fit, each brand’s predicted mean weekly sales are computed in the Scanner Data stores using the full model and are plotted against observed mean weekly sales (Figure 5). Each point is a brand in a WIC auction market. Because we target each store’s mean sales when estimating δ_s (Section 5.2.5), the figure shows the in-sample fit of these moments after aggregating stores to auction markets. The points lie close to the 45-degree line.

⁴³For reference, the WIC program’s monthly benefit amount for an infant up to 3 months old is 806 prepared ounces, each 13-oz can (the smallest powder package) is 90 prepared ounces, and a medium-sized tub (28-oz) is 193.8 prepared ounces.

⁴⁴A state agency can disqualify a vendor that violates its requirements for up to one year. If disqualification would impair participants’ access, the agency can instead impose a civil money penalty of 10 percent of the vendor’s average monthly WIC redemptions for each month of disqualification; see Code of Federal Regulations (n.d.b, §246.12(l)(1)(x) and (l)(2)(i)). We compute 10 percent of a year’s WIC formula redemptions, rounded to the nearest hundred dollars. This is a lower bound on the penalty because it excludes other WIC foods.

Table 1: Demand Estimates

(a) Big-Box: Consumer Preferences

	Estimate
<i>Price coefficient</i> $-\alpha_i = -(\bar{\alpha} + \alpha_D \log(\text{income}_i / \$10k))$	
Constant $-\bar{\alpha}$	-17.317 (0.317)
Log income $-\alpha_D$	2.347 (0.093)
<i>Product Characteristics</i>	
WIC brand (β_{WIC})	0.439 (0.002)
Mead Johnson ($\bar{\xi}_{\text{MJ}}$)	1.038 (0.009)
Nestlé ($\bar{\xi}_{\text{N}}$)	-1.628 (0.004)
<i>Implied (at the estimate)</i>	
Mean own-price elasticity	-2.113
Abbott / Mead Johnson / Nestlé	-1.943 / -2.572 / -1.825
Store×brand×month obs. (N)	664,849
Markets (ZIP3×month, T)	68,935

(b) Small-Box: Demand Shocks and Consumer Arrivals

	Point Estimate
SD of the private demand shocks $\sigma_{\xi j}$	
Abbott	2.44
Mead Johnson	2.95
Nestlé	1.59
Consumer arrival	
Expected daily arrivals / peak daily transactions	0.407
Voucher daily arrival probability ρ^{V*}	0.024
Private mkt. transaction size (prepared oz.)*	231.4
Voucher transaction size (prepared oz.)*	387


Notes: Standard errors in parentheses. *Panel (a):* Prices are instrumented with the Class III milk futures price and its interaction with the market's average log income. The sample excludes 2007–2009, when the milk futures prices are unreliable cost shifters because they are subject to the general conditions of the financial markets during the Great Recession. Store fixed effects are absorbed throughout, and Abbott is the omitted base brand. *Panel (b):* estimated jointly with the retailer costs in the Inventory Data (Table 2). Starred parameters are calibrated offline; expected daily arrivals are reported relative to a store's peak daily transactions, $P_s/\bar{y}$, because the level of N_s is a normalization (see Appendix Section C.2 for more details).

Table 2: Inventory Costs and WIC Stockout Penalties

Parameter	Description	Point Estimate
<i>Wholesale costs</i>		
κ_1	Abbott markup over wholesale (log)	0.293
κ_2	Mead Johnson	0.237
κ_3	Nestlé	0.250
$\sigma_{\tilde{\omega}}$	SD of the persistent price shock	0.396
σ_e	SD of transitory wholesale cost shock*	0.179
ζ_w	Daily probability of a $\tilde{\omega}$ update*	1/30
ρ_w	Persistence across $\tilde{\omega}$ updates*	0.943
<i>Ordering costs</i>		
$\gamma_{K,0}$	Constant	2.09
$\gamma_{K,1}$	Log store volume	1.02
	Implied K_s at the p25/p50/p75 store	3.77 / 6.25 / 9.82
σ_v	Scale of the ordering-action shock	29.8
λ_{WIC}	Penalty \$ per 1 p.p. stockout of WIC brand	40.13
β	Discount factor	$0.95^{1/30}$
# Small-box stores		544
# Store-days		515,703

Notes: Wholesale costs follow $\log w_{jst} = \log w_j + \log m_{js} - \kappa_j + \omega_{jst}$, where $\omega_{jst} = \tilde{\omega}_{jst} + e_{jst}$. Ordering costs satisfy $K_s = \exp\{\gamma_{K,0} + \gamma_{K,1} \log \text{store volume}_s\}$. Starred parameters are calibrated (see Appendix Section C.3 for more details). Store-level intercept coefficients are estimated jointly and reported in Appendix Table C.2, and the small-box demand parameters estimated in the same step are reported in Panel (b) of Table 1.

Figure 5: Model Fit for Brand-Level Sales



Notes: Each point is a brand's average weekly sales (voucher and private combined, in thousands of prepared ounces) in a WIC auction market: the model's prediction on the horizontal axis and the data on the vertical axis, both on log scales. Colors denote brands, and marker shapes denote the WIC contract brand (filled) versus the non-WIC brands (crosses). The dashed line is the 45-degree line of perfect fit.

6.3 Manufacturer Costs and Bidding Behavior

Each brand's production cost is recovered from its aggregate private-market quantity elasticity. The estimated cost share $\hat{\theta}_j = c_j/w_j$ implies that production costs are about 19 percent of wholesale price for Abbott, 17 percent for Mead Johnson, and 35 percent for Nestlé. Per 90 prepared ounces, production cost rises from \$2.09 to \$2.80 for Abbott, \$2.02 to \$2.61 for Mead Johnson, and \$3.96 to \$5.19 for Nestlé between 2006 and 2016 as posted wholesale prices rise. Figure 6 compares these costs with retail prices, posted wholesale prices, and bids. Across the 54 auctions in our estimation sample, most bids lie below estimated production costs, consistent with manufacturers accepting losses on WIC purchases in exchange for more sales in the private market.

We define the effective cost in bidding as $c_j + \mu_{j,\ell}^\eta$, production cost plus the mean type, as shown in Figure 6. Abbott's mean type is \$6.13 per 90 voucher prepared ounces in 2006–2009 and \$5.57 in 2010–2016; Mead Johnson's is \$7.39 and \$7.49. Nestlé's mean type changes from \$6.15 to $-\$2.20$. Thus, conditional on the production costs and spillovers implied by the demand and stocking model, Abbott and Mead Johnson offer lower rebates on average than bidding based on modeled profits alone would imply, while Nestlé's bidding becomes more aggressive between the two periods. Appendix Table C.4 reports estimates of the type distribution.

We interpret the large positive means primarily as systematic behavioral departures rather than invariant opportunity costs of supplying WIC. Three observations support this interpretation. First, winning bids are more aggressive after 2016, often below zero.⁴⁵ Second, the production costs recovered from wholesale-pricing first-order conditions, particularly for Abbott and Mead Johnson, are close to the cost levels suggested by the UK Competition and Markets Authority's manufacturer-accounting evidence (Competition and Markets Authority, 2025).⁴⁶ Attributing these positive mean types to production costs would imply much higher costs for Abbott and Mead Johnson. Third, the public investment disclosures we examined do not establish a contemporaneous expansion in production capacity for milk-powder products that would explain the shift toward lower contract prices as a reduction in the opportunity cost of serving WIC.⁴⁷


⁴⁵Calculated from USDA administrative data on winning bids: net prices are positive in all recorded auctions for contracts starting in 2006–2016. Of nine contracts starting in 2016, four have a zero net price and one a negative net price; 13 of 20 contracts starting in 2017–2020 have negative net prices.

⁴⁶The CMA reports aggregate manufacturer variable gross margins of 50–75 percent. An illustrative conversion using a £12 price for an 800g pack, a 20 percent retail margin, and the midpoint manufacturer margin of 62.5 percent gives a variable cost of £3.60 per pack. Using approximately 180 prepared ounces per pack and \$1.27 per pound sterling gives about \$2.29 per 90 prepared ounces.

⁴⁷Mead Johnson's 2016 Form 10-K attributes its 2015–2016 capital expenditures primarily to products for China and US liquid offerings (Mead Johnson Nutrition Company, 2017). MJ's 2019 investment at its Zeeland

Together, these observations favor a behavioral interpretation without identifying the source of the wedge. For example, mistaken strategic beliefs or bidding heuristics can be represented by departures from the modeled operating profit benchmark, even though the formal auction model rationalizes bids as a Bayesian Nash equilibrium in effective payoffs. We therefore remove the mean in the preferred counterfactual specification, while keeping the idiosyncratic types that generate heterogeneity in bidding. In the robustness exercise in Appendix Section D.3, we interpret the mean as an opportunity cost of supplying the contract and explore implications of this interpretation.

Figure 6: Production Costs, Effective Costs in Bidding, Prices, and Bids



Notes: Each panel plots, by WIC contract start date over 2006–2016, the recovered production cost $c_j = \theta_j w_j$, the effective cost in bidding $c_j + \mu_{j,\ell}^\eta$, the posted national wholesale price w_j , the brand’s net-of-rebate bids $b_{jm} = w_j - r_{jm}$, and observed retail prices (volume-weighted big-box and small-box combined) for each auction, all in dollars per 13-oz can (90 prepared ounces). The gap between the green and black lines is the mean type $\mu_{j,\ell}^\eta$, estimated separately for 2006–2009 and 2010–2016. The green line includes production cost but does not subtract modeled private-market spillovers or include the idiosyncratic bidding shock; these enter the full pseudo-cost in (21). Zero-rebate bids appear on the wholesale price line.

factory target hypoallergenic formula and export-market requirements (Michigan Department of Agriculture and Rural Development, 2019). These disclosures do not establish an expansion in non-specialty milk-powder products studied here that explains the onset of more aggressive bidding.

7 Evaluating Alternative Designs of Procurement Policies

The modeled is used to ask how procurement can broaden private-market consumers' choices while preserving the manufacturer payments that finance WIC. In particular, we ask whether an alternative design can raise private-market consumer surplus relative to the current design without increasing net program spending or reducing the quantity of formula that beneficiaries redeem through WIC.

For the national cross-section of contracts active in 2016, the designs in Table 3 are evaluated. As discussed in Section 4.5, procurement design affects the private market through three channels. The stocking requirement changes retailers' ordering and therefore private-market consumers' choice sets. Restricting eligibility to contracted brands makes manufacturers compete for WIC demand and the associated private-market spillovers, which determines their payments to the government. Both change how retailers' orders respond to wholesale prices, and therefore manufacturers' equilibrium wholesale prices.

To examine retailers' stocking incentives, the stocking requirement is removed while holding the observed contracts and wholesale prices fixed. To examine manufacturer competition for contracts, we change the auction format and the number of contracted brands. Contracting with more brands also extends the stocking requirement to those brands, so these designs affect private-market consumers' choice sets as well as manufacturer payments. We then consider three other changes to the program. We give beneficiaries a dedicated product that private-market consumers cannot buy, to examine the efficiency gain of pooling inventory across the two groups. Free formula is replaced with a coupon of fixed value, so that beneficiaries pay the remaining retail price. Lastly, WIC is expanded to include 50% more than the baseline pool of beneficiaries, which makes the contract more valuable to manufacturers and leaves fewer consumers in the private market.

Across all counterfactual designs, the set of stores consumers consider and how frequently they shop are held fixed. Consumer populations are held fixed at each store for WIC participants and private-market consumers except when explicitly noted. Except under designs in which WIC issues coupons, formula remains free for beneficiaries. For each design, the contract allocation and manufacturer payments are determined at observed wholesale prices, accounting for retailers' stocking responses. Manufacturers then adjust wholesale prices once in response to retailer demand under that allocation. As assumed in Section 4.3, retail prices change in proportion to wholesale prices, so that each store's ratio of retail to wholesale price stays fixed. Stocking and purchases are evaluated at the adjusted prices, while the contracts and corresponding payment terms to the government are held fixed.

Figure 8 reports the share of small-box consumers' choice sets that have zero, one, two, or all three brands. Figure 10 shows the change in consumer surplus from each design at observed prices and the additional change when manufacturers adjust wholesale prices. Figure 9 decomposes fiscal costs into gross government expenditure and manufacturer payments for contracts. Data underlying these figures can be found in Appendix Table D.1. Appendix Section D describes the measurement procedure for outcomes in each counterfactual equilibrium.

Below we explain the implementation of how WIC beneficiaries make choices if there are multiple WIC brands, how WIC contracts are allocated in different auction formats, and how program spending and private-market consumer surplus are measured relative to the status quo, in which contracts and rebates are held at their observed values.

Beneficiary Choices Some designs allow beneficiaries to choose among WIC brands or whether to redeem their benefits. These include designs in which vouchers cover more than one brand (“multi-sourcing” designs) and coupon designs in which the voucher is replaced by a price subsidy. In either case, we apply the demand parameters in (23) estimated from private-market consumers to beneficiary choices. Suppressing store and time indices, let a_i^d denote the eligible products available to beneficiary i under design d . Vouchers and restricted coupons cover the contracted brands; brand-neutral coupons cover all brands. Let $p_j^{d,WIC}$ denote the beneficiary's out-of-pocket price for an eligible product. It is zero under voucher designs; under coupon designs, it is the retail price net of the coupon value, or zero when the coupon value exceeds the retail price. The WIC coefficient β_{WIC} in (23), which we interpret as a labeling effect, steers consumer choices to the contracted brand. When there are multiple contracted brands, β_{WIC} is divided equally across them when benefits are restricted and across the store's carried brands otherwise.

