Bundling to save: Analyzing package size choices in South African grocery stores*

Andrea Szabó
Department of Economics
University of Houston
aszabo2@uh.edu

July 5, 2026

Abstract

Storable goods like laundry detergent are important recurring purchases for low-income households. Differences in unit prices across package sizes create opportunities to save. Do households take advantage of these opportunities? I estimate a dynamic demand model in South Africa using country-wide scanner data from the leading grocery chain. I find that households are responsive to nonlinear prices, including opportunities to "bundle" by buying multiple small packages when this is cheaper than a large package. Thus, small purchases need not reflect constraints or consumer mistakes. The model helps evaluate counterfactuals related to the expansion of small-format chain stores to low-income areas.

*For useful comments and suggestions I would like to thank Andrew Foster, Douglas Gollin, Igal Hendel, Jean-François Houde, Michael Keane, Robert Miller, Aviv Nevo, Romans Pancs, Amil Petrin, Gergely Ujhelyi, Dean Yang, and seminar and conference audiences at Michigan, ITAM, MWIEDC 2023, EEA 2023, NASMES 2023, IIOC 2023, and the 2023 STATA Empirical Microeconomics Conference. This work was completed in part with resources provided by the Research Computing Data Core at the University of Houston.

1 Introduction

Storable household goods, such as laundry detergent or non-perishable food items, play an important role in the budgets of low-income households in developing countries. These goods are purchased repeatedly, are difficult to substitute or home-produce, and are often among the most expensive items households buy in a typical grocery store. Products are usually sold in multiple package sizes with different unit prices, which creates an opportunity for households to save (e.g., by buying larger packages). However, the common view is that low-income households do not take advantage of this saving opportunity due to financial constraints, transportation costs, or simply a failure to optimize by ignoring unit prices (Attanasio et al. 2013; Dillon, De Weerdt and O'Donoghue 2021; Abubakari et al. 2025).

Do households actually purchase "too many" small quantities, and if so, why? The answer to this question matters for at least three reasons. First, it provides insights into the decision-making of low-income households in a dynamic context. Households face a potentially complicated inventory problem involving storage and transportation costs, and different unit prices. Previously, such models were estimated on consumers in developed countries (e.g., Erdem, Imai and Keane (2003); Hendel and Nevo (2006)), and it is useful to confirm if standard demand models are also applicable outside this setting (Szabo 2015).

Second, households' potential savings on the everyday activity of grocery shopping are large. In my data, over a year, the average household that always buys the largest package vs. the smallest package will save 11.1 USD on laundry detergents alone. By comparison, Karlan et al. (2016) found that a door-to-door campaign to motivate households to make deposits into their savings accounts resulted in 26.4 USD saved on average.[1] From a policy perspective, an improved understanding of how households buy storable products could help design alternatives to more costly interventions aimed at increasing savings.

Third, understanding household choices regarding storable goods is especially relevant for the expansion of formal grocery chains to low-income and rural areas. Traditionally, these areas were served mainly by informal stores, and high travel costs limited consumers' access to formal supermarkets. In South Africa and elsewhere, expansion of formal stores is taking place through small-format stores (often built from shipping containers or truck beds) with significant space constraints. What product sizes should these stores offer, and what is the impact of bringing stores closer to consumers when they provide a limited selection of products?

I study demand for storable goods using a scanner dataset on laundry detergent purchases

[1] This campaign took place in the Philippines, which has the closest mean household income to the South African setting among the countries in the Karlan et al. (2016) study.

from South Africa's leading grocery chain. Observing actual, as opposed to self-reported, purchases and prices is unique in the literature on developing countries, and allows me to document two important facts. First, small purchases are less common than previously thought. Second, in contrast to common belief, larger packages do not always have lower unit prices (a phenomenon that does not appear to be limited to South Africa). Thus, "bundling" by buying several smaller packages instead of one large package can be a cost-minimizing strategy for the consumer.

Building on these facts, I estimate a dynamic model of consumer choice, which I identify using auxiliary information collected from a specially designed household survey. The results indicate that households take advantage of bundling opportunities, and provide evidence against the view that developing country consumers ignore unit prices and simply purchase products with the lowest package price. The estimates quantify the importance of transportation and inventory costs in shaping package-size choices, and counterfactual analyses provide insights into some of the welfare effects of the on-going expansion of formal, small-format grocery stores into low-income areas.

The paper's analysis requires high-quality data on prices and quantities sold for an entire product category, linked to demographic characteristics of consumers. While scanner datasets with these features are readily available for the US, I am not aware of another such dataset from a developing country. Cooperating with Unilever South Africa, I obtained a unique dataset of supermarket scanner data with country-wide coverage for a 16-month period, from July 2011 to October 2012. The data contain monthly information on laundry detergent sales at the store level from all 330 stores of South Africa's leading grocery chain, Shoprite. In these scanner data, I observe all brands of Unilever detergents sold - these account for approximately 90% of all detergent sales. I focus on the three main brands of powdered hand-wash detergent, by far the most common category. Each brand is sold in several package sizes, ranging from 250 g to 5 kg, for a total of 14 brand-size combinations over the period of study. Using the geo-coded location of the stores, I match the scanner data to household demographics from the South African census.

The data allows me to document two important facts. First, small purchases are not that common. The smallest, 250 and 500 g packages account for 23% of all purchases. The most popular package size is 2 kg, which is one of the largest available packages. Second, in contrast to common belief, large packages sometimes have *higher* unit prices than small packages.[2] Thus, consumers may profitably "bundle" by buying several smaller packages instead of one large package.

[2]A cursory look at the most commonly used US Nielsen scanner dataset, or at one's local grocery store, reveals that such pricing is not unique to South African supermarkets.

To understand consumer choices, I estimate a dynamic model based on Erdem, Imai and Keane (2003) and Hendel and Nevo (2006). In the model, consumers choose which package size and brand to purchase, and how much inventory to hold in each period. Estimation of the model needs to take into account several features of the setting: products are differentiated by brand and package size, bundling opportunities are present in some markets, and purchases are observed only at the market level. Applying existing estimation strategies to my data is not straightforward. In particular, bundling opportunities violate the standard discrete-choice assumption, and I present a method to estimate the share of small packages that were purchased as part of a bundle. This is based on the idea that estimates from markets with no bundling opportunities can be used to infer consumer demand on markets with bundling opportunities if such opportunities were absent.

The estimation of dynamic models with unobserved state variables also raises inherent difficulties for identification. Typical applications must address the fact that only households' purchases are observed, while consumption and inventory are unobserved.[3] To address this challenge, I field an original survey that directly collects information on households' detergent consumption, inventories, and purchases in my study area. This information combined with the scanner data helps identify the value function in consumers' dynamic programming problem.

The estimates show that accounting for bundling opportunities has important implications for the interpretation of observed market shares. I estimate that, in the median market, there are over three times as many households who purchase 5 kg of detergent as households who purchase a single 5 kg package. In turn, allowing for such bundling produces more sensible parameter estimates and a better model fit, which is consistent with households taking advantage of these opportunities. These findings run against the common view that households in developing countries simply purchase products with the lowest package price rather than the lowest unit price. As the case studied here illustrates, looking purely at market shares can be misleading: a consumer might optimally purchase multiple small packages because this bundle is less expensive than a larger package.

The estimates also quantify the importance of transportation and inventory costs in shaping package-size choices. Among the demographic and location characteristics observed in the data, car ownership emerges as especially important. In addition, I find that low-income households face higher inventory costs, which makes purchasing large packages more difficult compared to high-income households all else equal.

I use the estimates to study counterfactuals motivated by the on-going expansion of formal

[3]For example, Hendel and Nevo (2006) simulate an initial distribution of inventories and compute optimal consumption based on the model (see also Erdem, Imai and Keane (2003)).

grocery chains in South Africa and other developing countries. The first counterfactual captures one aspect of the expansion of these stores: stores are located closer to households, and therefore transportation costs decrease. The simulations indicate that a reduction in transportation costs reduces demand for the smallest sizes, as consumers become better able to take advantage of the quantity discounts offered by larger packages.

The second exercise evaluates counterfactuals where small-format stores are constrained to carry only one size of detergent due to limited shelf space. I simulate consumers' dynamic choices under each of the six package sizes observed in the data and calculate the expected discounted present value of consumer utility in each case. The results show that offering one of the larger package sizes is optimal: the utility-maximizing package size is 3 kg in low-income areas and 2 kg in high-income areas. Offering only one of the smaller sizes yields lower utility. These results suggest that stocking only one of the smallest sizes, the current practice in many small-format stores, is not utility maximizing.

This paper continues a line of research on dynamic demand estimation that includes Lal and Rao (1997), Pesendorfer (2002), Erdem, Imai and Keane (2003), Arcidiacono and Miller (2011), Gowrisankaran and Rysman (2012), and Hendel and Nevo (2006, 2013). Unlike earlier studies, I consider a developing country application and explicitly account for the bundling opportunities present in some markets.

This paper also contributes to a large literature on the implications of quantity discounts in the context of poor consumers (Rao 2000; Beatty 2010; Attanasio and Pastorino 2020; Dillon, De Weerdt and O'Donoghue 2021; Abubakari et al. 2025). Attanasio and Pastorino (2020) estimate a model of quantity discounts in the context of the Progresa cash-transfer program in Mexico, and show that the program led to an increase in price discrimination (see also Attanasio et al. (2013)). Dillon, De Weerdt and O'Donoghue (2021) find that Tanzanian households often purchase goods in small quantities and forego the available bulk discounts. This appears to be driven by consumers' worry that a large inventory would increase the temptation to over-consume, and give rise to "social taxation" by family and friends. The current paper complements this literature by explicitly modeling consumers' dynamic choices regarding their inventory, by using scanner data, and by focusing on a non-food item (laundry detergent). The scanner data is useful because it contains direct information on prices and hence quantity discounts. In all other developing country studies I am aware of, quantity discounts are estimated (e.g., from household consumption diaries). Focusing on non-food items is useful because food items are especially prone to the temptation to overconsume (i.e., cravings) (Chandon and Wansink 2002; Dillon, De Weerdt and O'Donoghue 2021), feature heterogeneity in quality, and are associated with inherent preferences for variety, all of which could contribute to the frequency of small purchases. These motives are less important for

the laundry detergents considered here. My findings offer an alternative explanation for small purchases: rational bundling by cost-minimizing consumers.

Finally, this paper contributes to a growing literature on the retail sector in developing countries. In African countries, the literature has traditionally focused on small, often informal retailers and traders (e.g., Fafchamps and Minten (2001), Fafchamps and Hill (2008), Kremer et al. (2013)). Elsewhere, Atkin, Faber and Gonzalez-Navarro (2018) study the impact of Walmart on consumer welfare in Mexico, and Lagakos (2016) provides a theory where consumer transportation constrains the adoption of new retail technologies. I contribute to this literature by providing, for a storable product, counterfactual analyses designed to capture specific aspects of the expansion of formal stores on consumers, such as reduced transportation costs and limited product offerings.

2 Data and background

2.1 Hand-wash detergents in South Africa

At the time of the study, the South African laundry detergent market had two distinct types of product: hand-wash and automatic detergents. These products do not substitute eachother. Powdered hand-wash detergents, the product considered in this paper, are used for washing by hand as well as in low-efficiency washing machines typical in South Africa.[4] Automatic detergents are used in high-efficiency machines and are considered a niche product. Hand-wash detergents, account for about 90 percent of all sales during my period of study (Figure A.1).[5]

Among hand-wash detergents, Unilever had a stable 80-95 percent market share throughout 2011-2012 (Figure A.1 in the Appendix). The remaining products on this market are either small brands with mostly specialty detergents (such as detergents sold for baby clothes or dark clothes), or generic store brands. Small brands have at most a 2 percent share of the market, while generic brands have a combined share of at most 4 percent. There were no major product launches, mergers, or new competitors during this period.[6]

[4]Around 20% of South African households owned washing machines in 2011, most of which are low-efficiency machines. These drain into a bathtub or are used outside the house.

[5]The market share of bar soap and other low-cost cleaning agents has been small and declining.

[6]The second largest producer, Procter & Gamble, entered the market in 2013, after the end of my data, with up to 10 percent share of the market.

Table 1: Market shares and prices

	Sunlight		Sunlight Tropical		OMO	
	All	Low income	All	Low income	All	Low income
Market shares						
250g	0.04	0.03	0.03	0.01	0.02	0.01
500g	0.08	0.06	0.03	0.02	0.03	0.02
1kg	0.15	0.15	0.06	0.05	0.07	0.06
2kg	0.20	0.23	0.08	0.09	0.15	0.19
3kg	0.05	0.05				
5kg	0.02	0.02				
Prices						
250g	7.92	7.95	7.94	7.95	8.46	8.47
500g	15.30	15.37	15.35	15.39	16.66	16.74
1kg	24.66	24.62	24.82	24.76	26.25	26.20
2kg	39.41	39.50	39.50	39.54	40.69	41.06
3kg	55.00	55.30				
5kg	90.66	91.94				

Notes: Share of units sold of the given product of total units sold in all markets, or in low-income areas. Average prices in Rand across all markets, or in low-income areas. 5,255 markets in total, 912 markets in low-income areas.

2.2 Scanner data

I obtained a unique dataset of supermarket scanner data from all Shoprite stores in South Africa. Shoprite is a South-African supermarket chain that has grown considerably in the past decades and is known for expanding in all areas of the county, including low-income neighborhoods. Shoprite's presence in low-income neighborhoods distinguishes it from other formal grocery store chains in the country.

The data contains monthly information at the store level. This includes quantity sold and price per package separately for all types and brands of Unilever detergents. The data includes information on 330 stores over 16 month period, from July 2011 to October 2012. From now on, I refer to a store-month as a "market," and there are a total of 5,255 markets in my data (a few stores were not open during the entire sample period).

2.2.1 Products and package sizes

I focus on the three main hand-wash detergent brands produced by Unilever: Omo, Sunlight, and Sunlight Tropical. This group of products comprises 70-85% of all detergent sales across Shoprite stores, and a similar market share nationally.[7] Within this group, Sunlight accounts

[7]The highest market share of other Unilever hand-wash detergents is 3 percent. For non-Unilever hand-wash detergents, it is 1 percent.

for 54% of units sold. Omo, which is considered to be a higher-performance detergent with better cleaning ability, accounts for 27%. This brand is more popular in urban areas. Sunlight Tropical, which has a share of 20%, is a variation of the Sunlight brand - it is perceived to have a stronger, longer lasting scent. Perhaps because of this feature this is more popular in low-income and rural areas.[8]

These brands are sold in various package sizes, ranging from 250 g to 5 kg, for a total of 14 brand-size combinations over the period of study.[9] Table 1 lists these 14 products, along with their market shares and prices. The most popular size is the 2 kg package, which accounts for 43% of units sold. The second most popular size is 1 kg, with 28% of units sold. The smallest package sizes are not very popular: the 250 g and 500 g packages respectively account for 9 and 13% of all units sold.[10] Capturing consumers' actual purchases, the scanner data reveals that South African households do not in fact buy small packages very often.

Frictions related to liquidity constraints would suggest that low-income households should be especially likely to purchase the smallest packages. The data provide little support for this in the current setting (Table 1). Although incomes exhibit large differences across areas while prices of a given package size do not, low-income households do not disproportionately purchase the smallest package sizes. If anything, these households purchase relatively larger packages: the share of units sold from the smallest packages is always lower in low-income areas.[11]

2.2.2 Prices, quantity discounts, and promotions

There are some noteworthy features of prices in this context. Specifically, (1) they are not uniform across stores even though all stores belong to the same chain, (2) there is a substantial quantity discount embedded in the prices, and (3) promotions affect different package sizes at different times.

Non-uniform prices across stores and time. The average price of a given package size

[8] During my period of study, there is no clear trend in the market share of any of the sample products over time (Figure A.2). There are very few store-months when a particular product had zero sales (Table A.3). In these cases I cannot distinguish whether the product was physically in the store but had no sales, or was not available.

[9] There are no other differences between different sizes of a given brand, e.g., in packaging materials or user instructions.

[10] The share of these products in terms of revenue or volume sold is even smaller (2 and 7% or 1 and 4%, respectively), while the share of the two largest packages is larger (8 and 6% or 9 and 8%, respectively).

[11] Dillon, De Weerdt and O'Donoghue (2021) also find that liquidity constraints are not a major determinant of package size choice for food items in Tanzania, and Fehr, Fink and Jack (2022) find that poorer households are more likely to take advantage of profitable opportunities. If present, liquidity constraints are likely to matter over shorter periods than the monthly data observed here (Orhun and Palazzolo 2019). To the extent that liquidity constraints raise the cost of purchasing larger packages, they will be captured in the model through size-specific fixed costs of purchase.

is similar across brands (Table 1).[12] However, these similar average prices mask significant price variation both across stores and across time. This is in contrast to patterns documented in US data, for example in Hendel and Nevo (2013) and DellaVigna and Gentzkow (2019). On the top panel of Figure 1, I display the distribution of prices for selected products for a specific month to illustrate price dispersion of the same product across the 330 stores. The lower panel displays the same two product prices across all markets, which shows even more dispersion. Generally, smaller sizes have less price dispersion than larger sizes. Figure A.4 shows additional price variation: the evolution of prices over time for the most popular products in selected stores. These price differences across stores and over time provides the variation necessary to identify demand parameters.

Figure 1: Distribution of prices of selected products



Notes: Store level monthly prices in Rand. The upper panels show the price distribution across stores for two products in April 2012 (the middle of the sample period). The lower panels show the corresponding distributions over the entire sample (16 months, 330 stores).

Non-linear pricing. Table 2 shows the unit price of each package size relative to the 1 kg package of the same brand. The values shown are averages across markets. For example, based on the first column, purchasing 1 kg of Sunlight detergent as four 250 g packages is 25% more expensive than purchasing a single 1 kg package. Table 2 indicates that quantity

[12] Average prices are between 8 and 91 Rand. Average monthly household income in this period was 10,410 Rand.

Table 2: Quantity discounts

	Sunlight	Sunlight Tropical	OMO
250g	1.30	1.29	1.30
500g	1.25	1.24	1.27
1kg	1	1	1
2kg	0.80	0.80	0.78
3kg	0.74		
5kg	0.74		

Notes: Unit prices (per 1 kg) relative to a 1 kg package of the same brand. Based on regular (non-promotional) prices.

discounts across brands are very similar. In addition, although there are up to 5 different sizes from the same brand, there seem to be only three groups of unit prices. Smaller products are on average 25-30 percent more expensive, and larger products are 20-25 percent cheaper compared to a 1 kg package.[13]

There is variation in quantity discounts over time and across stores. To illustrate this, Figure 2 plots the quantity discount of a 2 kg package versus buying two 1 kg packages across the 330 stores for April 2012 (the middle of the period). The x-axis shows the relative prices (price of a 2 kg package divided by the price of two 1 kg packages). There are substantial (up to 35 percent) differences in quantity discounts. In most cases, these measures are below 1, as one would expect. However, there are also values above 1, where purchasing two 1 kg packages is less expensive than one 2 kg package. This feature arises because of temporary promotions.

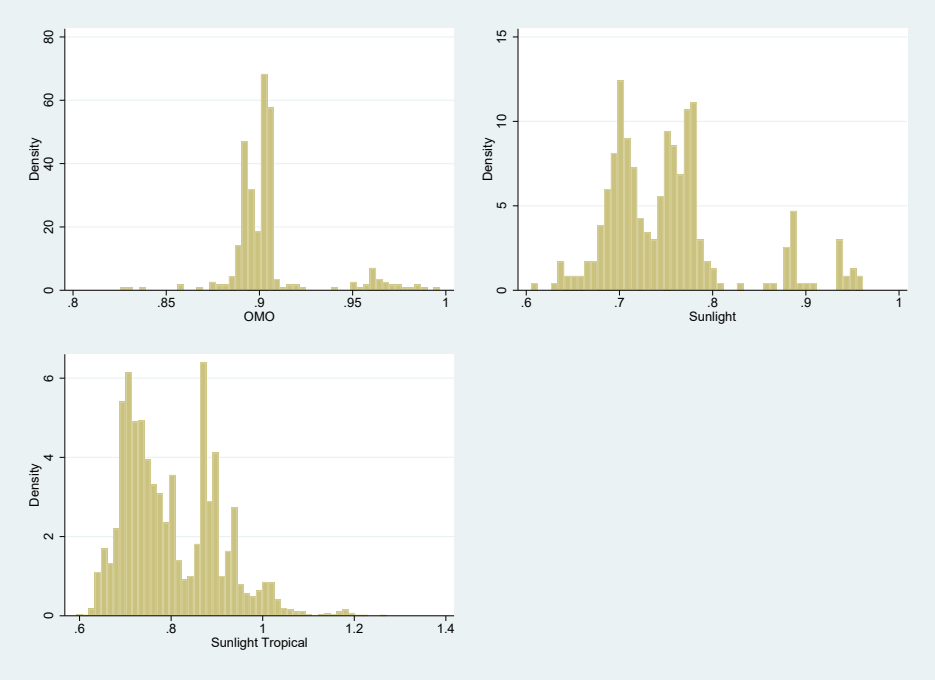
Temporary promotions. Pricing of laundry detergents is subject to temporary promotions.[14] On Figure 3, I plot the percentage of volume on promotion over time for three packages sizes of the most popular brand. Promotional periods typically involve only a few sizes. For example, on the week indicated with the vertical line on Figure 3, the 2 kg package was on promotion and the 5 kg package was not. Because the regular unit price of these packages tends to be similar (Table 2), a promotion on the 2 or 3 kg package creates periods where it is cheaper to buy 2+3 kg than 5 kg. This feature of the data will be discussed extensively in Section 3 below.