The beneficiaries' choice probability for each brand is as follows. Let v_{ij}^d be the utility index in (23), excluding ε_{ijst} and evaluated at $p_j^{d,WIC}$. Beneficiary i chooses product $j \in a_i^d$ with probability

$$\Pr_i^d(j | a_i^d) = \frac{\exp(v_{ij}^d)}{V_0 + \sum_{k \in a_i^d} \exp(v_{ik}^d)}.$$

When products are free for beneficiaries (under voucher designs), $V_0 = 0$. They choose among available eligible products without an outside option, as in the status quo; they do not redeem when no eligible products are available. When products require a copayment (under coupon designs), they may also choose not to buy, and $V_0 = 1$. We do not allow beneficiaries to purchase excluded brands at full price.⁴⁸

⁴⁸In small-box stores, we average beneficiary purchases over consumer arrivals and demand shocks, allowing

Contract Allocation and Payments In designs that restrict subsidy eligibility to a subset of brands, auctions are used to select the contracted manufacturers and determine their payments to the program. We consider a first-price auction for sole sourcing (having one WIC brand) and a Vickrey–Clarke–Groves (VCG) mechanism for both sole and dual sourcing (having two WIC brands). In the first-price auction, the manufacturer offering the lowest net wholesale price wins and pays its offered per-unit rebate.⁴⁹

For a contracted set of WIC brands a , let $Y_{jm}^V(a)$ and $Y_{jm}^{\text{priv}}(a)$ denote manufacturer j 's WIC and private-market sales. Using the type-adjusted WIC unit cost $\tilde{c}_{jm}^{\text{WIC}} = c_j + \eta_{jm}$ from (17), define the type-adjusted gross profit before auction payments as

$$\tilde{\Pi}_{jm}(a; \eta_{jm}) = \underbrace{(w_j - \tilde{c}_{jm}^{\text{WIC}})Y_{jm}^V(a)}_{\text{WIC gross profit}} + \underbrace{(w_j - c_j)Y_{jm}^{\text{priv}}(a)}_{\text{private-market gross profit}}. \quad (26)$$

Let $\mathcal{J}$ denote the set of J manufacturers. With n contracts, $A_m = \{a \subseteq \mathcal{J} : |a| = n\}$ contains all sets of manufacturers with cardinality n , where $n = 1$ under sole sourcing and $n = 2$ under dual sourcing. VCG selects the allocation that maximizes aggregate type-adjusted gross profit:

$$a_m^* \in \arg \max_{a \in A_m} \sum_j \tilde{\Pi}_{jm}(a; \eta_{jm}). \quad (27)$$

Let a_{jm}^0 solve the same problem subject to $j \notin a$. Manufacturer j pays the government

$$T_{jm} = \sum_{k \neq j} \left[\tilde{\Pi}_{km}(a_{jm}^0; \eta_{km}) - \tilde{\Pi}_{km}(a_m^*; \eta_{km}) \right]. \quad (28)$$

This payment is the other firms' aggregate profit loss relative to the best allocation that excludes j . Under this payment rule, losers always pay zero to the government. As in Assumption 4.2, expected sales under each set of contracted brands are assumed to be common knowledge, so each manufacturer's only private information is its type and truthful reporting is a dominant strategy.

Program Spending and Consumer Surplus We report annualized differences in net program spending and private-market consumer surplus relative to the status quo, which uses observed contract winners and rebates. Net program spending is gross government expenditure minus

availability to change as inventory is depleted. In big-box stores, all eligible brands are available; the recovered demand shocks ξ_{jst} are held fixed while integrating over consumer heterogeneity.

⁴⁹We re-solve the first-price auction in Appendix Section B.1 using the efficient grid search algorithm proposed by Kastl et al. (2025).

manufacturer payments. When calculating private-market consumer surplus, β_{WIC} is excluded, because the labeling effect does not generate consumption value from the product per se.⁵⁰ As specified in the description of how beneficiaries make choices under different WIC contract configurations, β_{WIC} continues to affect consumer purchases and, in turn, retailer and manufacturer behavior in equilibrium.

⁵⁰For design d relative to the status quo, denoted sq,

$$\Delta CS_i^d = \frac{\mathbb{E}[V_i^d(a_i^d) - V_i^{\text{sq}}(a_i^{\text{sq}})]}{|\alpha_i|},$$

where $V_i^d(a) = \log(1 + \sum_{j \in a} e^{\tilde{u}_{ij}^d})$ and $\tilde{u}_{ij}^d$ is utility excluding the logit shock ε_{ijst} and the allocated share of β_{WIC} . Here a_i^d is the choice set available to private-market consumer i under design d . The expectation is over demand shocks and the available choice sets generated by consumer arrivals and purchases under each design. This expression is in dollars per prepared ounce. We multiply by the transaction size $\bar{y}$ when aggregating across consumers.

Table 3: The Menu of Procurement Designs

Design	Brands bene. can buy	Beneficiary pays	Stockout penalty on	Contract award	Program pays	Δ Surplus (\$M)		Δ Fiscal (\$M)	
						Fixed	Adjusted	Fixed	Adjusted
Status quo	One	Zero	Winner	Observed first-price auction	Retail price – observed rebate	—	—	—	—
<i>Stocking mandate</i>									
Current allocation, no mandate	One	Zero	None	Observed winner	Retail price – observed rebate	–120	—	–75	—
<i>Allocation and payment rule</i>									
Sole-source (VCG)	One	Zero	Winner	VCG	Retail price – VCG payment	+70	+10	–322	–279
Profit-only bidding	One	Zero	Winner	First-price auction	Retail price – pay-as-bid rebate	+80	+20	–42	+6
<i>Number of contracted brands</i>									
Dual sourcing (VCG)	Two	Zero	Both winners	VCG over pairs	Retail price – VCG payments	+350	+340	–1	–1
Triple sourcing	All three	Zero	All brands	No auction	Retail price	+450	+540	+1,167	+1,123
<i>Distribution</i>									
Government brand	Dedicated WIC brand	Zero	Dedicated WIC brand	No auction	Production cost + retail margin + ordering costs	–110	+10	+220	+221
<i>Reimbursement</i>									
Brand-neutral coupon	All three	Price above credit	None	No auction	Coupon credit	–140	–50	+20	+31
One-brand coupon	One	Price above credit	Winner	VCG	Coupon credit – VCG payment	+90	–210	–1,028	–1,060
Two-brand coupon	Two	Price above credit	Both winners	VCG over pairs	Coupon credit – VCG payments	+350	+190	–711	–722
<i>Program scale</i>									
Sole-source, 1.5× WIC demand	One	Zero	Winner	VCG	Retail price – VCG payment	–210	–250	+36	+80

Notes: Each coupon provides a \$7.50 credit per 90 prepared ounces; beneficiaries pay any remaining price. The indicated brands carry a perceived stockout penalty in the retailer's inventory problem; availability depends on inventory choices and realized demand. Under first-price auctions, the manufacturer offering the lowest net wholesale price wins and pays its offered rebate. VCG awards and payments follow (27) and (28), respectively. Counterfactual auctions in this table set the mean of η to zero and include the idiosyncratic shocks. The government brand can be purchased only by WIC beneficiaries. Δ Surplus and Δ Fiscal report annual differences from the status quo in private-market consumer surplus and net program spending, respectively, in millions of dollars. Consumer-surplus differences exclude β_{WIC} and WIC beneficiaries. Fixed holds wholesale prices at their observed values; Adjusted allows one wholesale-price response with contract awards held fixed.

7.1 Retailer Stocking Incentives

The stocking requirement gives retailers a reason to hold inventory for WIC beneficiaries, but concentrating that inventory in one brand can discourage them from restocking substitutes. We first ask whether removing the requirement improves private consumers' choices. The requirement enters the retailer's problem as a perceived penalty on the probability that the WIC brand stocks out, with coefficient λ_{WIC} in (8). This penalty is set to zero while holding the observed contract winner, rebate, and wholesale prices fixed. Beneficiaries remain eligible to buy only the contracted brand at no out-of-pocket cost.

Removing the perceived stockout penalty for the WIC brand reduces WIC redemption by 21 percent and annual private-market consumer surplus by \$120 million. Without the stocking requirement, competing brands' private-market sales in small-box stores rise by 11 percent, but the WIC brand's own private-market sales fall by 15 percent. Stores have less incentive to hold buffer stock of the WIC brand, so less of its inventory is left for private-market consumers, including those whose preferred brand is out of stock. Total private-market sales fall by 6 percent, and the share of small-box private-market consumers who find no brand in stock rises from 1.6 to 4.0 percent. The requirement lowers the return to stocking competing brands, but the inventory it supports also serves private-market consumers. Eliminating it reduces availability for both groups; the accompanying fiscal saving reflects fewer redemptions.

This result motivates extending the stocking requirement to a second contracted brand. We next examine whether competition for contracts can finance this broader choice set for private-market consumers.

7.2 Competition for Contracts

We compare one, two, and three contracted brands, keeping formula free for beneficiaries and applying the perceived stockout penalty to every contracted brand.⁵¹ All comparisons below use outcomes after manufacturers adjust their wholesale prices.

Profit-only bidding Following the interpretation in Section 6.3, the mean type $\mu_{j,\ell}^\eta$ is set to zero in the counterfactual auctions, so bidder j 's type in market m is its idiosyncratic component $\tilde{\eta}_{jm}$. μ_j^η is the mean of bidder j 's cost specific to fulfilling WIC contracts and is common

⁵¹Each contracted brand carries its own penalty Λ_j in (8). An alternative rule would penalize the store only when all contracted brands are out of stock. Under that rule, a store can meet the requirement with a single brand. In a simulation of this rule, stores concentrate inventory in the more profitable contracted brand and stocks the other as it would a non-WIC brand, so contracting with a second brand does little to broaden private-market consumers' choice sets.

knowledge among bidders; we interpret it as a behavioral departure from bidding based on modeled profits alone.⁵² We re-solve the first-price auction is re-solved as implemented in the status quo and find that this change has little aggregate effect on spending and private-market consumer surplus. Contracts nevertheless shift away from Nestlé, which has been bidding much more aggressively relative to the benefits it enjoys from WIC contracts. We compare with the status quo contracts to establish a policy-relevant comparison.

Comparison between auction formats We first consider a change in the auction format, replacing the first-price auction with a VCG mechanism while keeping a single contracted brand. This design (“Sole-source (VCG)” in Table 3 and the figures) reduces annual net program spending from \$373 million under the profit-only first-price auction to \$89 million, as manufacturer payments rise from \$1.21 billion to \$1.53 billion (Figure 9).⁵³ Relative to the status quo contracts, it saves \$279 million annually. Switching to VCG thus lowers net program spending with little change in private-market consumer surplus.

Two contracted brands Under *dual sourcing*, WIC awards contracts to two brands using the VCG mechanism and applies the stocking requirement to both. Retailers are then incentivized to keep the second contracted brand available as well. The share of small-box consumers facing only one available brand falls from 22.3 percent in the status quo to 10.5 percent (Figure 8).

With two of three manufacturers winning, competition for contracts weakens, but each winner can still be replaced by the third manufacturer, so manufacturers continue to make payments to the government. In total, they pay \$1.18 billion a year, compared with \$1.53 billion when the same mechanism awards a single contract (Figure 9). Relative to that single-contract design, dual sourcing raises net program spending by \$278 million and annual private-market consumer surplus by \$330 million. Relative to the status quo contracts, it raises private-market consumer surplus by \$340 million for about the same net program spending, without reducing WIC redemption (Figure 7). Changing procurement rules can therefore deliver either substantial cost savings or greater consumer surplus within the existing program budget.

Three contracted brands Under *triple sourcing*, WIC contracts with all three brands and applies the stocking requirement to each. Two thirds of small-box consumers find all three brands

⁵²Appendix Section D.3 instead treats it as an invariant opportunity cost and reports results with type $\eta_{jm} = \mu_j^\eta + \tilde{\eta}_{jm}$.

⁵³VCG generates higher manufacturer payments both for the recovered realization of types and in expectation over the estimated type distribution. In expectation, annual manufacturer payments are \$60 million higher under VCG than under the first-price auction.

in stock, compared with about a third under dual sourcing (Figure 8). Relative to dual sourcing, the third contract raises annual private-market consumer surplus by about \$190 million and net program spending by \$1.12 billion, or \$5.84 in additional spending per dollar of private-market consumer surplus. Relative to the status quo contracts, private-market consumer surplus rises by \$540 million, of which \$90 million comes from manufacturers lowering wholesale prices (Figure 10).

From a practical perspective, preserving competition while broadening eligibility requires more bidders than contracted brands. GAO (2025) reports that only two manufacturers were bidding for WIC contracts after the 2022 infant formula shortage (which happened after our study period for procurement contracts), so contracting with both would remove the exclusion threat present in our three-bidder environment. Extending auction participation to more firms as well as encouraging entry to the industry can therefore create room to broaden eligibility as well as strengthen competition via rebates to the WIC program.


7.3 Other Program Designs

Separate distribution We then explore the advantage and disadvantage of using a private retailer who serves both WIC and private-market consumers. In particular, retailers can pay one ordering cost for the same pool of inventory that serves both types. To examine the efficiency gain of pooling inventory, we remove the pooling and give beneficiaries a dedicated product that private consumers cannot buy, labeled the *government brand* in the figures. The stocking requirement applies to that product; retailers choose private-brand inventories purely to maximize the discounted value of expected profit. WIC pays production and inventory costs, which removes manufacturer markups.⁵⁴ Public or nonprofit production of this kind has been proposed for markets in which manufacturers have market power, most prominently generic drugs and insulin (Liljenquist et al., 2018; Agrawal et al., 2025; Sherkow et al., 2023).

Even though the program buys formula at production cost, the government brand raises annual program spending by \$220 million relative to the status quo. WIC redemption is unchanged, and private-market consumer surplus is \$10 million higher. Under the status quo, most manufacturers already supply WIC and earn a per-unit revenue below production cost in exchange for private-market sales (Section 6.3). A dedicated product removes those sales, and with them manufacturers' reason to help finance WIC through contract payments. The program also pays the ordering costs of an inventory that serves beneficiaries alone, which accounts for

⁵⁴Production cost is set to the lowest recovered cost among the three manufacturers. The program also pays each market's average retail margin on WIC units and the dedicated product's ordering costs.