The figure also shows that in almost all cases, either zero or 100 percent of the volume

[13]This pricing behavior by the firm could be caused by price discrimination or cost differences across sizes - see, e.g., Cohen (2008). I also computed the same figures separately for different income areas, and I do not find any noticeable differences compared to Table 2.

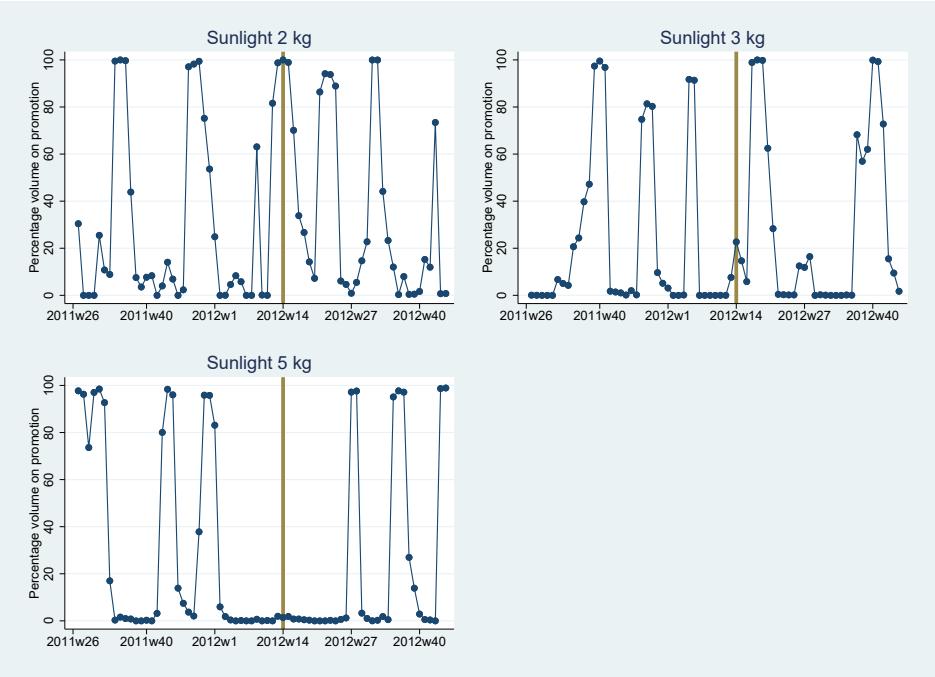
[14]I obtained data on weekly promotions - specifically, information on volume sold each week on temporary promotions at the national level.

Figure 2: Store level variation in quantity discounts



Notes: Prices across stores of one 2 kg vs two 1 kg packages of the same brand for April 2012 (middle of sample period).

Figure 3: Percentage of volume on promotion for selected products



Notes: Total across all stores. The vertical line indicates the same week on each panel.

was sold at promotional prices, which strongly suggests that promotions were conducted nationally by the chain, rather than individually at the store level.[15]

2.2.3 Store characteristics

The data allowed me to collect the GPS coordinates of each store using Google Maps. Knowing the exact location of the stores is an advantage of my dataset relative to, e.g., the popular US scanner database from Nielsen. There are Shoprite stores in all nine South African provinces (Figure A.6), and there are stores are located in urban, suburban and rural areas as well (Figure A.7).

Since I know the identity of each store, I was able to collect store characteristics from individual stores' websites. I collected the following information: store is located in a shopping mall, store is located in a city center, and Sunday opening hours (Table A.8). Approximately 9 percent of the stores are in a shopping mall. There is large variation in Sunday opening hours across stores, ranging from closed all day to open until 9 pm. I use these variables to proxy for unobservable market characteristics, such as popularity or accessibility of the store.

Unilever uses a Living Standard Measure (LSM) to categorize stores based on their location for internal marketing purposes. This measure is designed to approximate consumers' purchasing power, and places each store into one of three main categories.[16] For simplicity, in the rest of the paper I refer to low-, middle- and high-income areas based on this LSM definition. In my data, 17 percent of the stores fall into the low-income category, 44 percent into the middle-income category, and 39 in the high-income category (Table A.8).[17] Average household income in the lowest category is 2285 Rand (310 USD) per month, which makes this group similar to low-income households elsewhere in Africa.

2.2.4 Market size definition

Defining the market size has important practical implications. If the market size is defined too narrowly (smaller than the actual quantity of laundry detergent sold) estimation becomes infeasible. Conversely, if the market size is too broad, it complicates the interpretation of

[15] Price variation across stores is present both for products on promotion and for products not currently on promotion - see Figure A.5 in the Appendix.

[16] The Living Standard Measure is created by the South African Audience Research Foundation and is used widely by companies for marketing purposes. Based on multiple measures, it divides the population into 10 LSM groups, from 10 (highest) to 1 (lowest). In the data, stores are categorized into low LSM (1-4), medium LSM (5-6) and high LSM (7-10) areas. More details about LSM are available at http://www.saarf.co.za/lsm/lsms.asp

[17] The share of the adult population in South Africa in each of these categories is 32, 18 and 50 percent, respectively.

the outside option. The market size definition also affects which demographic characteristics should be included in the estimation. In South Africa, many areas are segregated, but substantially different populations live in relatively close proximity to each other. Thus, a more precise market size definition helps match the actual households that use particular stores. Similarly, having the right market size definition can be important to know which characteristics may affect consumers' product size choices. For example, it is common for the households considered here to walk to the store, which makes the cost of carrying a large package an important consideration in the dynamic model.

I define market size based on distance to the store.[18] Specifically, based on the Census, for each store, I compute the number of households N^r who live within r km, where $r = 1, 2, ..., 50$. I also compute the maximum number of laundry products the store sells in a given month over the time period observed in the data, L. I set the market size to be N^{r^*} where $r^* = \min\{r|L \cdot 1.1 \leq N^r\}$, i.e., the smallest radius where the number of households exceeds the total number laundry products sold, plus 10 percent. This is based on my survey, where 93.7 percent of households said that they purchased detergent once a month, and only 2 percent said that they made more frequent purchases.

Using this definition, the average market size is 11,432 households. The market size has a radius of 5 km or less from the store for 94.8% of the stores. The radius is 1 km for 43.9% of the stores, and it is 2 km for 32.7%. The distribution of the market radius is shown on Figure A.9. I also use this market size definition to create control variables in the estimation, specifically, the number of stores, and the distance from the store to its nearest neighboring store.[19]

2.3 Demographic data

I combine two additional datasets with the scanner data. First, I use the 2011 South African Census Community Profile dataset. The Census was undertaken exactly in the time period covered by the scanner data, and it allows matching households to very precise geographical areas. Specifically, the data contains an identifier called small area code which contains only 100-500 households (a few street blocks).

The Census contains household level information on annual income, type of main dwelling, population group (ethnicity) and gender of the household head, and ownership of various appliances. In the demand estimation I found it useful to include ownership of a wash-

[18] Commonly used market size definitions, such as equating markets with metropolitan areas, are not useful here. For example, the metropolitan area of Tshwane (the area around the capital city of Pretoria) has 26 stores, and almost all stores have a neighboring store within 2 kilometers.

[19] The median distance to the nearest store is 4.7 km, the mean is 17.8 km.

ing machine and ownership of a car. Ownership of a car can be important in explaining whether households choose larger package sizes that are more difficult to transport. Since I observe individual households, I can directly sample individual households on each market for estimation.

Table A.10 provides summary statistics and compares households on the market of Shoprite stores to the entire South African population. Although Shoprite stores are somewhat overrepresented in and around larger urban areas, 16% of the households on these stores' markets live in rural areas and 14% of them live in informal dwellings.

2.4 Survey data

The second source of information on households is from two rounds of a survey conducted specifically for this paper to collect data on how households purchase laundry detergents, and the amount of laundry detergent they keep at home (i.e., their inventory). The survey also collected descriptive information, such as the kinds of transportation households use when buying detergent, and the different costs of holding inventory (e.g., whether neighbors might ask to borrow detergent).

Details of the survey are described in Appendix A.2.2. I randomly selected 3 stores, one in each income area, and randomly sampled 100 households around each of these stores. The first round of the survey took place in December 2020 and the second in March 2022, about 16 months apart, which also corresponds to the length of the scanner data. Even though my survey was conducted in a different time period than the scanner data, it helps alleviate the identification challenges in estimating households' value function by providing direct information linking purchase, consumption, and inventory (see Appendix A.2.2).

The surveyors asked households to show them what kinds of laundry detergent they had at home. Surveyors recorded a description of the item(s) shown to them, including the package size and the amount left in each package - for example, "Sunlight 500 grams in a yellow plastic, dry, half empty" or "OMO washing powder regular 2kg just opened and dry." (No respondents refused to answer this question.) Based on this information, I compute the current inventory of detergents for each household.[20] There are 289 observations with nonmissing inventory and consumption in the first round of the survey, and 286 in the second round. Respondents without an inventory level showed either liquid detergents or laundry bars (soaps); all households had at least some detergent at home. Figure A.13 shows the distribution of inventories by income area.

[20] In some cases, the surveyor recorded that the package was "almost full" or "almost empty." In the first case, I multiply the package size by 0.9, and in the second case, by 0.1.

For consumption, I ask households how often they typically do laundry (daily, weekly, or monthly), and how many loads of laundry they typically do each time. Based on these, I compute how many loads of laundry a household typically does in a month, and the associated consumption of detergent (one load of laundry requires about 100 gram of powdered detergent, based on the directions on the package).

Table 3 shows the summary statistics of inventory and consumption. The data reveal little difference in consumption and inventory between income areas. If anything, consumption is higher in low-income areas, which is consistent with the fact that these households also buy larger packages (Section 2.2.1). Average inventory is similar to average consumption.

Table 3: Detergent inventory and consumption by income area in the survey

	Mean	Median	St.dev.	Min	Max	N
Inventory						
Low-income area	1035.92	1000	767.12	50	3600	190
Middle-income area	1181.80	1000	788.60	100	3800	172
High-income area	1185.83	1000	722.80	100	3800	180
All	1132.00	1000	761.58	50	3800	542
Consumption						
Low-income area	1223.96	1200	554.55	300	3200	96
Middle-income area	1148.19	1200	559.91	400	3200	83
High-income area	1113.48	1200	441.38	400	2400	89
All	1163.81	1200	521.66	300	3200	268

Notes: Values are in gram. Inventory information was collected in both survey rounds. Consumption information was only collected in the first round. Inventory shown is computed as described in Appendix A.2.2.2.

3 Savings due to bundling

As documented in the previous sections, the product features non-linear pricing, with similar unit prices within each of three broad size groups.[21] At the same time, temporary promotions, which usually target only some sizes, can substantially decrease prices.[22] In many markets, these two facts combined result in situations where it is cheaper to buy two smaller packages instead of a larger one to obtain the same total quantity. That is, the relatively flat quantity

[21] Prices in supermarkets are clearly displayed. I also checked this in the survey: to a question asking whether the prices of laundry detergent were clearly displayed, 92% of respondents indicated "always," 8% indicates "sometimes" and none indicated "almost never/never."

[22] From the firm's point of view, the use of promotions on some package sizes but not others of a given product is consistent with Aguirregabiria (1999).

discount within size groups combined with temporary promotions that affect only one or two sizes per brand induce non-monotonicity in the nonlinear pricing schedule.

In the data 14 products are being sold, where a "product" is a particular package size of a particular brand. By bundling smaller packages of the same brand, the quantities corresponding to these package sizes can be purchased in a total of 30 different combinations. For example, 1 kg of Sunlight Tropical could be purchased as four 250 g packages or two 500 g packages (as well as a non-bundled 1 kg package). Out of the 30 possible bundles, 16 are cost-minimizing at least once in the data, i.e., it is cheaper to purchase a bundle instead of a single package. Table 4 lists these product bundles.[23]

Table 4: Cost-minimizing product bundles

	Stores	Months	Markets	Price saving	Or. price
Sunlignt Tropical					
Buy 2×250g instead of 500g	265	16	988	0.20	16.32
Buy 2×500g instead of 1kg	54	1	54	0.96	23.70
Buy 2×1kg instead of 2kg	137	7	214	2.10	45.34
Buy 4×500g instead of 2kg	2	1	2	1.62	45.18
All	287	16	1207		
Sunlight					
Buy 2×250g instead of 500g	290	16	974	0.22	16.00
Buy 2×500g instead of 1kg	13	1	13	0.88	26.55
Buy 2×1kg instead of 2kg	140	6	220	1.46	45.30
Buy 1×2kg and 1×1kg instead of 3kg	308	7	967	3.46	66.37
Buy 3×1kg instead of 3kg	4	3	5	1.39	59.97
Buy 1×2kg and 2×500g instead of 3kg	11	1	11	1.48	66.08
Buy 1×2kg and 1×3kg instead of 5kg	329	16	1764	11.53	100.00
Buy 2×1kg and 1×3kg instead of 5kg	10	3	10	3.62	106.50
Buy 2×2kg and 1×1kg instead of 5kg	206	8	263	7.73	101.52
Buy 5×1kg instead of 5kg	2	1	2	0.33	95.82
All	330	16	3355		
OMO					
Buy 2×250g instead of 500g	283	16	1849	0.42	17.30
Buy 2×1kg instead of 2kg	3	1	3	0.40	43.84
All	283	16	1852		

Source: The table lists product bundles that are cost-minimizing at least once in the data. Entries indicate the number of times a specific bundle is the cost-minimizing option. Total number of stores is 330, total number of months is 16. Total number of markets is 5255. Price saving shows the average price saving in Rand from the bundle across the markets where that bundle is cost-minimizing. Or. price is the average price of buying the same total quantity as a single package on these markets.

[23]Note that a mixture of *different* small packages can never be cost-minimizing if the same total quantity can also be obtained as a bundle of identical packages (e.g., buying two 250g and one 500g package cannot be cost-minimizing).

The largest average price saving comes from purchasing a 2 kg and a 3 kg package instead of a 5 kg package. This is worth it in 1764 out of the 5255 markets in the data. There are markets where bundling can be cost-minimizing for multiple sizes and/or brands at the same time. For example, I see 363 markets where buying two 250 g packages of OMO instead of one 500 g package *and* buying a 3 kg and a 2 kg package of Sunlight instead of one 5 kg are both cost-minimizing bundles. In total, there are 57 possible combinations of these bundling opportunities in the data.

Bundling opportunities do not arise only in specific months or in specific stores. Based on Table 4, each brand has bundling opportunities in every month, Sunlight has bundling opportunities in every store, and Sunlight Tropical and OMO have bundling opportunities in 287 and 283 stores, respectively. I investigate any potential correlation between bundling opportunities and market characteristics, including whether the store is in a city center or shopping mall, area income and other average household characteristics. Table A.15 in the Appendix shows that none of these individual correlations are statistically significant. Table A.16 shows that the 14 characteristics I consider jointly explain no more than 1 percent of the variation on top of basic controls such as province and month fixed effects. Variables such as area income or store location appear to have little to no correlation with bundling opportunities.[24]

The presence of bundling opportunities has important implications in the context of African countries, where the common belief is that due to financial or other constraints, households often purchase products with the lowest package price rather than the lowest unit price. Looking purely at market shares can be misleading: a consumer might optimally purchase multiple small packages because this bundle is less expensive than a larger package.

In the next section, I address this issue in the context of demand estimation.

4 Model and estimation

Since laundry detergent is a storable product and prices change over time because of temporary discounts, the consumer may purchase the product in a dynamic manner. Besides these obvious features, the demand model should also incorporate the fact that the same laundry detergent is sold in different package sizes, and that each size has different brands (i.e., a differentiated product). The estimation needs to take into account of the fact that that there are bundling opportunities present on some, but not all markets, and that purchase

[24] Anecdotally, prices have three components: (i) a manufacturer-suggested price, (ii) adjustments by Shoprite, and (iii) adjustments by the local store manager. Bundling opportunities appear to arise from the idiosyncratic combinations of these components.

is observed at the market level rather than the individual level. The model and estimation below incorporate all these features.

4.1 Model setup

Consider the following model based on Erdem, Imai and Keane (2003) and Hendel and Nevo (2006). On a given market, consumer h makes monthly choices between buying brands $j = 1, ..., J$ sold in different package sizes x. The good is storable, and any quantity not consumed is stored as inventory. The consumer trades off the cost of holding an inventory with its benefits. The benefits arise on the one hand from the quantity discount embedded in purchasing a larger size, and on the other from the ability to exploit temporary price discounts.[25]

The per period utility is

$$\begin{aligned} U(c_{ht}, x_{jht}, i_{h,t+1}) &= u(c_{ht}, \nu_{ht}, \gamma) - C(i_{h,t+1}, \boldsymbol{\theta}^C) - F(x_{jht}) \\ &\quad + \boldsymbol{\beta}_x \mathbf{a}_{jxt} + \xi_{jxt} + \mu(\mathbf{a}_{jxt}, \mathbf{D}_h, \boldsymbol{\theta}^\mu_x) + \varepsilon_{hjxt} \end{aligned}$$

where u is the utility from consumption, c_{ht} is quantity consumed (from all brands) by the consumer at time t, ν_{ht} is randomness in the consumer's needs, $C(.)$ is the cost of holding inventory i, $F(.)$ is a size-specific cost of purchase, $\mathbf{a}_{jxt}$ are observed product characteristics (price and brand), ξ_{jxt} is the valuation of unobserved characteristics, $\mu(.)$ is the valuation of product characteristics as a function of consumer demographics $\mathbf{D}_h$, and ε_{hjxt} is a choice-specific shock. The parameters to be estimated are the marginal utility of consumption γ, inventory cost parameters $\boldsymbol{\theta}^C$, and the valuation $(\boldsymbol{\beta}_x, \boldsymbol{\theta}^\mu_x)$ of product characteristics.[26]

Following Erdem, Imai and Keane (2003), I assume that monthly consumption is exogenous and stochastic. This reflects the idea that each month, a household faces a specific need for laundry detergent, and does not get extra utility from using more than the needed amount (for example, by using more detergent for a load of laundry than what is recommended by the manufacturer, or by doing more loads of laundry than needed). This implies that detergent consumption is insensitive to the household's inventory, which is consistent

[25] Based on the survey, there is no evidence that households in this setting purchase large packages for something other than their own consumption. I asked whether anyone participated in a club where people bought detergent together in bulk, or whether they bought together with friends or neighbors. Everyone indicated "no" to the first possibility and only 1 respondent indicated "yes" to the second.

[26] As usual, the per-period utility function subsumes a budget constraint that the household allocates between consumption of the product in question (laundry detergent) and a composite good. Assumptions about whether the inventory cost C and the purchase cost F represent utility costs or monetary costs would lead to slightly different parametrizations that would not affect the estimation results. I follow Hendel and Nevo (2006) and work with the indirect utility function above.

with data from my survey (Appendix A.2.2.3).[27] The estimation below does allow for stock-out, i.e., the possibility that the inventory is insufficient to cover the consumer's needs, in which case consumption is equal to the inventory.

The inventory transition is given by

$$i_{h,t+1} = i_{ht} - c_{ht} + x_{jht} \tag{1}$$

The consumer's problem is:

$$V(\sigma_1) = \max_{\{j, x_{jht}|\sigma_t\}} \sum_{t=1}^{\infty} \delta^{t-1} E\left[U(c_{ht}, x_{jht}, i_{h,t+1})|\sigma_1\right] \tag{2}$$

s.t. $0 \leq i_{ht}$, $0 \leq c_{ht}$, $0 \leq x_{jht}$ and (1), where the state variables σ_t are $(i_{ht}, \nu_{ht}, c_{ht})$ as well as the vectors of product characteristics $\mathbf{a}_{jxt}$ and ξ_{jxt}, and the vector of choice-specific shocks.

I make similar assumptions regarding the distribution and evolution of consumption shocks, choice-specific shocks and prices as Hendel and Nevo (2006). Namely, the consumption shock ν_{ht} is independently distributed over time and across consumers, prices follow a first-order Markov process, and the choice-specific shocks are i.i.d. Type-I extreme value.

4.2 Towards estimation

Given the above assumptions, the choice of which brand of detergent to buy can be written as a static problem for the consumer. A useful summary measure of this choice is the "inclusive value,"which represents the expected maximum utility across brands conditional on choosing size x. Because the utility of brands conditional on size is distributed Type-I EV, the inclusive value for size x is given by

$$\omega_{hxt} = \log\left[\sum_j \exp(\boldsymbol{\beta}_x \mathbf{a}_{jxt} + \xi_{jxt} + \mu(\mathbf{a}_{jxt}, \mathbf{D}_h, \boldsymbol{\theta}_x^{\mu}))\right], \tag{3}$$

where the sum is over brands which come in size x.