Figure 7: Fiscal Cost and Changes in Private Market Consumer Surplus Across Procurement Designs



Notes: Both axes report annual changes from the status quo, in millions of dollars. The top axis also reports net fiscal cost in levels. Solid outlines denote outcomes with adjusted wholesale prices; the dashed outline denotes the current-allocation, no-mandate comparison at fixed wholesale prices. Circles include the perceived stockout penalty and squares omit it. Marker colors distinguish policy margins and marker sizes rank private-market volume. WIC percentages in the labels report changes in WIC volume. Consumer-surplus differences exclude WIC beneficiaries and each brand's allocated share of β_{WIC} .

most of the increase in spending. Dedicated inventory can no longer serve private-market consumers, so more of them find no brand in stock. At observed prices, the share of small-box private-market consumers in this situation rises from 1.6 to 4.3 percent, and private-market consumer surplus falls by \$110 million; manufacturers lower wholesale prices in response, and the average private-market transaction price falls by 4.2 percent, which offsets this loss. Because the contracted brand is also sold to private-market consumers, winning a contract raises a manufacturer's private-market sales, and competition for contracts leads manufacturers to give up part of this value to the program. A dedicated product removes this private-market spillover, so the program loses these payments in addition to paying for a separate inventory, while for private-market consumers lower wholesale prices roughly offset the loss in availability.


Coupons WIC cannot afford to provide formula free to all beneficiaries while purchasing it at market prices. However, at roughly the same costs, it can finance a coupon of \$7.50 per 90 prepared ounces that can be used toward any brand. The costs are a reduction in WIC redemption, by about 40 percent, and lower private-market consumer surplus. The consumers who are priced out of the program are those who are the poorest and are the most price-sensitive. Pricing the poorest out of the program contradicts the program's goal of equalizing access to nutrition.

In addition, we keep the coupon at the same per-unit subsidy and explore the effect of inducing manufacturer competition using restrictive contracts. The sole- and dual-sourcing VCG rules are applied to the coupon design. At observed wholesale prices, the two-brand coupon raises private-market consumer surplus by as much as dual sourcing. Manufacturers then raise wholesale prices, and the average private-market transaction price rises by 6.4 percent, which reduces the gain to \$190 million (Figure 10). Under vouchers, rebates adjust one for one with wholesale prices (Section 4.3), so a manufacturer's revenue on voucher units does not depend on its wholesale price, and it sets that price against private-market demand alone. Under coupons, this rule does not apply, and the contracted manufacturers earn their wholesale price on units bought with coupons. Beneficiaries can use the coupon only on the contracted brands, so their demand for those brands is less elastic than private-market demand, and manufacturers set higher wholesale prices. Manufacturer payments exceed the program's coupon outlays (Figure 9), but the design still prices out over a quarter of WIC beneficiaries.

Program scale Expanding WIC demand to 1.5 times its baseline level under sole-source VCG moves consumers from the private market into the program. Manufacturers compete to supply more WIC volume, but the private market from which contracts generate spillovers shrinks.


Per-consumer surplus changes little for those who remain in the private market.

Figure 8: Private Consumers' Choice Sets in Small-Box Stores Across Procurement Designs



Notes: Each row shows the consumer-weighted distribution of available-brand counts in small-box stores. The status quo uses observed wholesale prices; all other rows use adjusted wholesale prices. Counts describe the choice sets private consumers face during simulated purchases. Colors distinguish zero, one, two, and three brands. Percentages are rounded, and small segments are not labeled. The government brand is unavailable to private consumers and excluded from their choice sets. Consumer-surplus differences also account for brand identities and prices.

Figure 9: Gross Program Outlays, Manufacturer Payments, and Net Fiscal Cost



Notes: Each waterfall starts at gross program outlays, subtracts manufacturer payments, and ends at net fiscal cost. All amounts are annual national totals in millions of dollars, using adjusted wholesale prices except for the status quo and the current-allocation, no-mandate comparison, shown with pale bars. Manufacturer payments are observed rebates, pay-as-bid rebates, or VCG payments as specified by the design. For the government brand, outlays cover production and distribution costs. Negative net fiscal cost means manufacturer payments exceed program outlays.

Figure 10: Consumer-Surplus Differences and Wholesale-Price Responses



Notes: Red bars report private-market consumer-surplus differences from the status quo with wholesale prices held fixed. Green steps report the additional change when manufacturers adjust wholesale prices, and gray bars report the total difference from the status quo. The current-allocation, no-mandate design is reported only at fixed wholesale prices. All values are annual national totals in millions of dollars. Welfare calculations exclude each brand's allocated share of β_{WIC} in both the counterfactual and the status quo and exclude WIC beneficiaries.

8 Conclusion

Public procurement of in-kind benefits affects private-market consumers through the stocking decisions of retailers that distribute these benefits. Procurement designs can elicit competition between manufacturers of the targeted goods for procurement contracts, such that the government saves on fiscal costs relative to purchasing these goods at market prices. However, with private retail distribution of these benefits, the cost savings can come at the expense of private-market consumer surplus. In WIC's procurement of infant formula, which uses competitive bidding to solicit sole-source contracts and requires retailers to stock products of the contracted brand, the status quo procurement mechanism generates substantial cost savings. Using changes in WIC brands across different states, we demonstrate empirically that these contracts also have private-market spillovers in the form of substantially increasing a WIC brand's private-market share. These spillovers are driven by retailers' stocking decisions that steer private-market consumers to the WIC brand.

To quantify the tradeoff between government cost savings and private-market consumer surplus while keeping the private retailer distribution channel, we develop and estimate a model connecting consumer choices, retailers' dynamic inventory decisions, and manufacturer competition for WIC contracts and pricing in the private market. The novel component is the retailer's inventory problem, which extends the single-product (s, S) model of Arrow et al. (1951) to substitutable brands that are stocked jointly. Whereas empirical models of inventory typically abstract from interactions between products or capture them through cost parameters, in our model stocking one brand changes the return to replenishing its substitutes. The model rationalizes retailers' observed responses to WIC contracts as a profit-maximizing outcome. Because stocking responds to procurement rules in the model, we can evaluate designs that change the number of contracted brands and the stocking requirement. Counterfactual simulations from the model show that contracting with two of the three manufacturers and requiring retailers to stock both brands can increase private-market consumer surplus by broadening consumers' choice sets while preserving competition for WIC contracts. The third manufacturer provides a threat of replacement for either winner, so manufacturers continue to make substantial payments to avoid losing access to WIC demand and its associated private-market sales.

Many in-kind benefits rely on private intermediaries for product or service distribution, such as insurers for public health insurance plans, hospitals for health care services, pharmacies for drugs, and grocery stores for other types of nutritional assistance. When upstream suppliers have market power, procurement rules used for fiscal cost minimization can change these intermediaries' incentives and affect the balance between fiscal savings and private-market con-

sumer surplus. The policy question is therefore how to balance the two after internalizing the private distributors' incentives. Exploration of other channels through which private distributors' incentives can be affected by public procurement designs provides a promising avenue for future research. Application of the retailer stocking mechanism to other settings has the potential to improve policy designs.

References

- Abito, Jose Miguel, Kathleen Hui, Yuval Salant, and Kosuke Uetake**, “Unintended Consequences of Public Procurement: Evidence from the WIC Program,” Technical Report CFDP 2555, Cowles Foundation for Research in Economics 2026.
- Agarwal, Nikhil and Paulo J Somaini**, “Demand analysis under latent choice constraints,” Technical Report 29993, National Bureau of Economic Research 2022. Revised November 2025.
- Agrawal, Sahil, Melissa Barber, Amy Kapczynski, and Trudel Pare**, “Drug Dealing: Making Public Pharma Work,” *Washington University Law Review*, 2025, 103 (1), 49–102.
- Aguirregabiria, Victor and Francis Guiton**, “Decentralized Decision-Making in Retail Chains: Evidence from Inventory Management,” 2023. Working paper.
- American Academy of Pediatrics Committee on Nutrition**, “The Use of Whole Cow’s Milk in Infancy,” *Pediatrics*, 1992, 89 (6), 1105–1109.
- An, Yonghong, David Davis, Yizao Liu, and Ruli Xiao**, “Procurement in welfare programs: Evidence and implications from WIC infant formula contracts,” 2023.
- Andreyeva, Tatiana, Joerg Luedicke, Ann E. Middleton, Michael W. Long, and Marlene B. Schwartz**, “Changes in Access to Healthy Foods after Implementation of the WIC Food Package Revisions,” Contractor and Cooperator Report 66, U.S. Department of Agriculture, Economic Research Service April 2011.
- Arrow, Kenneth J, Theodore Harris, and Jacob Marschak**, “Optimal inventory policy,” *Econometrica: Journal of the Econometric Society*, 1951, pp. 250–272.
- Arthur, David and Sergei Vassilvitskii**, “k-means++: The Advantages of Careful Seeding,” in “Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms” 2007, pp. 1027–1035.
- Berry, Steven, James Levinsohn, and Ariel Pakes**, “Automobile Prices in Market Equilibrium,” *Econometrica*, 1995, 63 (4), 841–890.
- Byrne, Anne T., Xiao Dong, Jessie Handbury, Erik James, Katherine Meckel, and Andrés C. Rovira**, “Welfare Implications of Increased Retailer Participation in SNAP,” Technical Report 34467, National Bureau of Economic Research 2025.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Centers for Disease Control and Prevention**, “Breastfeeding Report Card: United States, 2022,” Technical Report, Centers for Disease Control and Prevention 2022.

- , “Natality Information,” n.d. <https://wonder.cdc.gov/natality.html>. Accessed September 19, 2026.
- CME Group**, “Class III Milk,” n.d. <https://www.cmegroup.com/markets/agriculture/dairy/class-iii-milk>. Accessed September 19, 2026.
- Code of Federal Regulations**, “Title 7, Section 246.16a: Infant Formula and Authorized Foods Cost Containment,” 2011.
- , “Title 21, Part 106: Infant Formula Requirements Pertaining to Current Good Manufacturing Practice, Quality Control Procedures, Quality Factors, Records and Reports, and Notifications,” n.d. <https://www.ecfr.gov/current/title-21/chapter-I/subchapter-B/part-106>. Accessed August 31, 2026.
- , “Title 7, Section 246.12: Food Delivery Methods,” n.d. <https://www.law.cornell.edu/cfr/text/7/246.12>. Accessed August 31, 2026.
- Competition and Markets Authority**, “Infant Formula and Follow-on Formula Market Study: Final Report,” Technical Report, Competition and Markets Authority February 2025.
- Conlon, Christopher and Jeff Gortmaker**, “Best practices for differentiated products demand estimation with PyBLP,” *The RAND Journal of Economics*, 2020, 51 (4), 1108–1161.
- Conlon, Christopher T and Julie Holland Mortimer**, “Demand estimation under incomplete product availability,” *American Economic Journal: Microeconomics*, 2013, 5 (4), 1–30.
- Currie, Janet and Firouz Gahvari**, “Transfers in Cash and In-Kind: Theory Meets the Data,” *Journal of Economic Literature*, 2008, 46 (2), 333–383.
- Davis, David E.**, “WIC Infant Formula Rebate Database: Description, Data, and State-Alliance Key,” SDSU Working Papers in Progress 12008, South Dakota State University, Department of Economics September 2016.
- de Chaisemartin, Clément and Xavier D’Haultfœuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, 110 (9), 2964–2996.
- Dean, Emma B., Josh Feng, and Luca Maini**, “Stocking Under the Influence: Spillovers from Commercial Drug Coverage to Medicare Utilization,” 2026. Working paper, May 3.
- DellaVigna, Stefano and Matthew Gentzkow**, “Uniform Pricing in U.S. Retail Chains,” *The Quarterly Journal of Economics*, 2019, 134 (4), 2011–2084.
- Duggan, Mark and Fiona M. Scott Morton**, “The Distortionary Effects of Government Procurement: Evidence from Medicaid Prescription Drug Purchasing*,” *The Quarterly Journal of Economics*, 02 2006, 121 (1), 1–30.

- FDA**, “Status Update on FDA’s Infant Formula Response Activities,” U.S. Food and Drug Administration December 2023. <https://www.fda.gov/food/infant-formula-guidance-documents-regulatory-information/status-update-fdas-infant-formula-response-activities>.
- Goldin, Jacob, Tatiana Homonoff, and Katherine Meckel**, “Issuance and incidence: Snap benefit cycles and grocery prices,” *American Economic Journal: Economic Policy*, 2022, 14 (1), 152–178.
- Gosliner, Wendi, Daniel M. Brown, Betty C. Sun, Gail Woodward-Lopez, and Patricia B. Crawford**, “Availability, Quality and Price of Produce in Low-Income Neighbourhood Food Stores in California Raise Equity Issues,” *Public Health Nutrition*, 2018, 21 (9), 1639–1648.
- Grummer-Strawn, Laurence M., Kelley S. Scanlon, and Sara B. Fein**, “Infant Feeding and Feeding Transitions During the First Year of Life,” *Pediatrics*, 2008, 122 (Supplement 2), S36–S42.
- Handbury, Jessie and Sarah Moshary**, “School food policy affects everyone: Retail responses to the national school lunch program,” Technical Report, National Bureau of Economic Research 2021.
- Hastings, Justine and Ebonya Washington**, “The First of the Month Effect: Consumer Behavior and Store Responses,” *American Economic Journal: Economic Policy*, 2010, 2 (2), 142–162.
- Huang, Rui and Jeffrey M. Perloff**, “WIC Contract Spillover Effects,” *Review of Industrial Organization*, 2014, 44 (1), 49–71.
- Institute for Local Self-Reliance**, “How Big Are Big-Box Stores?,” December 2008. <https://ilsr.org/article/independent-business/how-big-are-bigbox-stores/>.
- Jaravel, Xavier**, “What Is the Impact of Food Stamps on Prices and Products Variety? The Importance of the Supply Response,” *AEA Papers and Proceedings*, 2018, 108, 557–561.
- Jehiel, Philippe and Benny Moldovanu**, “Auctions with Downstream Interaction Among Buyers,” *RAND Journal of Economics*, 2000, 31 (4), 768–791.
- Kastl, Jakub, Eric Richert, and Jesper Rüdiger**, “Efficient Grid Search to Solve Static Games with Private Information,” 2025. Working paper.
- Kaufman, Phillip, Charles Handy, Edward W. McLaughlin, Kristen Park, and Geoffrey M. Green**, “Understanding the Dynamics of Produce Markets: Consumption and Consolidation Grow,” Agriculture Information Bulletin 758, USDA Economic Research Service August 2000. <https://www.ers.usda.gov/publications/42295>.
- Kuehn, Joseph**, “Estimating Auctions with Externalities: The Case of USFS Timber Auctions,” *Review of Economics and Statistics*, 2019, 101 (5), 791–807.

- Lee, Munseob**, “Government Purchases and Firm Growth,” 2021. Working paper, University of California, San Diego.
- Leung, Justin H. and Hee Kwon Seo**, “How Do Government Transfer Payments Affect Retail Prices and Welfare? Evidence from SNAP,” *Journal of Public Economics*, 2023, 217, 104760.
- Liljenquist, Dan, Ge Bai, and Gerard F. Anderson**, “Addressing Generic-Drug Market Failures — The Case for Establishing a Nonprofit Manufacturer,” *New England Journal of Medicine*, 2018, 378 (20), 1857–1859.
- Mahajan, Siddharth and Garrett van Ryzin**, “Stocking Retail Assortments Under Dynamic Consumer Substitution,” *Operations Research*, 2001, 49 (3), 334–351.
- Mahardika, Ivan**, “The Evolution, Implementation, and Evidence of Indonesia’s Food Assistance Program: From Rice Subsidy to Electronic Food Vouchers,” in Ugo Gentilini, ed., *Social Protection, Food Security and Nutrition: An Update of Concepts, Evidence and Select Practices in South Asia and Beyond*, Washington, DC: World Bank, 2022, chapter 6.
- Maskin, Eric and John Riley**, “Asymmetric Auctions,” *Review of Economic Studies*, 2000, 67 (3), 413–438.
- Mead Johnson Nutrition Company**, “Annual Report on Form 10-K for the Year Ended December 31, 2016,” Technical Report, U.S. Securities and Exchange Commission 2017.
- Meckel, Katherine**, “Is the cure worse than the disease? unintended effects of payment reform in a quantity-based transfer program,” *American Economic Review*, 2020, 110 (6), 1821–1865.
- , **Maya Rossin-Slater, and Lindsey Uniat**, “Efficiency versus Equity in the Provision of In-Kind Benefits: Evidence from Cost Containment in the California WIC Program,” *Journal of Human Resources*, 2023, 58 (2), 363–392.
- Michigan Department of Agriculture and Rural Development**, “Mead Johnson and Company, LLC: Food and Agriculture Investment Fund,” Technical Report, Michigan Commission of Agriculture and Rural Development July 2019. July 18 meeting materials, pp. 57–58.
- Ministerio de Economía, Fomento y Turismo**, “Bolsillo Familiar Electrónico: Guía para el comercio,” May 2023. Government of Chile. May 12, 2023.
- Nevo, Aviv**, “Measuring Market Power in the Ready-to-Eat Cereal Industry,” *Econometrica*, 2001, 69 (2), 307–342.
- NHS Business Services Authority**, “Retailers,” n.d. Healthy Start. Accessed August 31, 2026.
- Oliveira, Victor and Mark Prell**, “Sharing the Economic Burden: Who Pays for WIC’s Infant Formula,” USDA Economic Research Service, Amber Waves September 2004. <https://ers.usda.gov/amber-waves/2004/september/sharing-the-economic-burden-who-pays-for-wic-s-infant-formula>.

- , **Elizabeth Frazão, and David M Smallwood**, “Rising infant formula costs to the WIC program: recent trends in rebates and wholesale prices,” 2010.
- Pelletier, Jennifer E, Liana RN Schreiber, and Melissa N Laska**, “Minimum stocking requirements for retailers in the special supplemental nutrition program for women, infants, and children: disparities across US states,” *American journal of public health*, 2017, 107 (7), 1171–1174.
- Rojas, Christian A. and Hongli Wei**, “Spillover Mechanisms in the WIC Infant Formula Rebate Program,” *Journal of Agricultural & Food Industrial Organization*, 2019, 17 (2), 1–14.
- Sanders, Robert Evan**, “Dynamic Pricing and Organic Waste Bans: A Study of Grocery Retailers’ Incentives to Reduce Food Waste,” *Marketing Science*, 2024, 43 (2), 289–316.
- Sherkow, Jacob S., Eli Y. Adashi, and I. Glenn Cohen**, “Assessing—and Extending—California’s Insulin Manufacturing Initiative,” *JAMA*, 2023, 329 (7), 533–534.
- Smilowitz, Jennifer T. et al.**, “Cross-Sectional Analysis of Infant Diet, Outcomes, Consumer Behavior and Parental Perspectives to Optimize Infant Feeding in Response to the 2022 U.S. Infant Formula Shortage,” *Nutrients*, 2024, 16 (5), 748.
- U.S. Census Bureau**, “American Community Survey Data,” n.d. <https://www.census.gov/programs-surveys/acs/data.html>. Accessed September 19, 2026.
- U.S. Department of Health and Human Services**, “Prior HHS Poverty Guidelines and Federal Register References,” Office of the Assistant Secretary for Planning and Evaluation n.d. <https://aspe.hhs.gov/topics/poverty-economic-mobility/poverty-guidelines/prior-hhs-poverty-guidelines-federal-register-references>. Accessed September 19, 2026.
- U.S. Government Accountability Office**, “WIC Infant Formula: Single-Supplier Competitive Contracts Reduce Program Costs and Modestly Increase Retail Prices,” Report GAO-25-106503, U.S. Government Accountability Office January 2025.
- USDA Food and Nutrition Service**, “SNAP Retailer Locator Data,” n.d. <https://www.fns.usda.gov/snap/retailer-locator/data>. Accessed September 19, 2026.
- , “WIC Program Data,” n.d. <https://www.fns.usda.gov/pd/wic-program>. Accessed September 19, 2026.
- World Food Programme**, “Nigeria,” n.d. Accessed August 31, 2026.

A Data and Descriptive Evidence

A.1 Data Construction

Auction markets and purchasing alliances Some states procure formula jointly through purchasing alliances. We define an auction market m as a state or alliance in a given contracting period. The USDA contract records identify each state agency, its alliance, and the contract's start and expiration dates. For example, Arkansas, New Mexico, and North Carolina share an alliance, with common contract dates and winners throughout 2006–2020. We combine these records with the market identifiers in Davis (2016), reconcile alliance names across the two sources, and assign membership by state and month. This allows membership to change over time and gives winning and losing bids the same market definition. In the auction model, we aggregate projected store quantities across the member states, so a joint auction's bids and associated WIC and private-market sales enter as one market.

Bids across formula forms Manufacturers submit separate rebates for powdered, liquid-concentrate, and ready-to-use formula, which states evaluate using standardized quantities of each form.⁵⁵ We do not observe these weights. Because powder is the dominant form, we abstract from the other forms and model the auction using powdered formula alone.

Approved Product Lists We supplement contract-to-UPC matches with Approved Product Lists from Idaho, Iowa, Maine, Massachusetts, Missouri, Ohio, and Washington, obtained through public records requests. To exclude additions associated with the 2022 formula shortage, we require APL products to have a recorded authorization before 2022.

Consumer Panel household classification The WIC participation indicator is missing both for households reporting no participation and for households that do not respond. We therefore combine participation reports with household income and household size to identify private-market households.

We exclude household observations from the first year in which the household reports current or past WIC participation onward. We also exclude households if, in any observed year, the upper bound of their reported income bracket is at or below 185 percent of the federal poverty guideline for their household size. We match the income report to the guideline for two years before the panel year. Income brackets that straddle the threshold are not excluded by this income screen.

⁵⁵See Code of Federal Regulations (2011, §246.16a(c)(3)–(4)).

Allocating sales between private and voucher purchases For model estimation, we use the Consumer Panel to estimate the private-purchase share of eligible-product sales of the local contract brand. We restrict panel purchases to products identified by the contract records or Approved Product Lists whose brand holds the local WIC contract. Within each brand and retail channel, we divide the prepared ounces purchased by households passing the private-market classification above by the prepared ounces purchased by all households. These shares pool purchases nationally across years and product forms and use purchase volumes, weighted by the households' sampling weights.

Let $\hat{a}_{jc}$ denote this private-purchase share for brand j in channel c , and let A_{jst} denote store s 's eligible-product sales of that brand in period t , set to zero when j does not hold the local contract. If Y_{jst} denotes the brand's total sales, the allocation is

$$\hat{Y}_{jst}^V = (1 - \hat{a}_{j,c(s)})A_{jst}, \quad (\text{A.1})$$

$$\hat{Y}_{jst}^{\text{priv}} = Y_{jst} - \hat{Y}_{jst}^V. \quad (\text{A.2})$$

Thus, private purchases include the panel-implied private share of the contract brand's eligible-product sales and all other sales. This allocation is used for estimation; the descriptive analysis isolates private purchases using products that are never WIC-eligible.

Additional data sources We use household income distributions from the American Community Survey (U.S. Census Bureau, n.d.) to represent demand heterogeneity and CME Class III milk futures prices (CME Group, n.d.) as instruments for retail prices. We construct store projection weights from USDA records of SNAP-authorized retailers (USDA Food and Nutrition Service, n.d.a), since the stores in our Scanner Data are not a random sample of the retail universe. USDA WIC participation data (USDA Food and Nutrition Service, n.d.b) and CDC Natality data (Centers for Disease Control and Prevention, n.d.) provide infant and program population measures, which we also use to test if manufacturers benefit from economies of scale from WIC contracts. Finally, we use HHS poverty guidelines (U.S. Department of Health and Human Services, n.d.) for the Consumer Panel household income screen.

A.2 Alternative Explanations for Private Market Spillovers

Labelling and retail visibility Consumers may favor the WIC brand, holding availability fixed, because of perceived government endorsement or greater visibility. Our stocking mechanism does not require these demand-side channels to be absent. When we introduce the model

in Section 4, we explicitly allow for a preference for the WIC brand that can reflect labeling or perceived endorsement.

WIC consumers supplementing their benefits It is possible that WIC benefits do not cover the full demand of WIC consumers and therefore, WIC consumers may purchase more of the WIC brand out-of-pocket to satisfy their needs, driving our observed spillover effects. In Appendix Figure A.10, we use the same Consumer Panel data to show that WIC households likely stay within the limits of their benefit packages; thus, spillovers of this magnitude are unlikely to be due to WIC beneficiaries exhausting their benefit package and supplementing with products from the same brand.

Manufacturer's economies of scale Winning a WIC contract raises a manufacturer's production volume. If marginal cost falls with volume, the contract lowers the winner's cost of serving the private market. Therefore, we might expect the winner to serve more private market consumers due to the benefits of scale. However, we find little response of private-market retail prices to contract changes (Appendix Figure A.7), providing little evidence of cost savings passed through to consumers. To examine manufacturer economies of scale more directly, we test whether a firm's bids respond to increases in the volume of existing WIC contracts. A proper test requires exogenous variation in a firm's production quantity, in the spirit of Lee (2021). We use two such shifters, measured in the firm's other markets: growth in births and growth in fully-formula-fed WIC participation. In Appendix Table A.3, we find that a firm's bid does not respond to plausibly exogenous growth in its volume in other markets.

Comparison with existing evidence Broadly, other studies of WIC's procurement of infant formula, including Huang and Perloff (2014), Rojas and Wei (2019), and An et al. (2023), also find large private-market spillovers of WIC contracts. We add evidence on the mechanism behind these spillovers, which matters for procurement design because the effect of a procurement rule depends on retailers' endogenous responses. Abito et al. (2026) document a roughly 30-percentage-point increase in the new WIC supplier's share of regular-size formula purchases among non-WIC households and also study the mechanism, pinpointing retailer behavior as the driver. We build on their findings by documenting intermittent stockouts of competing brands in small-box stores and modeling the stocking incentives behind them.


To estimate spillover effects, Abito et al. (2026) utilize data which records WIC benefit use and links purchases to retailers' authorization status, providing a more direct and complete measure of private-market spillovers. To construct a comparison point, we use the NielsenIQ Con-

sumer Panel to construct state-level quarterly brand market shares among households classified as private-market consumers based on self-reported WIC participation, income, and household size. We estimate the market share event study using these purchases and find that the winning brand gains 21.2 percentage points four quarters after the contract change (95% CI: 4.2–38.3 percentage points), an interval that includes their estimate (Appendix Figure A.3).

Abito et al. (2026)’s event studies show little response for bulk-size products. Applying their greater-than-16-ounce cutoff for bulk size within products never eligible to WIC redemptions, we find positive spillovers, particularly in small-box stores (Appendix Figure A.4 and Appendix Figure A.5). The difference between our result and theirs could be driven by the selection of (higher-income) panelists in the Numerator data, or be driven by low statistical power, since their state-level monthly market share data come from the aggregation of purchases by a smaller sample of consumers, compared to the much more vast coverage of NielsenIQ’s store-level data.


Abito et al. (2026) also find that spillovers are concentrated in physical retail and persist among retailers selling all three national brands, and argue that product availability does not seem to play a role in driving the spillovers. However, in our Scanner Data sample, 97% of small-box stores sell all three brands at some point during a contract period. Therefore, we argue that their finding is consistent with our interpretation of intermittent stockouts in small-box stores, who need not keep brands continuously available.

Figure A.1: Event-Study Sample




Notes: The figure shows the timing of WIC contract changes and sample coverage by state and brand during 2006–2020. Red denotes brands that gain a contract, blue denotes brands that lose a contract, and gray denotes brands whose contract status does not change. Filled dots mark the month of the contract change. Blank rows denote brand–state pairs that both gain and lose a contract during the sample period and are excluded from both directional event studies.

Figure A.2: WIC's Effect on Private Market Product's Market Share, Alternative DID Estimators




Notes: Each panel re-estimates the market-share event study for private-market products in Figure 2, by store channel (big-box vs. small-box) and contract outcome (winning vs. losing brand), using three estimators: two-way fixed effects (TWFE), the Callaway and Sant'Anna (2021) estimator with not-yet-treated units as controls (CS), and the de Chaisemartin and D'Haultfœuille (2020) estimator (dCDH). The horizontal axis is months relative to the state's WIC contract change. Vertical bars are 95% confidence intervals.