Assuming that the transition process of inclusive values satisfies $\Pr(\omega_{hx,t+1}|\boldsymbol{a}_t, \boldsymbol{\xi}_t, \boldsymbol{\mu}_{ht}) = \Pr(\omega_{hx,t+1}|\boldsymbol{\omega}_{ht})$, the dynamic programming problem simplifies to

$$\begin{aligned} V(i_{ht}, \nu_{ht}, c_{ht}, \boldsymbol{\omega}_{ht}, \boldsymbol{\varepsilon}_{ht}) = \quad & \max_{x_{ht}|\sigma_t} u(c_{ht}, \nu_{ht}, \gamma) - C(i_{h,t+1}, \boldsymbol{\theta}^C) - F(x_{ht}) + \omega_{hxt} + \varepsilon_{hxt} \\ & + \delta E\left[V(i_{h,t+1}, \nu_{h,t+1}, \boldsymbol{\omega}_{h,t+1}, \boldsymbol{\varepsilon}_{h,t+1})|i_{ht}, \nu_{ht}, \boldsymbol{\omega}_{ht}, \boldsymbol{\varepsilon}_{ht}, x_{ht}\right]. \end{aligned} \tag{4}$$

[27] Using data from France, Chandon and Wansink (2002) also find that detergent consumption is unrelated to a household's inventory, in contrast to food items like orange juice and cookies.

The dynamic decision on package size is conditioned on the inclusive values that summarize the static brand choice. Thus, consumers choose package size as a function of their state variables: inventory, inclusive values, consumption shocks, and choice-specific utility shocks.

For estimation, I need to specify functional forms for the objects defined above. Following Hendel and Nevo (2006), I specify

$$u(c_{ht}, \nu_{ht}, \gamma) = \gamma \log(c_{ht} + \nu_{ht})$$

$$C(i_{ht}, \boldsymbol{\theta}^C) = \theta_1^C i_{ht} + \theta_2^C i_{ht}^2 + \theta_3^C i_{ht}^3$$

The distribution of ω_{hxt} is assumed to be Normal, with standard deviation σ_x and mean

$$\iota_0^{(x)} + \sum_{x'} \iota_{x'}^{(x)} \omega_{hx',t-1},$$

where the ι's are parameters to be estimated.

Figure A.2 in the Appendix presents an overview of the estimation steps. To proceed with the dynamic estimation, one needs to compute inclusive values for all markets. The inclusive values in (3) contain marginal utilities of the prices and product characteristics from the static demand. Since individual purchase data is not available, a BLP (Berry, Levinsohn and Pakes 1995) type discrete choice estimation is the natural approach to recover these marginal utilities. However, the discrete choice assumption is violated in markets where there is a potential bundling opportunity. I now describe a possible solution to this issue.

4.3 Estimating static parameters and correcting the market shares

First, consider markets with no bundling opportunities (i.e., where based on the observed prices, it is not cost-minimizing to purchase multiple smaller packages instead of a bigger package). I observe a large number of such markets: for each size, I have over 1000 markets where there are no bundling opportunities (Table A.14).

For these markets, one can proceed with a conventional discrete choice demand model. In particular, I estimate $\boldsymbol{\beta}_x$, $\boldsymbol{\xi}_{xt}$ and $\boldsymbol{\mu}_{ht}$ (that is, all the ingredients for the inclusive values in (3)) using standard BLP. A detailed discussion of the choice of linear variables, the specification of individual heterogeneity through nonlinear terms, and the instruments used to address price endogeneity and identify the nonlinear parameters is provided in Appendix A.2.1. Since the consumer's static brand choice is conditional on size, a separate BLP estimation is run for each set of products of a given size x. In each case, the interpretation

of the outside option is "a choice other than size x."[28]

Next, I consider markets with bundling opportunities. On these markets, I first need to estimate the share of packages that were purchased as part of a bundle. I then use these estimates to "correct" the market shares, and compute all necessary ingredients and proceed with the estimation of dynamic demand.

I begin by computing the total price of buying each cost-minimizing bundle (for example, the price of buying a 2 kg and a 3 kg package instead of a 5 kg package of the same brand when the latter is more expensive). Call this the "effective price" of the given total quantity of a given brand (in this example, the effective price of 5 kg of this brand).

Using these effective prices together with the BLP parameter estimates obtained above, it is possible to compute the "corrected" market share of each size x of the product - the share of purchases that contain a total of x kg of the product. To do this, in the above example I use the parameter estimates to compute the demand increase Δ_{j5t} for a 5 kg package when its price drops to the effective price. This is the extra demand for 5 kg of the product that will be fulfilled through the purchase of bundles.[29] The corrected market share for $x = 5$ is then $\hat{s}_{j5t} = S_{jt}^{(5)} + \Delta_{j5t}$, where $S_{jt}^{(5)}$ is the observed market share of the 5 kg packages sold. To compute the corrected market share of $x = 2$ and $x = 3$, I consider how many packages of these sizes are needed to assemble the extra 5 kg bundles. In this case, the number of packages needed is simply 1 of each size per bundle. The corrected market share for 2 kg and 3 kg will thus be $\hat{s}_{j2t} = S_{jt}^{(2)} - (\hat{s}_{j5t} - S_{jt}^{(5)})$ and $\hat{s}_{j3t} = S_{jt}^{(3)} - (\hat{s}_{j5t} - S_{jt}^{(5)})$, respectively.

What remains is the computation of the ξ's on these markets with bundling opportunities. To obtain these, I use the BLP parameter estimates $\boldsymbol{\beta}_x$ and $\boldsymbol{\mu}_{ht}$ and solve systems of equation of the form $s_{jxt}(\boldsymbol{a}_{jxt}, \boldsymbol{\beta}_x, \boldsymbol{\xi}_{xt}, \boldsymbol{\mu}_{ht}) = \hat{s}_{jxt}$, i.e., equating the model-predicted market shares to the corrected market shares. For sizes not involved in bundling (for example, packages other than the 2, 3, and 5 kg in the example above), I simply use the market shares observed in the data, $S^{(x)}$, and solve $s_{jxt}(\boldsymbol{a}_{jxt}, \boldsymbol{\beta}_x, \boldsymbol{\xi}_{xt}, \boldsymbol{\mu}_{ht}) = S_{jt}^{(x)}$.

This procedure provides all the information necessary to assemble the inclusive values ω, and provides corrected market shares for markets with bundling opportunities, making it possible to include these markets in the dynamic problem.

[28] Note that any differences in the utility of the outside option across these specifications will be captured by the size-specific dummy variables $F(x_{ht})$ included in the per-period utility function in the dynamic estimation.

[29] To compute Δ_{j5t}, I draw 1000 values of the ξ_{j5t} estimated for the markets with no bundling. Using these values, I compute the market shares predicted by the model under the original and the effective prices for the market under study, and take the average difference.

4.4 Estimation of the dynamic parameters

As described above, the inclusive values in (3) allow simplifying the dynamic problem, which now requires keeping track only of expected utility conditional on size. Since the (simulated) households have specific characteristics, the inclusive values will be household-specific as well. Below, I start by estimating a single process for all households; I then relax this assumption and allow separate processes for household groups by income area.

The estimation proceeds by computing the consumer's value function (4) for a given trial of the dynamic parameters. This requires identifying the dependence between purchased package size, inventory, and consumption. Typical applications must address the identification challenge that only purchases are observed, while consumption and inventory are unobserved (Hendel and Nevo 2006; Erdem, Imai and Keane 2003). My survey helps address this challenge by collecting information on these elements. This information is used in approximating the value function (see Appendix A.2.2 for details).

With the estimated value function, the dynamic choice probabilities of each simulated consumer can be computed. Specifically, for each package size x the model implies

$$\Pr(x_{ht} = x) = \frac{\exp(\omega_{hxt} + M(c_{ht}, i_{h,t+1}, \nu_{ht}, \boldsymbol{\omega}_{ht}, x))}{\sum_{x'} \exp(\omega_{hx't} + M(c_{ht}, i_{h,t+1}, \nu_{ht}, \boldsymbol{\omega}_{ht}, x'))}$$

where

$$\begin{aligned} M(c_{ht}, i_{h,t+1}, \nu_{ht}, \boldsymbol{\omega}_{ht}, x) &= u(c_{ht}, \nu_{ht}, \gamma) - C(i_{h,t+1}, \boldsymbol{\theta}^C) - F(x) \\ &\quad + \delta E\left[V(i_{h,t+1}, \nu_{h,t+1}, \boldsymbol{\omega}_{h,t+1}, \boldsymbol{\varepsilon}_{h,t+1}) | i_{ht}, \nu_{ht}, \boldsymbol{\omega}_{ht}, x\right] \end{aligned}$$

Households' consumption needs c_{ht} are assumed to be fixed and are computed from the data. Starting inventories are set equal to the inventory levels obtained after simulating the model for a number of periods, following Hendel and Nevo (2006) and Erdem, Imai and Keane (2003). For each simulated consumer I compute choice probabilities for 10 vectors of the consumption shock ν_{ht} and take the average. See Appendix A.2.3 for details.

Since the scanner data is at the store level, simulated consumers are aggregated within markets to obtain predicted market shares. Parameters are computed using a simple simulated minimum distance estimator, minimizing the squared distance of model-predicted and observed market shares. Let N_H denote the number of simulated consumers for a given store, and define

$$Q(\boldsymbol{\theta}) = \sum_{x,t} \left(S_t^{(x)} - \frac{1}{N_H} \sum_{h=1}^{N_H} \Pr(x_{ht} = x) \right)^2$$

The estimator minimizes $Q(\boldsymbol{\theta})$ summed across stores. Because the parameters of the inclusive value process enter the dynamic estimation in a complex way, and because the estimation involves multiple simulated objects, it is practical to use a bootstrap to estimate standard errors. I draw datasets of the same size with replacement, and repeat the estimation 30 times. This includes resampling the dataset on which the inclusive value process is estimated.

4.5 Identification

The identification of the static demand parameters follows the identification of the standard BLP procedure, with instruments discussed in Appendix A.2.1. This section describes the identification of the dynamic problem.

Estimating the dynamic parameters, i.e. the parameters of the utility from consumption and the parameters of the inventory cost function, requires solving the dynamic programming problem for each parameter trial. This requires estimating the value function, which is done using value function approximation. I approximate the value function using 32 polynomial terms of the consumer's state variables. Besides the constant, I include a cubic function of log inventory, and quadratic functions of inclusive values, choice-specific extreme value shocks, and consumption shocks.

More formally, approximate the value function in (4) as

$$V_t = \boldsymbol{\varphi}\mathbf{r}_t, \tag{5}$$

where $\mathbf{r}_t$ is a vector of polynomial terms of $(i_t, v_t, \boldsymbol{\omega}_t, \boldsymbol{\varepsilon}_t)$, and $\boldsymbol{\varphi}$ are the parameters to be estimated. The Bellman equation can be written as

$$\boldsymbol{\varphi}\mathbf{r}_t = m_t(x_t) + F(x_t) + \delta\boldsymbol{\varphi}\mathbf{r}_{t+1}, \tag{6}$$

where

$$m_t(x_t) = \gamma \ln(c_t + v_t) - (\theta_1^C(\log i_t) + \theta_2^C(\log i_t)^2 + \theta_3^C(\log i_t)^3) + \omega_{xt} + \varepsilon_{xt}.$$

Identification relies on data from the household survey, which provides information on per-period inventory, consumption, and purchase. From the survey, x_t, c_t, and i_t are all observed (and thus i_{t+1} can be computed).

The inclusive values obtained from the static demand estimation are matched to surveyed households, and thus to their inventory, consumption and purchased size as described in Appendix A.2.3. Note that because inclusive values are estimated from different BLP

specifications that are normalized to different outside options, each ω_x is identified only up to a constant.

For a given realization of shocks and a given draw of dynamic parameters, (6) can be written as

$$m_t(x_t) = \boldsymbol{\varphi}(\mathbf{r}_t - \delta \mathbf{r}_{t+1}) - F(x_t)$$

where $m_t(x_t)$, $\mathbf{r}_t$ and $\mathbf{r}_{t+1}$ are all known. Thus, the parameters $\boldsymbol{\varphi}$ can be estimated using OLS on $(\mathbf{r}_t - \delta \mathbf{r}_{t+1})$ and size-dummies. The coefficient estimates on the size-dummies are equal to $F(x_t)$ plus the normalization constant for the inclusive value of size x. With both inventory and consumption known, identification of the remaining dynamic parameters γ and $(\theta_1^C, \theta_2^C, \theta_3^C)$ follows standard arguments (see Hendel and Nevo (2006, p.1653)).

In practice, the identification of the consumption parameter γ comes from two sources of variation in the data, both of which produce variation between planned and realized consumption levels. First, there is a consumption shock to planned consumption. Second, some purchases do not allow for the full planned consumption, because consumption cannot be higher than the current inventory of the households.

5 Estimation results

5.1 Static demand estimates

I begin by describing the parameter estimates of the static choice between brands, conditional on size. Table A.18 shows the results from a separate BLP estimation for the choice of brands given each package size. As described above, these estimates use data from markets where there are no bundling opportunities. Each of these specifications passes the J-test for the validity of the moment conditions, and the Newey and West (1987) D-test always rejects the null that the nonlinear parameters are jointly 0.

The price coefficient has the expected negative sign and is statistically significant in each specification.[30] Households headed by a male have less elastic demand. This could reflect the fact that these households are wealthier, or it could reflect differences in preferences between men and women when shopping for laundry detergents.[31] Consumers have higher utility from buying larger packages in stores that are open on Sunday: the corresponding coefficient is negative for smaller packages and positive for larger packages.

[30] I do not impose equality of the price coefficient across different product sizes x because consumers' (unmodeled) choice of other products in the consumption basket could depend on x.

[31] In general, the lower price elasticity of male shoppers is consistent with the findings of Fitzpatrick (2017) in Uganda.

Coefficients on the (constant × car only) interaction term are positive and significant for the larger packages (1 kg and above) and switch signs for smaller packages. This indicates that large packages are valued particularly by consumers who own a car but do not own a washing machine.[32] This makes sense, as these are the consumers who use hand-wash detergents more *and* can more easily transport the larger packages. Coefficients on (constant × no car or washm) are positive and significant for the smaller packages and switch signs for larger packages (2 kg and above). This too makes sense, as these are the consumers who use hand-wash detergents but face higher costs of transporting larger packages.

The significance of car ownership in explaining demand for storable products, and laundry detergent in particular, is well-documented in marketing and economics (Blattberg et al. 1978; Talukdar 2008; Griffith et al. 2009). The estimates also mirror patterns in the survey, which specifically asked respondents how they typically get to the store when they buy detergent. In high-income (low-income) areas, 48% (1%) indicated using their own vehicle, and 1% (22%) indicated taking a taxi.[33]

I compute the inclusive values based on the static parameter estimates. Table A.19 shows the estimated processes when a single process is estimated for all simulated individuals and Table A.20 shows the adjusted R^2 from other specifications. First, I estimate separate processes by income area or car/washing machine ownership. Dividing households into these groups does not improve the fit of the regressions. In each case, for every size x, the largest coefficient estimate is on ω^x_{t-1}, i.e., the own lagged inclusive value, similarly to specification shown in Table A.19. To relax the first-order Markov process assumption, I also estimate the inclusive value process in two alternative ways: adding the second lag of each size, and adding the sum of five lags of each size, as in Hendel and Nevo (2006). As shown in the bottom section of Table A.20, neither of these improves the fit meaningfully compared to the baseline specifications.

5.2 The effect of bundling opportunities

Table 5 shows the effect of accounting for possible bundling. The upper panel describes markets with bundling opportunities. Each column corresponds to a total quantity purchased by a consumer at a given time. "Quantity sold" shows the average number of packages sold of the corresponding size in the data. "Corrected quantity sold" shows the estimated number of consumers purchasing that quantity, either by buying one package of the corresponding size, or through bundling. For example, while on average 51.3 units are sold of the 5 kg packages,

[32] The excluded category is composed of consumers who own both a car and a washing machine.

[33] In high-income (low-income) areas, walking to buy detergent was indicated by 60% (81%) of respondents. No-one indicated taking public transportation.

Table 5: Effect of buying in bundles

	5 kg	3 kg	2 kg	1 kg	500 g	250 g
Markets with bundling opportunities						
N of markets	2040	2816	3057	1352	2752	2709
Quantities sold	51.30	290.87	2319.79	1476.44	613.18	385.26
Corrected quantity sold	154.09	241.44	2241.21	1408.01	634.93	349.58
Average ratio	5.46	1.28	0.96	0.94	1.06	0.85
Median ratio	3.02	0.88	0.97	0.97	1.02	0.93
Market share	0.010	0.056	0.401	0.277	0.123	0.084
Corrected market share	0.037	0.045	0.391	0.265	0.130	0.077
Markets without bundling opportunities						
N of markets	1029	1049	1088	1088	1088	1064
Quantities sold	119.73	164.28	1592.165	1508.52	750.03	514.30
Market share	0.025	0.033	0.305	0.336	0.176	0.130

Notes: The table shows summary statistics of each size, separately for markets with a bundling opportunity affecting that size, and for markets without any bundling opportunities. Quantities sold are the number of packages observed in the data, averaged across the relevant markets. Corrected quantities sold applies the bundling correction described in the text. The ratio of quantities sold and corrected quantities sold is computed on each market, and the table shows the average and the median of these ratios. Market share and corrected market share are the corresponding market shares.

I estimate that on average 154.09 consumers buy a total of 5 kg. That is, I estimate that, on average, 102.79 ($= 154.09 - 51.3$) consumers find it profitable to buy 5 kg of detergent by creating a bundle of smaller packages. Similarly, although the average number of 250 g packages sold is 385.26, some of these are purchased as part of a bundle - accordingly, I estimate that the average number of consumers buying a total of 250 g is only 349.58.

I calculate the ratio of Corrected sold and Sold separately for each market; the table shows the average and the median of these measures. On the median market, there are over 3 times as many households who purchase 5 kg of detergent as households who purchase a 5 kg package. This suggests that looking purely at the number of packages sold can be misleading regarding the share of households who purchase larger quantities.

In different markets the corrected market share of a particular package size can be larger or smaller than the market share observed in the data. In some cases, more consumers are buying 2 kg of detergent than 2 kg packages, and in other cases consumers are using 2 kg packages to create bundles, resulting in more 2 kg packages sold. As shown in Figure A.17, there is indeed substantial variation in the ratio of corrected/original quantity sold for each size. Looking purely at the number of packages sold can thus lead to misleading conclusions in either direction regarding the quantity that households prefer to purchase.

For comparison, the lower panel of Table 5 shows the average number of units sold on

markets with no bundling opportunities in the data. For some package sizes, average quantity sold differs substantially from the top panel, which is likely explained at least in part by the presence vs. absence of bundling opportunities. For example, there are many more 5 kg packages sold in markets with no bundling opportunities than on markets with bundling opportunities (119.73 vs 51.30). The Corrected sold values indicate that this difference could be due to the fact that, on the latter, it is profitable to buy 5 kg detergent in other ways than as a 5 kg package.

5.3 Results of the dynamic model

Table 6 reports the dynamic parameters from various specifications. Column (1) is estimated using all stores, and columns (2)-(4) are estimated separately for the three income areas. Implied inventory costs are plotted in Figure 4.

Inventory costs rise from around 7 Rand for 500 g to between 11-13 Rand for 1.5 kg.[34] Beyond that, I consistently find declining inventory costs until about 4 kg, after which costs rise again. (As shown in the Appendix, adding a fourth-degree term to the inventory costs does not affect these patterns.) To understand where the estimated non-monotonicity comes from, note that, in the data, there is no or very little price discount for packages larger than 2 kg. In fact, in some cases, the unit price of a 5 kg package is higher (see Figure A.21 in the Appendix for the distribution of unit prices across markets). At the same time, 3 and 5 kg packages account for about 14 percent of the market. Because larger packages tend to result in higher inventory levels, the non-monotonicity in inventory costs helps rationalize the observed purchases of these larger packages. When interpreting inventory costs, it is important to note that, in a developing country setting, these reflect several considerations beyond the physical cost of storage that are typically the focus of US studies.[35] My survey explicitly asked about the importance of four types of costs: difficulty of storage, inventory getting damaged, inventory getting stolen, and neighbors asking to borrow detergent.[36] For each of these, the share of respondents who indicated that these considerations were "important" / "somewhat important" is 0.32/0.29, 0.38/0.28, 0.22/0.25, 0.32/0.25, respectively. Overall, between half and two-thirds of households indicated that each of these considerations was important or somewhat important. While storage costs are

[34] To interpret the inventory cost parameters, suppose that inventory at the beginning of the period is 250 g and consumption is 750 g. Then according to column (1), buying a 500 g package vs. a 1 kg package results in an inventory cost of 0 vs. 6.83. For comparison, the average savings due to non-linear pricing from buying the 1 kg rather than the 500 g package is 6.19.

[35] The literature using US data typically relies on measures such as house size or urban/suburban location to proxy for these physical storage costs (e.g., Bell and Hilber (2006)).

[36] Note that neighbors borrowing detergent could be either a cost or a benefit (e.g., through gift-exchange).

Table 6: Dynamic parameter estimates

	(1) All	(2) Low-income	(3) Middle-income	(4) High-income	(5) w/o bundling
Inventory cost function					
θ_1^C	175.385	199.109	200.376	175.994	212.224
	(10.003)	(11.323)	(8.486)	(7.168)	(10.532)
θ_2^C	-828.143	-946.125	-936.677	-809.971	-982.192
	(54.355)	(55.805)	(43.601)	(43.529)	(53.997)
θ_3^C	1044.984	1192.942	1166.802	985.299	1297.072
	(77.154)	(72.219)	(60.379)	(68.198)	(72.910)
Utility from consumption					
γ	654.168	732.775	897.914	999.749	-968.707
	(107.001)	(115.752)	(116.798)	(86.836)	(125.573)

Notes: Estimates of the (cubic) inventory cost function and consumption utility. Specifications also contain six fixed costs parameters, one for each size. Column (1) is for all stores, columns (2)-(4) are for stores in low, middle and high income areas, respectively. Column (1) assumes a single inclusive value process for all (simulated) individuals. Columns (2)-(4) estimate separate inclusive value processes for each income area. Column (5) repeats column (1) without correcting for bundling. Standard errors are bootstrapped as described in the text.

likely to increase with inventory, these other considerations could result in non-monotonicities (e.g., smaller inventories may be easier to steal).