Figure A.3: WIC's Effect on Brand Market Shares among Private-Market Consumers




Notes: Event-study estimates from the NielsenIQ Consumer Panel, restricted to households classified as private-market consumers using the participation and income screens described above. Market shares are measured in prepared ounces within each state and quarter, using all milk-based powder products purchased by these households, including WIC-eligible products. Event time is measured in quarters relative to the state's WIC contract change. Estimates are reported for the winning and losing brands using two-way fixed effects (TWFE), the Callaway and Sant'Anna (2021) estimator with not-yet-treated controls (CS), and the de Chaisemartin and D'Haultfœuille (2020) estimator (dCDH). Vertical bars are 95% confidence intervals; standard errors are clustered by state.

Figure A.4: WIC's Effect on Private-Market Brand Shares, by Container Size




Notes: Two-way fixed-effects event-study estimates for milk-based powder products that were never WIC-eligible, pooling store types. Container size is measured in ounces of powder per container, rather than per multipack; products with missing size are excluded. Market shares are measured in prepared ounces within each state, month, and size category, among the three national brands. The larger-container category uses the greater-than-16-ounce cutoff for bulk-size products in Abito et al. (2026); both panels retain our restriction to never-WIC-eligible products. The winning and losing series follow the brands gaining and losing the state's WIC contract. Shaded bands are 95% confidence intervals, with standard errors clustered by state. Source: NielsenIQ Retail Scanner, 2006–2020.

Figure A.5: WIC's Effect on Private-Market Brand Shares, by Container Size and Store Type




Notes: Two-way fixed-effects event-study estimates using the product sample and container-size definitions in Appendix Figure A.4, estimated separately for big-box stores (mass merchandisers) and small-box stores (drug and grocery). Market shares are measured in prepared ounces within each state, month, size category, and store type, among the three national brands. The winning and losing series follow the brands gaining and losing the state's WIC contract. Shaded bands are 95% confidence intervals, with standard errors clustered by state. The vertical scale is common across store types within each size category but differs between the two size categories. Source: NielsenIQ Retail Scanner, 2006–2020.

Figure A.6: Winning the WIC Contract and Private-Market Product Availability




Notes: Event-study estimates of the probability that a private-market product of the winning brand has zero sales in a store-month. Estimates are shown separately for big-box (mass merchandiser) and small-box (drug and grocery) stores, by month relative to the WIC contract change (vertical dashed line). Shaded bands are 95% confidence intervals.

Figure A.7: WIC's Effect on Private Market Product Prices




Notes: Event-study estimates of the average retail price of private-market (never WIC-eligible) milk-based powder products, in dollars per 90 prepared ounces, by month relative to a state's WIC contract change (vertical dashed line). The *winning* series follows the brand that gains the state's WIC contract; the *losing* series follows the brand that loses it. Shaded bands are 95% confidence intervals.

Figure A.8: Distribution of Households' Big-Box Share




Notes: Sample is the NielsenIQ Consumer Panel restricted to private-market households with at least one milk-based powder formula purchase in 2006–2020 at a Discount Store, Grocery, Warehouse Club, or Drug Store. For each household we compute the share of its milk-powder quantities purchased at big-box stores (Discount + Warehouse Club) over the sample period and plot the cross-household distribution. The dashed line marks the mean. Quantities are in prepared ounces.

Figure A.9: Households' Within-Episode Brand Loyalty




Notes: Same consumer-panel sample as Appendix Figure A.8. A formula-consumption episode is a sequence of purchases with no more than 100 days between successive purchases, spanning less than 365 days. For each household episode, we identify the most frequently purchased brand and plot its share of total prepared ounces purchased.

Figure A.10: WIC Beneficiaries' Formula Purchases Relative to Benefit Limits



Notes: Monthly prepared ounces of WIC-contracted formula purchased by WIC beneficiaries in the NielsenIQ Consumer Panel, 2006–2020. A benefit month is dated from the start of a consecutive purchase episode (a new episode begins after 180 days without a formula purchase), with a month defined as 30.437 days. The sample is household-benefit-months with positive purchases of WIC-contracted products ($N = 593$ households). The vertical dashed lines mark the benefit-package limits: 314 prepared ounces (partial breastfeeding package, age 6–11 months) and 884 prepared ounces (fully formula-fed package, ages 4–5 months, which is the largest package).

Figure A.11: National Wholesale Prices Across WIC Auctions



Notes: Each dot is one firm-auction observation, 2006–2017: the firm’s posted national wholesale price at the time of the auction, in dollars per can (13-oz milk-based powder). Each firm maintains a single national wholesale price, revised infrequently in discrete jumps.

Table A.1: Balance of Household Characteristics by Store Type

	BB-type HHs	SB-type HHs	Diff
<i>Demographics</i>			
White	0.73	0.74	0.0025 (0.0245)
Black	0.08	0.10	0.0020 (0.0173)
Asian	0.07	0.05	0.0061 (0.0156)
Other	0.11	0.11	-0.0106 (0.0198)
Hispanic	0.17	0.17	-0.0106 (0.0259)
<i>Income</i>			
Annual household income < \$50,000	0.14	0.15	0.0085 (0.0203)
[\$50,000, \$70,000)	0.19	0.20	0.0151 (0.0238)
[\$70,000, \$100,000)	0.31	0.29	-0.0140 (0.0285)
> \$100,000	0.37	0.36	-0.0096 (0.0313)
<i>Employment</i>			
Any head employed	0.94	0.92	-0.0170 (0.0169)
Both heads employed full-time	0.44	0.38	-0.0720** (0.0301)
Female head employed full-time	0.55	0.52	-0.0324 (0.0313)
<i>F-stat (joint balance): 0.920 ($p = 0.513$)</i>			
Observations	7,293	4,159	

Notes: Source: NielsenIQ Consumer Panel, 2006–2020. Each household is classified once as “big-box type” (BB) if its big-box share of milk-powder quantities over the sample period is at least 0.75, and “small-box type” (SB) otherwise. The unit of observation is household $\times$ panel year. The “Diff” column reports the coefficient on an SB-type indicator from a bivariate regression with ZIP-code fixed effects and household-clustered standard errors, weighted by the NielsenIQ projection factor; the outcome is indicated in each row. The joint balance F -statistic tests the null that all 10 non-reference covariates jointly have zero coefficient in a regression of the BB-type indicator on the full covariate set (reference categories: race = Other, income > \$100,000). Employment indicators use NielsenIQ HMS head-of-household codes: “any head employed” = at least one head in codes 1–3, “full-time” = codes 2–3 (≥ 30 hrs). Quantities are in prepared ounces. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.2: Consumers' Brand and Store Choice After WIC Contract Change

	Switch brand (New WIC share) (1)	Change channels (Big-box share) (2)
Treated $\times$ Post	0.045 (0.055)	0.146* (0.085)
Treat $\times$ Post $\times$ small-box shopper	0.236* (0.124)	
Sample Observations	All HHs 432	Baseline SB 148

Notes: Coefficients from household $\times$ purchasing spell $\times$ pre/post-period panel regressions around changes in a state's WIC contracted brand. Each purchasing spell spans a contract change, and the pre- and post-periods are the spell's purchases before and after the change. The "treated" group is households whose pre-period purchases were mostly on the outgoing WIC brand. "Small-box shopper" (in the triple interaction) indicates households whose post-contract switch transactions were mostly at small-box stores. "Baseline SB" in column (2) restricts to households whose pre-contract switch transactions were mostly at small-box stores. The sample includes private-market households and purchases of powder milk-based infant formula; transactions are restricted to Discount, Grocery, Warehouse Club, and Drug. Regressions are weighted by pre-period quantities. Outcomes: (1) share of the household's period quantities on the incoming WIC brand; (2) share purchased at big-box stores (Discount + Warehouse Club). Quantities are in prepared ounces. All specifications use episode fixed effects and cluster standard errors at the household level; column (1) also includes the lower-order terms of the triple interaction. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.3: Manufacturer Scale Economies: Bids and Other-Market Volume Growth

	Bid		Log-bid	
	(1)	(2)	(3)	(4)
Exposure: other-market births growth	0.0911 (0.5505)		-5.4699 (13.4292)	
Exposure: other-market WIC participation growth		-0.3582 (0.3133)		-7.1029 (7.8499)
Obs	165	78	165	78

Notes: The outcome is a firm's bid, measured as the net wholesale price in dollars per prepared ounce (columns 1–2) or its log (columns 3–4). In columns 3–4, the coefficient measures the elasticity of the bid with respect to other-market scale. The regressor is the log-change over a two-year window in the total volume of the firm's frozen leave-one-out portfolio of *other* markets, measured by births growth (columns 1 and 3) or by fully-formula-fed WIC participation growth (columns 2 and 4, available 2009–2017, hence the smaller sample). All specifications include firm, market/alliance, and year fixed effects, and the sample is restricted to auctions in which all three firms submitted a bid. Robust standard errors are in parentheses.

B Model Appendix

B.1 The WIC Auction: Derivations and Equilibrium Characterization

Fix a market and suppress its subscript m . We work with each firm's expected order quantity $\mathcal{Q}_j(w, a)$ when firm a holds the contract, common knowledge under Assumption 4.2, with $\mathcal{Q}_j(w, j) = Y^V + Y_{j|WIC=j}^{\text{priv}}$ and $\mathcal{Q}_j(w, k) = Y_{j|WIC=k}^{\text{priv}}$. Orders are valued at the posted wholesale price w_j , so the payoffs of Section 4.3 are

$$\Pi_j^{\text{win}}(b) = (w_j - c_j) \mathcal{Q}_j(w, j) - (w_j - b) Y^V, \quad \Pi_j^{\text{lose}}(k) = (w_j - c_j) \mathcal{Q}_j(w, k),$$

with the bid entering only the win profit, through the rebate $r_j = w_j - b$: $\partial \Pi_j^{\text{win}} / \partial b = Y^V$.

The bidding first-order condition With two rivals k and l , firm j 's probability of winning at net price b is $\bar{G}_j(b) = (1 - G_k(b))(1 - G_l(b))$, with density $\bar{g}_j(b) = -\bar{G}'_j(b) = g_k(b)(1 - G_l(b)) + g_l(b)(1 - G_k(b))$. The probability that j loses to rival k specifically — k is the overall lowest —

integrates over k 's bid below b while l bids higher still,

$$P_k(b) = \mathbb{P}(b_k < b, b_k < b_l) = \int_{-\infty}^b (1 - G_l(x)) g_k(x) dx, \quad P'_k(b) = (1 - G_l(b)) g_k(b),$$

and $P'_k + P'_l = \bar{g}_j$. At $\eta_j = 0$, firm j chooses b to maximize

$$\mathbb{E}\Pi_j(b) = \bar{G}_j(b) \Pi_j^{\text{win}}(b) + P_k(b) \Pi_j^{\text{lose}}(k) + P_l(b) \Pi_j^{\text{lose}}(l).$$

Differentiating, using $\partial \Pi_j^{\text{win}} / \partial b = Y^V$ and $\bar{G}'_j = -\bar{g}_j$, and dividing each P'_k by $\bar{g}_j$ gives the externality weights, $\lambda_{k|j} = P'_k / \bar{g}_j = h_k / (h_k + h_l)$, and the first-order condition

$$\bar{g}_j(b) \left[\Pi_j^{\text{win}}(b) - \sum_{k \neq j} \lambda_{k|j}(b) \Pi_j^{\text{lose}}(k) \right] = \bar{G}_j(b) Y^V. \quad (\text{B.1})$$

The weight is the share of the marginal loss in j 's winning probability attributable to rival k . The probability that k wins conditional on any loss by j is instead $P_k(b) / (1 - \bar{G}_j(b))$.

Modeled incremental profit and the pseudo-cost Expanding the bracket of (B.1) with the payoffs above,

$$\Pi_j^{\text{win}}(b) - \sum_{k \neq j} \lambda_{k|j} \Pi_j^{\text{lose}}(k) = \underbrace{(w_j - c_j) \Delta \mathcal{Q}_j}_{\equiv v_j} - (w_j - b) Y^V,$$

where the λ -weighted incremental quantity from holding the contract is

$$\Delta \mathcal{Q}_j \equiv \mathcal{Q}_j(w, j) - \sum_{k \neq j} \lambda_{k|j} \mathcal{Q}_j(w, k).$$

The *modeled incremental profit* $v_j = (w_j - c_j) \Delta \mathcal{Q}_j$ is the prize, in dollars, from holding the contract relative to the expected rival winner. Splitting the quantity into voucher and private parts under Assumption 4.3 gives $\Delta \mathcal{Q}_j = Y^V + \Delta \bar{Y}_j^{\text{priv}}$, with the latter term being the expected private market spillover of (20).

Absent type η , the pseudo-cost is the net price at which the rebate outlay equals the prize, $(w_j - \tilde{c}_j^0) Y^V = v_j$. With $\eta_j = \mu_j^\eta + \tilde{\eta}_j$ as in Assumption 4.4, the effective bidding objective replaces the winning payoff by $\Pi_j^{\text{win}} - \eta_j Y^V$. Its pseudo-cost is

$$\tilde{c}_j = \tilde{c}_j^0 + \eta_j, \quad \tilde{c}_j^0 = w_j - (w_j - c_j) \frac{\Delta \mathcal{Q}_j}{Y^V}.$$

Substituting $v_j = (w_j - \tilde{c}_j^0)Y^V$ into the first-order condition for this effective objective makes the payoff difference $(b - \tilde{c}_j)Y^V$. The resulting procurement condition is $b - \tilde{c}_j(b) = \bar{G}_j(b)/\bar{g}_j(b)$. At a fixed bid, the pseudo-cost rises one for one with η_j ; its dependence on the bid comes through the weights in $\Delta\mathcal{Q}_j(b)$. Inverting at an interior bid with $\Delta\mathcal{Q}_j \neq 0$ gives

$$w_j - [r_j + \bar{G}_j(b_j)/\bar{g}_j(b_j)] \frac{Y^V}{\Delta\mathcal{Q}_j} = c_j + \eta_j \frac{Y^V}{\Delta\mathcal{Q}_j}. \quad (\text{B.2})$$

Appendix Section C.6 recovers c_j from wholesale pricing and then estimates the mean and variance of η_j from bids. The weights $\lambda_{k|j} = h_k/(h_k + h_l)$ depend on rivals' equilibrium bids, so the bid distributions and the pseudo-costs must be mutually consistent.