Inventory costs are consistently higher in low-income areas than high-income areas. This is also consistent with my survey: compared to respondents in high-income areas, respondents in low-income areas were more likely to view as an important consideration difficulty of storage (0.25 vs 0.31), inventory getting stolen (0.17 vs 0.32) and neighbors asking to borrow detergent (0.18 vs 0.43).[37] All else equal, higher inventory costs make it more difficult for these households to purchase large packages.

For comparison, Column (5) reports the dynamic parameter estimates when the static demand stage ignores bundling opportunities, so both prices and market shares are taken directly from the raw data without correction. This specification clearly produces a worse fit: for the same set of markets, the minimized objective function value is 99.22, compared with 80.92 for the specification in column (1). More importantly, ignoring bundling yields a negative utility on consumption. To see why, note that when bundling is ignored, on average, larger quantities appear systematically less popular, while smaller quantities appear more popular. The dynamic model must then rationalize why consumers buy small quantities, and this is achieved through a negative estimate of the marginal utility of consumption. In other words, the model infers that consumers dislike consuming more.[38] Overall, allowing

[37] Inventory getting damaged was viewed similarly in the two income areas (0.40 vs 0.41).

[38] Ignoring bundling also results in higher estimated inventory costs, as some of the misclassified small

for bundling produces more sensible parameter estimates and fits the data better.

Robustness. Figure A.22 shows inventory cost estimates from alternative specifications of the dynamic model. First, I replace the cubic inventory cost with a more flexible quartic inventory cost function. Second, I double the potential consumption shock (setting the mean to 100 g). I do not find any noteworthy changes in the dynamic parameter estimates in either specification.

Third, I restrict the maximum inventory level to 3.8 kg, which is the maximum observed in the survey data. In this case, the model is forced to rule out large inventory quantities. As a result, the predicted inventory cost increases sharply at this inventory level relative to the baseline specification.

Fourth, I note that the time period covered by the scanner data contains one month of December. This includes the Christmas holidays, when purchase patterns change somewhat (see Figure A.23). I include two dummy variables to control for a Christmas effect: a shift in the utility of the outside option and a shift in the utility of purchasing any detergent. The estimated coefficients are positive, indicating higher purchase propensity in this month. The inventory cost parameters do not change, indicating that this one-time demand shift does not affect intertemporal purchase patterns. Including these controls does not have a large impact on model fit.

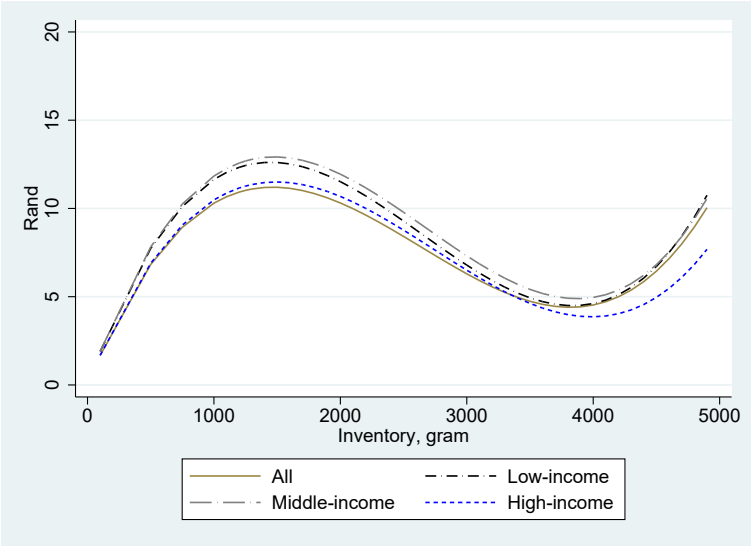
Fifth, recall that, in the main specification, I impose the constraint that households choose quantities that (together with their inventory) cover their consumption requirement. An alternative approach is to allow for stock-outs and estimate a stock-out cost. Implementing this, I estimate a large stock-out cost (171.71 Rand), while the remaining parameters remain unchanged. These two complementary approaches result in similar dynamic parameter estimates.

Finally, recall that the estimation computes starting inventories from a pre-period simulation starting from an inventory of zero. Instead of zero, I alternatively assign the average inventory from the survey (1.3 kg) to all households, or draw random starting inventories conditional on purchase from the survey. In both cases, there are no important changes relative to the main specification. This indicates that the pre-period simulation is sufficiently long and that the results are not sensitive to the initial inventory distribution.

Model performance To investigate the performance of the model, I use the estimated parameters to plot the model-implied purchase patterns over time. I simulate consumer choices for each market over 16 months. Figure A.23 shows the fit of the model separately for each package size. The model appears to capture well the over-time fluctuations in market

purchases make it appear that consumers avoid holding large inventories.

Figure 4: Inventory cost estimates



Notes: Inventory cost estimates from columns 1-4 in Table 6.

shares. Restricting attention separately to low or high income areas yields a similar picture.

6 Simulating the effect of small-format stores

6.1 Small-format stores in South Africa and elsewhere

South Africa has seen an emerging trend of large grocery chains opening small-format stores in low-income and rural areas.[39] Some of these stores are created out of shipping containers or large trucks instead of brick-and-mortar locations. Traditionally, these areas were served only by a collection of informal stores, and high travel costs meant that consumers had limited access to formal supermarkets, not unlike the situation of "food deserts" in the US (see Marshall and Pires (2018)). Similar models have been implemented globally, including mobile grocery vans serving rural populations in India, Hungary, Japan, Germany, and the US, as well as container-style stores introduced in the UK, China, Belgium, and the Netherlands in recent years.[40] Thus, this pattern is not specific to the South African setting,

[39]Shoprite CEO Pieter Engelbrecht described this as moving toward "formats closer to customers' homes" adding "These small-format stores offer a limited range of basic foods at everyday low prices to lower income consumers and are often located in previously underserved communities in South Africa." (https://supermarket.co.za/index.php/store-openings/4344-shoprite-plans-smaller-stores,-closer-to-shoppers-and-more-shops-in-containers).

[40]For example, https://www.esmmagazine.com/retail/lidl-hungary-launches-mobile-shops-to-serve-rural-communities-299497, https://progressivegrocer.com/kroger-brings-grocery-store-wheels-louisville-neighborhoods, https://nextshark.com/japan-familymart-launches-towable

but reflects a broader shift toward expanding retail access. This makes the analysis below potentially relevant for a wide range of settings.

Two common features of small-format stores are that (i) they are located closer to consumers, and (ii) they operate under tight space constraints. From consumers' perspective, the presence of these stores can reduce transportation costs, but this benefit is only realized if the store carries the product the consumer wants to buy. Since these small-format stores have limited space, they typically do not carry a full selection of products, and often restrict the package sizes offered. It is common for stores to carry the smallest package sizes based on the notion that consumers look to these stores to satisfy their temporary demand. However, this ignores the dynamic trade-offs highlighted by the model above, between quantity discounts, transportation costs, and inventory costs.

To study the impact of small-format stores, I perform 3 sets of counterfactual simulations. First, I study a scenario where consumers' transportation costs decrease. Second, instead of the full range of package sizes currently offered, I simulate the impact of offering only one of these sizes. Third, I simulate the combined impact of reduced transportation costs and limited product selection.

6.2 Counterfactual: reducing transportation costs

In this section, I consider a counterfactual that captures one aspect of the expansion of small-format stores, namely that stores are located closer to households and therefore transportation costs are lower.

In the model, transportation costs are partly captured through household demographics, in particular through the ownership of a car.[41] The majority of South African households do not own a car, and car ownership is highly heterogeneous. Based on the 2011 census, 67.83% of households do not own a car (61.23% in areas covered by Shoprite stores - see Table A.10). In low and high income areas these figures are, respectively, 66.99 and 49.51%. In my survey, 48% of high-income area respondents indicated using their own vehicle specifically when buying laundry detergent. The corresponding figure for low-income area respondents was only 1%.

To model a reduction in transportation costs, I consider what would happen if all households owned a car. I recompute the inclusive values using the estimated demand parameters, and re-solve the dynamic programming problem. Purchase probabilities are predicted for 16 months, keeping everything else, including prices, unchanged. I run this counterfactual

-convenience-store

[41] See Lagakos (2016) for a detailed analysis of the importance of households' car ownership for retailers' investment decisions.

separately for low-income areas (using the estimates in Table 6, column 2), and high-income areas (using the estimates in Table 6, column 4).

Figure 5 compares the distribution of market shares across package sizes in the counterfactual and at baseline. Since low-income households own cars at a much lower rate, the change in market shares between the baseline and the car-access counterfactual is larger for low-income households. Considering the two most popular sizes at baseline, the 1 kg and the 2 kg packages, the market share of the former decreases and the market share of the latter increases throughout. This makes sense as the quantity discount is much larger for the 2 kg package (Table 2): once the the transportation cost is reduced, households can more easily buy this larger package. Similarly, lower transportation costs also facilitate dynamic substitution between sizes. For instance, car access has a smaller impact on the market share of the 2 kg package when that package is relatively less attractive, as reflected by the lower baseline market share (month 7), and a larger impact in months when it is more attractive.

We also see that the market shares of smallest sizes decline. This finding goes against the common practice of small-format stores stocking primarily the smallest sizes in order to target low-income consumers. Based on my estimates, the reduction in transportation costs achieved by these stores *decreases* demand for the smallest sizes, as consumers become better able to take advantage of quantity discounts offered by larger packages.

At larger sizes, the pattern is more nuanced. The 3 kg size becomes more popular when transportation costs decrease, while the 5 kg size loses market share. This can be explained with the fact that the 5 kg size does not offer a substantial quantity discount relative to other large sizes, while it creates higher inventory costs. When transportation becomes less costly, consumers can shop more frequently and substitute toward the 3 kg size rather than purchasing the largest package.

Table A.24 shows the associated consumption and inventory changes under this counterfactual. With access to cars, inventory held by the households increases by 91 (34) percent in low (high) income areas compared to the baseline. Note however that I did not impose any constraint on inventory accumulation. As such, this overlooks factors like limited storage space, a preference for fresh products to avoid spoilage, etc. Inventory decisions by the households may involve additional considerations than those in the model - particularly for inventory levels not observed in the data.