The equilibrium bid distribution This characterization connects bid distributions to the distribution of the type η . Let firm j use a strictly increasing strategy $b = b_j(\eta_j)$ on its interior bid support, with inverse $\varphi_j = b_j^{-1}$. Thus a firm bidding b has $\eta_j = \varphi_j(b)$, of which only $\varphi_j(b) - \mu_j^\eta$ is private information. Write $F_{\eta,j}(x) = F_{\tilde{\eta},j}(x - \mu_j^\eta)$ for the distribution including the publicly known mean, suppressing the market and time indices. The induced bid distribution on this interior support is

$$G_j(b) = F_{\eta,j}(\varphi_j(b)), \quad g_j(b) = f_{\eta,j}(\varphi_j(b))\varphi_j'(b), \quad h_j(b) = \frac{g_j(b)}{1 - G_j(b)}.$$

Since $\bar{g}_j/\bar{G}_j = \sum_{k \neq j} h_k$, the first-order condition is

$$\sum_{k \neq j} h_k(b) = \frac{1}{b - \tilde{c}_j(b)}, \quad j = 1, 2, 3, \quad (\text{B.3})$$

with the pseudo-cost and the weights built from the same bid distribution, and $\eta_j = \varphi_j(b)$:

$$\tilde{c}_j(b) = w_j - (w_j - c_j) \frac{\Delta\mathcal{Q}_j(b)}{Y^V} + \varphi_j(b), \quad \lambda_{k|j}(b) = \frac{h_k(b)}{\sum_{n \neq j} h_n(b)}.$$

With $G_j = F_{\eta,j}(\varphi_j)$, (B.3) gives three coupled equations in inverse bidding strategies. It has the structure of the asymmetric first-price system of Maskin and Riley (2000), with the additional dependence of the pseudo-cost on the weights in $\Delta\mathcal{Q}_j(b)$. These are interior first-order conditions; candidate strategies must also satisfy global best-response conditions and the bid bound $b_j \leq w_j$.

If a firm's losing payoff is the same whichever rival wins, its weighted losing payoff does not

depend on $\lambda_{k|j}$. Under full symmetry of costs, wholesale prices, downstream sales, and type distributions, a symmetric equilibrium additionally has $\lambda_{k|j} = 1/2$. Then the private spillover is a common shift in pseudo-cost and (B.3) reduces to the symmetric three-bidder case of $b - \tilde{c} = (1 - G(b))/(2g(b))$. In our application, losing payoffs depend on the winner's identity, so the weights matter for the incentive to outbid a stronger or a weaker rival.

C Estimation Appendix

C.1 Demand

We estimate the demand system with the software implementation of Conlon and Gortmaker (2020). A market is a ZIP3-month, and the products are brand-store pairs, so that consumers' choice sets of big-box stores include all stores in the ZIP3 and the outside good. The linear characteristics are the brand fixed effects $\bar{\xi}_j$ and the WIC indicator, with the store intercepts δ_s absorbed as store fixed effects. Price heterogeneity loads entirely on income: the price coefficient interacts with $D_i = \log(\text{income}_i/10,000)$, integrated over the American Community Survey's ZIP3 household-income brackets. The endogenous regressors are the price and its income interaction. We instrument them with the national CME Class III milk futures price, a shifter of the formula input cost. Milk prices shift production costs, which shift retail prices through the markup over wholesale prices, but not the demand shocks. We drop 2007–2009, when milk future prices moved with macroeconomic condition of the financial crisis rather than formula cost.

C.2 Consumer Arrivals and Market Size

The retailer model starts from store-level daily arrival distributions of private and voucher consumers, $N_{st} \sim \text{Binomial}(N_s, \rho)$. The data record prepared ounces of formula sold. Therefore, we calibrate $\bar{y}$ using the average number of prepared ounces a consumer buys per formula transaction.

From the Consumer Panel's formula transactions we calibrate the arrival probability as the reciprocal of the mean duration between a household's successive formula purchases. We calibrate $\rho^V = 0.024$ per day (a purchase roughly every six weeks) for households currently redeeming WIC vouchers. We use the private-market consumers' purchase frequency of 0.021 to derive the market size as detailed below. The corresponding basket size is $\bar{y}^V = 387$ prepared ounces per voucher transaction and $\bar{y} = 231.4$ prepared ounces per private market transaction.

C.2.1 Inventory Data

We then calibrate the number of potential voucher and private-market consumers in each store. Voucher consumers do not choose the outside option, so we set N_s^V to the store's mean daily voucher volume divided by $\rho^V \bar{y}^V$. Private-market consumers may choose the outside option, so we set N_s from the upper tail of the store's daily private-market sales.⁵⁶ Let P_s be the larger of the store's maximum daily private-market sales and its 95th percentile divided by 0.75.⁵⁷ We set $N_s = P_s / (0.021, \bar{y})$, where 0.021 is the daily purchase frequency of private-market consumers. Expected daily arrivals are then $N_s \rho = (\rho / 0.021), P_s / \bar{y}$. Because arrivals enter only through $N_s \rho$ and we estimate ρ , the factor 0.021 is a normalization. We therefore report $\rho / 0.021$, the ratio of expected daily arrivals to peak daily transactions $P_s / \bar{y}$.

C.2.2 Scanner Data

The Scanner Data reports weekly private-market sales. To apply the daily market-size calibration, we use the Inventory Data to estimate a correction for weekly aggregation. At each frequency, define P as the larger of maximum sales and the 95th percentile divided by 0.75. We approximate the corresponding daily sales level by

$$\hat{P}_d = \frac{P_w}{7} m(\text{size}),$$

where size is mean weekly private-market sales and

$$m(\text{size}) = \frac{1 + z_d \sqrt{g^2 / \text{size} + \text{CV}_\infty^2}}{1 + z_w \sqrt{(g^2 / \text{size} + \psi \text{CV}_\infty^2) / 7}}. \quad (\text{C.1})$$

This specification allows variation from individual purchases to decline with store size while retaining variation shared across consumers.

We calibrate the five constants using the 544 Inventory Data stores. Regressing the squared coefficient of variation of daily private-market sales on inverse store size gives slope g^2 and intercept CV_∞^2 . The corresponding weekly regression's intercept determines ψ , equal to seven times the weekly-to-daily intercept ratio. At each frequency, z_d and z_w are the across-store medians of $(P - \mu) / \sigma$, where μ and σ are mean sales and their standard deviation.

⁵⁶The 2022 formula-shortage period is excluded throughout so that the quantiles are not affected by sales patterns in an abnormal period.

⁵⁷Among U.S. infants born in 2019, 24.9% were exclusively breastfed through six months (Centers for Disease Control and Prevention, 2022).

The calibrated values are $g = 45.8$, $CV_\infty = 0.73$, $\psi = 1.75$, $z_d = 6.2$, and $z_w = 3.4$. We convert the corrected daily sales level into market size using the purchase-frequency and transaction-size adjustments described above.

C.3 The Wholesale Price Process

For each brand and store, we fix the retail price at its average list price, excluding promotional discounts.

Wholesale price variation For each delivery, we subtract the national posted log wholesale price from the store’s log wholesale price, then subtract the average difference for that brand across deliveries to all stores. Denote this adjusted log price by $\hat{\omega}_{jst}$:

$$\hat{\omega}_{jst} = \log w_{jst} - \log w_{jt} - \bar{\Delta}_j, \quad (\text{C.2})$$

where w_{jst} is the average wholesale price the store paid, w_{jt} is the national posted wholesale price, and $\bar{\Delta}_j$ is the average of $\log w_{jst} - \log w_{jt}$ across all deliveries of brand j to all stores. The standard deviation of $\hat{\omega}_{jst}$ is $\hat{\sigma}_\omega = 0.207$. We calculate correlations between adjusted log prices for deliveries of the same brand to the same store, grouping delivery pairs by the number of days between them. Appendix Table C.1 reports these correlations.

Table C.1: Wholesale Price Correlations by Time Between Deliveries

Horizon	Mean distance (days)	Pairs	Correlation	Fitted
8–30 days	18.7	1,598,449	0.261	0.239
1–3 months	59.1	4,301,967	0.239	0.222
3–6 months	133.6	5,792,214	0.194	0.193
6–12 months	268.0	9,886,838	0.137	0.149
Beyond a year	622.1	18,072,575	0.078	0.076

Notes: Pairs are two deliveries of the same brand to the same store, at the given calendar distance. Correlation is the Pearson correlation between the adjusted log wholesale prices defined in (C.2) on the earlier and the later delivery, computed across the horizon’s pairs. Fitted values are $s e^{-\zeta_w(1-\rho_w)\bar{\tau}}$ at the horizon’s mean distance $\bar{\tau}$, with $\zeta_w = 1/30$ and the fitted $\rho_w = 0.943$.

Autocorrelation implied by the process We calibrate wholesale-price persistence using the relationship

$$\text{Corr}(\hat{\omega}_{jst}, \hat{\omega}_{js,t+\tau}) \approx s e^{-\zeta_w(1-\rho_w)\tau}, \quad (\text{C.3})$$

where τ is the number of days between deliveries and s is the fitted correlation amplitude. The rate at which correlation declines identifies $\zeta_w(1-\rho_w)$. We fix the daily update probability at $\zeta_w = 1/30$ and estimate ρ_w from this relationship.

Fit We regress the log correlation in each interval on the mean number of days between deliveries, weighting by the number of delivery pairs. With $\zeta_w = 1/30$, the fitted decay rate implies $\rho_w = 0.943$. We use the fitted amplitude to calibrate the transitory standard deviation as $\sigma_e = \hat{\sigma}_\omega \sqrt{1-s} = 0.179$.

Estimating $\sigma_{\tilde{\omega}}$ We estimate $\sigma_{\tilde{\omega}}$ by the method of simulated moments. The identifying moment is the correlation of log wholesale prices across deliveries of the same brand to the same store no more than seven days apart.

C.4 Retailer Costs

Ordering shocks Ordering and waiting receive independent Type-I extreme-value payoff shocks with scale σ_v . Let $\bar{w}_{jst} = \mathbb{E}[w_{jst} \mid \tilde{\omega}_{jst}]$. The net value of ordering is

$$\Delta_{js}(q_j, \tilde{\omega}_j) = [V_{js}(S_j, \tilde{\omega}_j) - K_s - \bar{w}_{jst}(S_j - q_j)] - V_{js}(q_j, \tilde{\omega}_j). \quad (\text{C.4})$$

The probability of ordering is

$$\text{Pr}(\text{order} \mid q_j, \tilde{\omega}_j) = \frac{1}{1 + \exp\{-\Delta_{js}(q_j, \tilde{\omega}_j)/\sigma_v\}}. \quad (\text{C.5})$$

Moment construction At each Inventory Data store, we match 32 moments. Nine are computed separately for each brand: the mean and coefficient of variation of shipments and sales; the mean and variance of log wholesale prices on delivery days; ordering and zero-sales frequencies; and mean shipments divided by the ordering frequency. The remaining five are the mean and coefficient of variation of WIC-brand private-market sales, the WIC brand's share of private-market sales, total private-market sales, and the correlation of log wholesale prices across same-brand deliveries no more than seven days apart.

The 20 estimated parameters are reported in Table 2, Appendix Table C.2, and Panel (b) of Table 1.

Table C.2: Store-Level Intercept Coefficients

Parameter	Description	Value
$\gamma_{0,0}$	Intercept	0.645
$\gamma_{0,1}$	Log median income	-0.424
$\gamma_{0,2}$	Log population density	-0.309
$\gamma_{0,3}$	Chain 3 fixed effect	-1.72
$\gamma_{0,4}$	Chain 4 fixed effect	0.687
$\gamma_{0,5}$	Chain 5 fixed effect	0.046
$\gamma_{0,6}$	Share of Black residents	5.11
$\gamma_{0,7}$	Share of children under five	8.01

Notes: Coefficients of the store-level intercept $\delta_s = \gamma'_0 X_s^0$, estimated jointly with the parameters in Table 2. As in (23), δ_s enters the utility of the inside goods: a positive coefficient raises the appeal of buying formula at the store.

C.5 Computing Sales under Inventory Constraints

This appendix develops the algorithm that computes the distribution of sales under the sequential choice process of Section 4.1 and the Poisson arrival specification of Section 5.2.

Fix a demand draw ξ , holding prices and shocks fixed so that the choice probabilities $d_j(a)$ depend only on the available set a . Consumers arrive as a Poisson stream of rate μ (the expected arrivals $N_s \rho$ of Section 5.2), and each buys one transaction of $\bar{y}$ ounces of the brand it chooses. Each brand's demand in ounces lies on the grid $\{0, \bar{y}, 2\bar{y}, \dots\}$. Write $M_j = \lceil Q_j / \bar{y} \rceil - 1$ for the largest buyer count the inventory serves without running out, $f(\cdot; \mu)$ and $F(\cdot; \mu)$ for the Poisson PMF and CDF, and $\mathbb{E}_\mu[\cdot]$ for an expectation taken with $D/\bar{y} \sim \text{Pois}(\mu)$.

The algorithm proceeds in rounds of demand allocation and reallocation, indexed by $\tau \geq 1$. Let $a^{(\tau)}$ denote the available set in round τ , with $a^{(1)} = \{j : Q_j > 0\}$. Several brands may exhaust their inventories within a round.

Round 1. For brands $j \in a^{(1)}$, demands are independently distributed with $D_j^{(1)} / \bar{y} \sim \text{Pois}(\mu_j^{(1)})$ at rate $\mu_j^{(1)} = \mu d_j(a^{(1)})$. Each brand's stockout probability is $P_j^{(1)} = \mathbb{P}(D_j^{(1)} \geq Q_j) = 1 - F(M_j; \mu_j^{(1)})$. For any subset $B \subseteq a^{(1)}$, let $O_B^{(1)}$ denote the event that exactly the brands in B

stock out in round 1. Since the round-1 demands are being approximated as independent,

$$\mathbb{P}(O_B^{(1)}) = \prod_{j \in B} P_j^{(1)} \prod_{k \in a^{(1)} \setminus B} (1 - P_k^{(1)}).$$

If no brand runs out, which happens with probability $\prod_j (1 - P_j^{(1)})$, the recursion terminates and each brand's sales is a realization of its demand. If all brands run out, the recursion also terminates. Otherwise, conditional on $O_B^{(1)}$ for a nonempty proper subset B , each $j \in B$ has a pool of unsatisfied consumers, the mean size of which is

$$\bar{N}_j^{(1)} = \mathbb{E} \left[\frac{D_j^{(1)} - Q_j}{\bar{y}} \mid D_j^{(1)} \geq Q_j \right] = \frac{\mathbb{E}_{\mu_j^{(1)}}[(D - Q_j)^+]}{\bar{y} P_j^{(1)}}.$$

The recursion advances to round 2, whose state depends on the subset B that triggered the branch. Conditional on $O_B^{(1)}$, the available set is $a^{(2)}(O_B^{(1)}) = a^{(1)} \setminus B$ and the round-2 arrival stream is a Poisson stream at rate $\sum_{j \in B} \bar{N}_j^{(1)}$,⁵⁸ so a brand $k \in a^{(2)}(O_B^{(1)})$'s round-2 demand increment is an independent Poisson count at rate $(\sum_{j \in B} \bar{N}_j^{(1)}) d_k(a^{(2)}(O_B^{(1)}))$. Its cumulative demand up to and including round 2 is then $\bar{D}_k^{(2)}(O_B^{(1)})/\bar{y} \sim \text{Pois}(\mu_k^{(2)}(O_B^{(1)}))$ at the cumulative rate $\mu_k^{(2)}(O_B^{(1)}) = \mu_k^{(1)} + (\sum_{j \in B} \bar{N}_j^{(1)}) d_k(a^{(2)}(O_B^{(1)}))$.

Round τ . Round τ is reached along some history of stockouts $\ell^{(\tau)} = (B^{(1)}, B^{(2)}, \dots, B^{(\tau-1)})$ inherited from prior rounds. The available set $a^{(\tau)}(\ell^{(\tau)}) = a^{(1)} \setminus \bigcup_{r < \tau} B^{(r)}$ and the cumulative Poisson rates $\{\mu_j^{(\tau)}(\ell^{(\tau)})\}_{j \in a^{(\tau)}(\ell^{(\tau)})}$ are built up by the previous rounds of recursions over the unsatisfied consumers at stockout events. Suppressing $\ell^{(\tau)}$ from the arguments where no ambiguity arises, the stockout probability of product j up until round τ is $P_j^{(\tau)} = \mathbb{P}(\bar{D}_j^{(\tau)} \geq Q_j) = 1 - F(M_j; \mu_j^{(\tau)})$. Conditional on j having stayed available through round $\tau - 1$, its conditional probability of running out during this round is

$$h_j^{(\tau)} = \frac{P_j^{(\tau)} - P_j^{(\tau-1)}}{1 - P_j^{(\tau-1)}},$$

⁵⁸With income-varying price sensitivity, this treats the rationed consumers as fresh population draws rather than tracking the rationed pool's income composition.

with the convention $P_j^{(0)} = 0$, so $h_j^{(1)} = P_j^{(1)}$. For each subset $B \subseteq a^{(\tau)}$, the event $O_B^{(\tau)}$ that exactly the brands in B stock out in round τ has probability

$$\mathbb{P}(O_B^{(\tau)} | \ell^{(\tau)}) = \prod_{j \in B} h_j^{(\tau)} \prod_{k \in a^{(\tau)} \setminus B} (1 - h_k^{(\tau)}).$$

The expected number of consumers unable to buy brand j , conditional on its inventory being exhausted in round τ , is

$$\begin{aligned} \bar{N}_j^{(\tau)} &= \mathbb{E} \left[\frac{\bar{D}_j^{(\tau)} - Q_j}{\bar{y}} \middle| \bar{D}_j^{(\tau)} \geq Q_j, \bar{D}_j^{(\tau-1)} < Q_j \right] \\ &= \frac{\mathbb{E}_{\mu_j^{(\tau)}}[(D - Q_j)^+] - \mathbb{E}_{\mu_j^{(\tau-1)}}[(D - Q_j)^+] - \bar{y}(\mu_j^{(\tau)} - \mu_j^{(\tau-1)})P_j^{(\tau-1)}}{\bar{y}(P_j^{(\tau)} - P_j^{(\tau-1)})}. \end{aligned}$$

If B is empty or equals $a^{(\tau)}$, the recursion terminates. Otherwise, conditional on $O_B^{(\tau)}$, the recursion advances to round $\tau + 1$ along the extended history $\ell^{(\tau+1)} = (\ell^{(\tau)}, B)$, with available set $a^{(\tau+1)} = a^{(\tau)} \setminus B$ and cumulative rates

$$\mu_k^{(\tau+1)} = \mu_k^{(\tau)} + \left(\sum_{j \in B} \bar{N}_j^{(\tau)} \right) d_k(a^{(\tau+1)}).$$

Each continuing round removes at least one brand, so the recursion terminates within J rounds.

Averaging over sequences. A *sequence* ℓ is a nested chain of survivor sets $a^{(1)} \supseteq a^{(2)} \supseteq \dots$, occurring with probability

$$\mathbb{P}(\ell) = \prod_{j \text{ survives}} F(M_j; \mu_j^{(\ell)}) \prod_{k \text{ sold out in } \tau_k} (P_k^{(\tau_k)} - P_k^{(\tau_k-1)}).$$

Along a sequence ℓ each brand has either sold out or survived. A brand that sold out sells $Y_k = Q_k$ with leftover $q'_k = 0$. A brand j that survived carries one Poisson count at its final rate $\mu_j^{(\ell)}$, conditioned on having stayed at or below M_j buyers; truncating and renormalizing the Poisson and using $x f(x; \mu) = \mu f(x - 1; \mu)$,

$$\mathbb{E}[Y_j | \ell] = \bar{y} \mu_j^{(\ell)} \frac{F(M_j - 1; \mu_j^{(\ell)})}{F(M_j; \mu_j^{(\ell)})},$$

which lies below the unconditional mean $\bar{y} \mu_j^{(\ell)}$ because not running out is evidence that demand was low. Averaging over sequences by the law of total expectation gives the objects that serve as inputs to the retailer's problem,

$$\begin{aligned}\mathbb{E}[Y_j] &= \sum_{\ell: j \text{ sold out}} \mathbb{P}(\ell) Q_j + \sum_{\ell: j \text{ survives}} \mathbb{P}(\ell) \bar{y} \mu_j^{(\ell)} \frac{F(M_j - 1; \mu_j^{(\ell)})}{F(M_j; \mu_j^{(\ell)})}, \\ \mathbb{P}(Y_j = Q_j) &= \sum_{\ell: j \text{ sold out}} \mathbb{P}(\ell), \\ \mathbb{E}[V(q'_j)] &= \left(\sum_{\ell: j \text{ sold out}} \mathbb{P}(\ell) \right) V(0) + \sum_{\ell: j \text{ survives}} \mathbb{P}(\ell) \sum_{x=0}^{M_j} \frac{f(x; \mu_j^{(\ell)})}{F(M_j; \mu_j^{(\ell)})} V(Q_j - \bar{y} x).\end{aligned}$$

The stockout probability $\mathbb{P}(Y_j = Q_j)$ is the object in the WIC stocking penalty Λ_j of (8), and the store's expected total revenue is $\sum_k p_k \mathbb{E}[Y_k]$.

C.6 Manufacturer Costs and Bidding Behavior

The projection weights NielsenIQ's sample of stores relative to the universe varies by place and retail channel. For accurate national accounting of quantities, each store must be projected according to appropriate sampling weights. We build the universe from the census of SNAP-authorized retailers from the USDA, restricted to the formula-relevant food, drug, and mass-merchandise channels, with each store active in the years spanned by its SNAP authorization window. Within each ZIP3-channel-year cell we compare the n_z formula-carrying Scanner Data stores with the cell's universe of $|U_z|$ stores and compute the cell's *coverage share* $\chi_z = n_z / |U_z|$. The shares are computed for each calendar year; auctions before 2007 are assigned their cell's 2007 share due to a data-availability issue for the SNAP-authorized retailer dataset. A sampled store's projection weight is the reciprocal of its cell's coverage share, $\phi_s = 1/\chi_z$, and weighting sampled-store quantities by ϕ_s in (C.6) sums the sample up to the universe.

Aggregating orders We aggregate the downstream model's predicted orders to each auction market m , a state or purchasing alliance in a given contracting period. For each candidate WIC winner a , expected orders of brand j are

$$\mathcal{Q}_{jm}(w, a) = \sum_{s \in m} \phi_s \mathbb{E} \left[\sum_t (Q_{jst} - q_{jst}) \right], \quad (\text{C.6})$$

where m indexes the auction market, a is the candidate WIC winner, and ϕ_s projects sampled stores to the retailer population in their ZIP3, channel, and year. For small-box stores, the expectation is over orders from the inventory model. Under complete availability, big-box orders equal demand-system sales and require no inventory simulation. Thus, the aggregation includes expected orders for both store formats. We value both at the posted national wholesale price w_j .

Retail prices and manufacturer profits In small-box stores, the retail price p_{js} is fixed at the store's average list price within each inventory problem. When we vary the national wholesale price w_j to calculate manufacturers' demand elasticities, we adjust p_{js} proportionally, holding m_{js} fixed, and solve the inventory problem again. Daily shocks to the store's acquisition costs do not change its retail price.

In big-box stores, demand estimation uses observed average retail prices for each brand-store-month. To implement the proportional price response in Section 4.3, we construct $m_{jst} = p_{jst}/w_{jt}$ from these retail prices and the national wholesale price. When calculating manufacturers' demand elasticities, we hold m_{jst} fixed as we vary wholesale prices.

For a fixed WIC allocation, a wholesale price change affects big-box quantities through consumer demand and small-box quantities through both demand and stocking. Changing the WIC allocation at fixed wholesale prices changes big-box demand through the WIC indicator, with no additional retail price adjustment.

When brand j holds the contract, its expected profit from store s is

$$\pi_{js}^{\text{win}}(r_j) = \mathbb{E} \left[\sum_t \{ (w_j - c_j)(Q_{jst} - q_{jst}) - r_j Y_{jst}^V \} \right]. \quad (\text{C.7})$$

For big-box stores, we replace orders in this expression with sales. Summing these profits with the projection weights defined above gives the modeled market-level profits under each possible winner.

Step 1 for auction: the bid distributions Given production costs, the bidding estimation has two steps. First, we fit each firm's net-bid distribution jointly across markets as a lognormal, upper-censored at w_j :

$$\log b_{jm} = \min\{\mu_{jm} + \sigma_{j\ell(m)}\varepsilon_{jm}, \log w_j\}, \quad \mu_{jm} = \alpha_j + \delta_{\ell(m)} + Z'_{jm}\gamma, \quad \varepsilon_{jm} \sim N(0, 1). \quad (\text{C.8})$$

Based on the equilibrium bid-distribution characterization in Appendix Section B.1, we include in Z_{jm} the increment from winning and the spread across rival winners in firm j 's order revenue and order quantity, along with the posted wholesale price w_j and voucher volume Y_m^V . Each covariate is divided by its standard deviation. Market m belongs to period $\ell(m) \in \{2006\text{--}2009, 2010\text{--}2015\}$, the same split used for types in Step 2: the location shifts by δ_ℓ and the scale $\sigma_{j\ell}$ is specific to a brand and a period. Writing $\mathcal{P} = \{(j, m) : r_{jm} > 0\}$ for the interior bids, the log-likelihood over all firm–market bids is

$$\mathcal{L}_1(\{\alpha_j\}, \{\delta_\ell\}, \gamma, \{\sigma_{j\ell}\}) = \sum_{(j,m) \in \mathcal{P}} \log \frac{1}{\sigma_{j\ell(m)}} \phi\left(\frac{\log b_{jm} - \mu_{jm}}{\sigma_{j\ell(m)}}\right) + \sum_{(j,m) \notin \mathcal{P}} \log \left[1 - \Phi\left(\frac{\log w_j - \mu_{jm}}{\sigma_{j\ell(m)}}\right) \right]. \quad (\text{C.9})$$

An interior bid contributes the density; a zero-rebate bid $b_{jm} = w_j$, $r_{jm} = 0$, contributes the upper-tail mass that its latent optimal bid exceeds the wholesale price. From the fitted distributions the win-probability objects follow: firm j 's probability of winning at b is $\bar{G}_j(b) = \prod_{k \neq j} (1 - G_k(b))$, the hazards $h_j = g_j / (1 - G_j)$ are nonnegative by construction, the inverse hazard is $\bar{G}_j / \bar{g}_j = 1 / \sum_{k \neq j} h_k$, and the weights on rival winners at the margin are $\lambda_{k|j} = h_k / (h_k + h_l)$ with l the remaining rival, all evaluated at firm j 's own bid b_j . Appendix Table C.3 reports the estimates.

Step 2 for auction: types We estimate the mean and variance of the bidder's type separately for 2006–2009 and 2010–2015. From the fitted bid distributions we form the increment in firm j 's order quantity from winning, $\Delta \mathcal{Q}_j = \mathcal{Q}_j(w, j) - \sum_{k \neq j} \lambda_{k|j} \mathcal{Q}_j(w, k)$. We estimate the cost share θ_j from the pricing problem and hold it fixed here. National posted wholesale prices change over time, so the cost $c_j = \theta_j w_j$ entering a bid's optimality condition is evaluated at the posted price of that market's contracting period. Holding c_j fixed, the bid's first-order condition pins down the bidder's type,

$$\eta_j = (w_j - c_j) \frac{\Delta \mathcal{Q}_j}{Y^V} - \left(r_j + \frac{\bar{G}_j(b_j)}{\bar{g}_j(b_j)} \right). \quad (\text{C.10})$$

A firm that bids at zero rebate reveals that offering no rebate is optimal, which happens when $\eta_j \geq \bar{\eta}_j$. To value a deviation to a net price $b < w_j$, note that the contract changes hands only when the lowest rival net price M falls in (b, w_j) . Conditional on that, rival k takes the contract

Table C.3: Bid Distribution Estimates

	1 SD	2006–2009	2010–2015
<i>Firm intercepts α_j and period shift</i>			
Abbott		−0.665 (1.115)	
Mead Johnson, relative to Abbott		−0.192 (0.175)	
Nestlé, relative to Abbott		−0.183 (0.314)	
2010–2015, relative to 2006–2009		−0.403 (0.251)	
<i>Downstream fundamentals γ, per standard deviation</i>			
Order revenue, increment from winning	\$174k/day	0.160 (0.469)	
Order revenue, spread across rival winners	\$10.3k/day	0.408 (0.312)	
Order quantity, increment from winning	1.15M prepared oz./day	−0.777 (0.687)	
Order quantity, spread across rival winners	94.5k prepared oz./day	−0.326 (0.264)	
Posted wholesale price w_j	\$0.0144 per prepared ounce	−0.253 (0.121)	
Voucher volume Y_m^V	796k prepared oz./day	0.393 (0.458)	
<i>Scale $\sigma_{j,\ell}$</i>			
Abbott		0.520 (0.076)	1.135 (0.154)
Mead Johnson		0.405 (0.061)	1.102 (0.166)
Nestlé		1.983 (0.551)	0.906 (0.131)
Firm–auction bids (N)		75	87
Of which at the reserve (censored)		13	5
R^2 among positive-rebate bids		0.269	

Notes: Standard errors in parentheses, from the observed information of (C.9). The outcome is the log net bid $\log b_{jm}$; a bid at the reserve ($r_{jm} = 0$) enters as censored above at $\log w_j$. Intercepts and fundamentals are common to both periods and are reported once; the scale is estimated separately for each brand and period. Nestlé’s 2006–2009 scale rests on 12 interior bids against 13 at the reserve, which is why its standard error is the largest in the table. The fundamentals are firm j ’s expected order revenue and order quantity over the contract period under each candidate winner, aggregated as in (C.6): the increment is the own-win value minus the average of the two rival-winner values, and the spread is the absolute difference between the two rival-winner values. Each fundamental is scaled by its standard deviation across firm–auction cells, so its coefficient is the effect of a one-standard-deviation change on the mean of the log net bid; the “1 SD” column reports that magnitude in natural units, with revenue and quantity measured per retained contract day and volumes in prepared ounces. R^2 compares the fitted mean with the log net bids of the positive-rebate observations.

with probability

$$\Lambda_k(b) = \frac{\int_b^{w_j} \bar{G}_j(m) h_k(m) dm}{\bar{G}_j(b) - \bar{G}_j(w_j)} = \Pr(k \text{ wins} \mid b < M < w_j).$$

The threshold is the best interior deviation,

$$\begin{aligned} \bar{\eta}_j &= \sup_{b < w_j} \left[u_j^c(b) - c_j \frac{\Delta \mathcal{Q}_j^c(b)}{Y^V} \right], \\ \Delta \mathcal{Q}_j^c(b) &= \mathcal{Q}_j(w, j) - \sum_{k \neq j} \Lambda_k(b) \mathcal{Q}_j(w, k), \\ u_j^c(b) &= w_j \frac{\Delta \mathcal{Q}_j^c(b)}{Y^V} - (w_j - b) \frac{\bar{G}_j(b)}{\bar{G}_j(b) - \bar{G}_j(w_j)}, \end{aligned} \tag{C.11}$$

evaluated at the cost $c_j = \theta_j w_j$.

For each brand and period we estimate $(\mu_{j,\ell}^\eta, \sigma_{\eta,j,\ell})$ by maximizing

$$\mathcal{L}_2(\mu_{j,\ell}^\eta, \sigma_{\eta,j,\ell}) = \sum_{m \in \ell: r_{jm} > 0} \log \frac{1}{\sigma_{\eta,j,\ell}} \phi\left(\frac{\eta_{jm} - \mu_{j,\ell}^\eta}{\sigma_{\eta,j,\ell}}\right) + \sum_{m \in \ell: r_{jm} = 0} \log \left[1 - \Phi\left(\frac{\bar{\eta}_{jm} - \mu_{j,\ell}^\eta}{\sigma_{\eta,j,\ell}}\right) \right]. \tag{C.12}$$

A zero-rebate bid enters as a censored term: it reveals only that the firm's η draw lies at or above the threshold $\bar{\eta}_{jm}$.

Appendix Table C.4 reports the estimates. For Abbott and Mead Johnson the mean $\mu_{j,\ell}^\eta$ is positive in both periods, so their mean types raise their pseudo-costs relative to those implied by production costs and modeled spillovers alone. Nestlé turns negative in 2010–2015. Figure 6 in the main text compares production costs and effective costs in bidding with prices and bids.

Table C.4: Mean Type and Idiosyncratic Shock Scales

	Mean $\mu_{j,\ell}^\eta$		Scale $\sigma_{\eta,j,\ell}$	
	2006–09	2010–15	2006–09	2010–15
Abbott	6.13 (0.62)	5.57 (0.88)	3.10 (0.44)	4.67 (0.64)
Mead Johnson	7.39 (0.79)	7.49 (1.09)	3.95 (0.56)	5.82 (0.83)
Nestlé	6.15 (2.07)	−2.20 (0.54)	8.26 (1.95)	2.92 (0.40)

Notes: Standard errors in parentheses. Estimates of the mean $\mu_{j,\ell}^\eta$ and scale $\sigma_{\eta,j,\ell}$ of the unobserved type η by brand and period, from the censored maximum likelihood (C.12), in dollars per 90 voucher prepared ounces (one 13-oz can). These estimates exclude production costs and modeled private-market spillovers, which enter the pseudo-cost separately in (21). Zero-rebate bids enter censored: one for Abbott, three for Mead Johnson, and 14 for Nestlé.

D Counterfactual Appendix


For each allocation and set of wholesale prices, we solve small-box retailers’ dynamic stocking problems and simulate average daily outcomes under the resulting stationary policies. These outcomes, together with big-box demand and the bidder types, determine the type-adjusted gross profits in (26). We aggregate daily outcomes across stores and multiply by 365 to report annual spending and consumer-surplus differences.

Retail prices respond proportionally to wholesale-price changes as specified in Section 4.3. We hold the small-box multipliers m_{js} and big-box multipliers m_{jst} fixed across designs, using the construction in Appendix Section C.6.

D.1 Representative Stores and Markets

We evaluate the 25 procurement markets with contracts in force in 2016. The store population is the same as in the auction inversion: 23,799 small-box store–auction cells and 747 big-box ZIP3 markets (Appendix Section C.6). Each small-box cell uses wholesale costs anchored to observed retail prices through the markups κ_j , cost covariates based on store volume, and its projection weight ϕ_s . Its demand intercept δ_s is recovered to match observed sales at the estimated parameters. Solving every cell’s inventory problem under every design is computationally infeasible. We therefore solve the inventory problems for representative stores and weight their outcomes to represent the full population.

Figure D.1: K -Means Objective by Number of Types



Notes: Each panel plots the within-cluster sum of squares of the K -means clustering against the number of clusters K . Panel (a) clusters the small-box (store, auction) cells on the five store characteristics, pooled across strata; the marked point is the selected number of clusters, $K = 110$. Panel (b) clusters the big-box ZIP3 markets; the selected number is $K = 80$.

We cluster the small-box cells by K -means (Arthur and Vassilvitskii, 2007) on five store characteristics: the intercept δ_s , the level of the store's price multipliers m_{js} , store volume, voucher sales, and ZIP3 income. We cluster within groups defined by the WIC brand and the set of brands the store carries, so each representative has the same status quo WIC allocation as the stores it represents. For each of the 110 types, we select the sampled store nearest the cluster centroid. We assign its simulated outcomes to each member cell and apply that cell's projection weight ϕ_s .

For big-box stores, we use representative ZIP3 markets because consumers choose among all big-box stores in their ZIP3. Within each auction, we cluster markets on brand-level measures combining retail prices and demand elasticities, together with sales-weighted δ_s , market size, voucher volume, income, and total store count. This reduces the 747 ZIP3 markets to 80 representative markets. For each cluster, we select a market near the centroid with high NielsenIQ coverage. We assign its predicted shares and prices to each member market. We then calculate sales using the member market's own market size and voucher volume and apply its projection weights.

Appendix Figure D.1 plots the within-cluster sum of squares against the number of clusters. We select the number of representative stores and markets at the elbow of each curve.

D.2 Fiscal Cost and Market Outcomes

Appendix Table D.1 reports annual national outcomes at observed wholesale prices and with adjusted wholesale prices for the procurement designs in Section 7.

We calculate manufacturers' operating profits implied by the model as $\sum_{j,m}[(w_j - c_j)Y_{jm} - T_{jm}]$, where T_{jm} is the manufacturer's payment to the government under the relevant design. The bidder's type affects contract awards and payments but is excluded from this operating-profit measure, consistent with (14)–(15).

Table D.1: Fiscal Cost and Market Outcomes Across Procurement Designs

Design	Observed wholesale prices (w)					Adjusted wholesale prices (w^*)					Δp (%)
	Fiscal (\$M/yr)	Δ Surplus (\$M/yr)	Q_{WIC} (bil. prepared oz.)	Q_{Priv} (bil. prepared oz.)	Mfr. (\$M/yr)	Fiscal (\$M/yr)	Δ Surplus (\$M/yr)	Q_{WIC} (bil. prepared oz.)	Q_{Priv} (bil. prepared oz.)	Mfr. (\$M/yr)	
Status quo	367	—	7.58	14.99	1,680	—	—	—	—	—	—
<i>Changing allocation and payment rule</i>											
Sole-source (VCG)	45	+70	7.58	15.17	1,320	89	+10	7.58	14.88	1,360	+1.8
Profit-only bidding	326	+80	7.58	15.17	1,570	373	+20	7.58	14.89	1,610	+1.6
<i>Changing number of eligible and mandated brands</i>											
Dual sourcing (VCG)	366	+350	7.58	16.08	1,770	367	+340	7.58	16.07	1,770	−0.0
Triple sourcing	1,534	+450	7.58	16.27	2,940	1,491	+540	7.58	16.66	2,890	−2.6
<i>Changing stocking requirement</i>											
Current allocation, no mandate	293	−120	5.99	14.26	1,630	—	—	—	—	—	—
<i>Changing distribution arrangement</i>											
Government brand	588	−110	7.58	13.89	1,800	589	+10	7.58	14.55	1,790	−4.2
<i>Changing reimbursement rule</i>											
Coupon 7.5, brand-neutral, no penalty	387	−140	4.48	13.73	2,280	399	−50	4.51	14.16	2,230	−3.9
Coupon \$7.50, 1 brand	−660	+90	4.64	15.29	1,440	−692	−210	4.26	13.75	1,510	+12.2
Coupon \$7.50, 2 brands	−344	+350	5.73	16.08	1,890	−354	+190	5.60	15.33	1,960	+6.4
<i>Changing program scale</i>											
Sole-source, 1.5× WIC mkt.	403	−210	11.32	13.73	1,190	447	−250	11.32	13.55	1,230	+1.0

Notes: Dollar amounts and sales volumes are annualized national totals. Fiscal cost is WIC reimbursement net of manufacturer payments. “ Δ Surplus” is the change in private-market consumer surplus from the status quo; it excludes β_{WIC} and WIC beneficiaries. Q_{WIC} and Q_{Priv} are sales volumes in billions of prepared ounces. “Mfr.” reports operating profits implied by the model across WIC and private sales, net of production costs and manufacturer payments and excluding the bidder’s type. The first panel uses observed WIC and private sales, net of production costs and manufacturer payments with one or two eligible brands, the eligible brands and their award-time payments remain fixed. Δp is the percentage change in the national private-market transaction price from observed to adjusted wholesale prices within each design. Fiscal cost and Δ Surplus are the coordinates in Figure 7.

D.3 The Mean Type as an Opportunity Cost

The preferred specification removes μ_j^η in line with the interpretation in Section 6.3. Here we instead treat the mean as part of the firm's opportunity cost of supplying the contract and hold it fixed across counterfactual designs. We evaluate $\tilde{\Pi}_{jm}(a; \eta_{jm})$ at $\eta_{jm} = \mu_j^\eta + \tilde{\eta}_{jm}$ and recompute allocations and payments at observed wholesale prices. We keep the same idiosyncratic bidding shocks, demand parameters, and cost parameters. We then let manufacturers adjust wholesale prices under the resulting awards, as in the main analysis.

Retaining the mean type raises net spending under both sole and dual sourcing and changes their allocations (Appendix Table D.2). Nestlé, whose mean is negative, wins 13 of the 25 sole-source contracts instead of none, and private-market consumer surplus under sole sourcing falls below the status quo. With adjusted wholesale prices, dual sourcing generates approximately \$1.60 of additional private-market consumer surplus per additional fiscal dollar relative to the status quo: \$299 million in surplus for \$184 million in spending. Dual sourcing still raises private-market consumer surplus, but it now requires higher program spending than the status quo.

This comparison excludes WIC beneficiaries' welfare and changes in manufacturer profits. The exercise also assumes that the mean opportunity cost remains fixed when awards change. If it reflects capacity constraints, it may instead depend on the volume or composition of contracts.

Table D.2: Sensitivity to the Interpretation of the Mean of η

	Wholesale prices fixed				Wholesale prices adjusted			
	Net spending (\$M)		Δ Surplus (\$M)		Net spending (\$M)		Δ Surplus (\$M)	
	Removed	Included	Removed	Included	Removed	Included	Removed	Included
Status quo	-	367	-	0			—	
Sole-source (VCG)	45	170	+71	−98	89	251	+7	−68
Dual sourcing (VCG)	366	559	+346	+287	367	551	+344	+299

Notes: Annual national outcomes, in millions of dollars. Surplus differences are relative to the observed status quo and exclude β_{WIC} and WIC beneficiaries, as in the main analysis. Valuations use $\tilde{\eta}_{jm}$ with the mean removed and $\mu_j^\eta + \tilde{\eta}_{jm}$ with the mean included. Both VCG designs apply the perceived stockout penalty to every contracted brand. The left panel holds wholesale prices at observed values; the right panel lets them adjust under each design's awards.