Figure 5: Counterfactual: reduced transportation costs



Notes: Market shares over time at baseline and in the counterfactual when all consumers have access to a car. Average values shown for low-income and high-income areas separately. Simulations use the specification in Table 6 column 2 (low-income areas) and column 4 (high-income areas).

6.3 A planner's problem under limited shelf space

In this section, I use the estimated model to answer the following question: Assuming that a store can carry only one size of laundry detergent, what size would a social planner choose?[42]

I consider a social planner's problem of choosing one package size to maximize the sum of households' utility based on Equation (2). Let $V_h(\sigma_1|x)$ denote the solution of problem (2) for consumer h with the added constraint that $x_{jht} \in \{x, 0\}$. That is, the consumer is restricted to choosing a given package size x or the outside option of not buying the product. The planner then simply chooses x from the package sizes currently being sold (0.25, 0.5, 1, 2, 3, or 5 kg) to solve

$$\max_x \sum_h V_h(\sigma_1|x).$$

To compute $V_h(\sigma_1|x)$, I calculate consumers' utility streams over a period of 16 months using the prices observed in the data.

I consider the planner's problem separately by income area in order to model the possibility that a store could decide to tailor the package size to local demand. The potential value of such customization has recently been highlighted by Klopack (2024) in the context of US fast-food chains.

The results are in Table 7. Each column corresponds to a different scenario with the consumer's choice set restricted to the specific size (or the outside option). The table shows the average monthly values of consumption, inventory, purchase probability and utility level.

Not surprisingly, if only the smallest packages were offered, consumers' purchase probability would be very close to 1 in order to cover their consumption. Even then, consumption would be limited, and shifting to larger packages would therefore allow consumption to rise. We observe consumption increasing up to 2 kg, after which consumption is essentially flat, while inventory increases and purchase probabilities decrease, indicating that these larger sizes are purchased less frequently.

The baseline purchase probability for the most common 2 kg size is 0.23 in low-income and 0.18 in high-income areas. When this is the only size offered, this increases to 0.32 and 0.31, respectively. For larger sizes (3 kg and 5 kg), the baseline purchase probabilities are below 5 percent in both income groups, and increase substantially in the corresponding counterfactuals.

The results show that the offering only one of the smaller sizes yields lower utility, while

[42] The social planner perspective on the choice of package size could be different from the question of which package size maximizes profits. On the latter, see Björnerstedt and Verboven (2016) in the context of merger analysis.

the 2 kg, 3 kg, and 5 kg sizes provide very similar, and higher, utility levels in both areas. Stocking only one of the smallest sizes, the current practice in many small-format stores, is not utility maximizing.[43]

What would be the utility maximizing package size if the transportation cost went down? This corresponds to a situation where small-format stores simultaneously increase access but limit product offerings. To answer this question, I repeat the previous exercise assuming that all households have access to a car. Table 8 shows the results.

In low-income areas, car access increases purchase probabilities especially when 2 or 3 kg are the only sizes available (0.33 and 0.24 in Table 8 compared to 0.27 and 0.20 in Table 7). The corresponding changes in high-income areas are smaller. The changes in purchase probabilities translate into changes in expected utility. Accordingly, when stores are limited to offering a single size, the 2 and 3 kg packages provide the highest utility. The utility maximizing package size in low-income areas is 3 kg, while in high-income areas it is 2 kg. This reflects the higher consumption needs of low-income households as well as their higher taste for larger packages observed in the data. Overall, I find that the utility maximizing package size offered in a small-format store is similar to the most popular package size purchased in regular stores.

[43]This exercise maintains the discrete choice assumption that consumers purchase at most one product (once a month). In reality, if only the smallest sizes were offered, consumers might compensate by making more trips to the store. In this sense, the above exercise underestimates consumer utility from the smallest sizes. Note however that buying only small packages still means that the consumer foregoes the quantity discounts of larger packages. Therefore, buying twenty 250 g packages (for instance) will still yield lower utility than buying a 5 kg package.

Table 7: Counterfactual: only one package size offered

	Low-income area						High-income area					
	250 g	500 g	1 kg	2 kg	3 kg	5 kg	250 g	500 g	1 kg	2 kg	3 kg	5 kg
Consumption												
Average	23.70	42.92	58.93	64.08	66.41	66.46	23.78	42.26	58.57	61.93	61.64	61.00
Median	25.00	50.00	53.80	64.12	65.77	65.73	25.00	50.00	55.77	57.10	56.48	55.86
Inventory												
Average	1.19	5.18	29.86	168.40	178.45	201.16	3.18	12.49	73.75	209.93	197.99	219.18
Median	0.00	0.00	9.46	136.18	178.48	198.01	0.00	0.00	39.21	223.61	197.69	219.30
Purchase probability												
Average	0.95	0.86	0.59	0.32	0.22	0.13	0.95	0.85	0.59	0.31	0.20	0.12
Median	1.00	0.95	0.64	0.27	0.20	0.08	0.98	0.88	0.57	0.28	0.19	0.10
Average expected utility (normalized)												
Average	-1.21	-1.14	-1.07	-0.87	-0.87	-0.83	-1.11	-1.08	-1.02	-0.93	-0.94	-0.93

Notes: Each column corresponds to a different scenario where the consumer's choice set is restricted to the given size or the outside option. The simulations span a period of 16 months, with 50 individuals per store. Consumption and inventory are measured in 10 g. Expected utility is normalized relative to the average across all individuals and time periods in the given income area (low or high).

Table 8: Counterfactual: only one package size offered and reduced transportation costs

	Low-income area						High-income area					
	250 g	500 g	1 kg	2 kg	3 kg	5 kg	250 g	500 g	1 kg	2 kg	3 kg	5 kg
Consumption												
Average	23.55	42.17	58.59	66.90	70.06	66.77	23.73	41.87	58.15	62.86	62.97	61.22
Median	25.00	50.00	52.85	67.85	68.97	65.96	25.00	50.00	54.83	57.94	57.94	56.04
Inventory												
Average	0.59	3.07	28.08	250.74	269.79	204.22	1.54	8.31	59.72	263.93	237.74	220.37
Median	0.00	0.00	8.76	285.06	272.54	201.22	0.00	0.00	34.75	281.47	236.71	220.20
Purchase probability												
Average	0.94	0.84	0.59	0.34	0.23	0.13	0.95	0.84	0.58	0.31	0.21	0.12
Median	1.00	0.94	0.63	0.33	0.24	0.09	0.98	0.87	0.56	0.29	0.19	0.10
Average expected utility (normalized)												
Average	-1.44	-1.36	-1.14	-0.65	-0.63	-0.78	-1.09	-1.07	-1.02	-0.93	-0.94	-0.95

Notes: Each column corresponds to a different scenario where the consumer's choice set is restricted to the given size or the outside option, and all consumers have access to a car. The simulations span a period of 16 months, with 50 individuals per store. Consumption and inventory are measured in 10 g. Expected utility is normalized relative to the average across all individuals and time periods in the given income area (low or high).

7 Conclusion

This paper analyzes consumer choices between different package sizes of a storable household good in South Africa. Using countrywide scanner data on laundry detergents from the leading grocery chain, I document that purchases of the smallest packages are not especially common. Furthermore, large packages sometimes have higher unit prices, which means that some of the observed small purchases may actually reflect consumers bundling to save.

I estimate a dynamic demand model that allows for the fact that households may optimally exploit bundling opportunities. An original household survey on detergent consumption, inventories, purchases, and shopping behavior helps identify and interpret the model parameters.

I find that bundling matters: for example, in the median market there are over three times as many households who purchase 5 kg of detergent as households who purchase a single 5 kg package. This contrasts with the common view that households in developing countries tend to ignore unit prices and simply purchase products with the lowest package price. Model estimates also highlight the importance of transportation costs (car ownership) and inventory costs in shaping consumers' package-size choices, and this is consistent with descriptive patterns from my survey.

Using the estimated model to study the impact of small-format store expansion, I first consider a scenario in which increased proximity lowers consumers' transportation costs. I find that demand for the smallest sizes decreases, as consumers become better able to take advantage of the quantity discounts offered by larger packages. I also evaluate counterfactuals where small-format stores are constrained to carry only one size of detergent. I find that stocking only one of the smallest sizes, the current practice in many of these stores, is not utility maximizing. When stores are closer to consumers but have limited shelf space, larger package sizes may generate higher utility by allowing households to benefit from quantity discounts while facing lower transportation costs. From a policy perspective, this suggests unexploited opportunities for complementing the many existing interventions designed to help consumers save.

References

Abubakari, Sulemana, KP Asante, Misbath Daouda, B Kelsey Jack, Darby Jack, Flavio Malagutti, and Paulina Oliva. 2025. "Targeting subsidies through price menus: Menu design and evidence from clean fuels." Working paper.

Aguirregabiria, Victor. 1999. "The dynamics of markups and inventories in retailing firms." *The Review of Economic Studies*, 66(2): 275–308.

Arcidiacono, Peter, and Robert A Miller. 2011. "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity." *Econometrica*, 79(6): 1823–1867.

Atkin, David, Benjamin Faber, and Marco Gonzalez-Navarro. 2018. "Retail globalization and household welfare: Evidence from Mexico." *Journal of Political Economy*, 126(1): 1–73.

Attanasio, Orazio, and Elena Pastorino. 2020. "Nonlinear Pricing in Village Economies." *Econometrica*, 88(1): 207–263.

Attanasio, Orazio, Vincenzo Di Maro, Valérie Lechene, and David Phillips. 2013. "Welfare consequences of food prices increases: Evidence from rural Mexico." *Journal of Development Economics*, 104: 136–151.

Beatty, Timothy K.M. 2010. "Do the Poor Pay More for Food?: Evidence from the United Kingdom." *American Journal of Agricultural Economics*, 92(3): 608–621.

Bell, David R, and Christian AL Hilber. 2006. "An empirical test of the Theory of Sales: Do household storage constraints affect consumer and store behavior?" *Quantitative marketing and economics*, 4: 87–117.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4): 841–890.

Björnerstedt, Jonas, and Frank Verboven. 2016. "Does merger simulation work? Evidence from the Swedish analgesics market." *American Economic Journal: Applied Economics*, 8(3): 125–164.

Blattberg, Robert, Thomas Buesing, Peter Peacock, and Subrata Sen. 1978. "Identifying the deal prone segment." *Journal of Marketing Research*, 15(3): 369–377.

Chandon, Pierre, and Brian Wansink. 2002. "When are stockpiled products consumed faster? A convenience–salience framework of postpurchase consumption incidence and quantity." *Journal of Marketing research*, 39(3): 321–335.

Cohen, Andrew. 2008. "Package size and price discrimination in the paper towel market." *International Journal of Industrial Organization*, 26(2): 502–516.

DellaVigna, Stefano, and Matthew Gentzkow. 2019. "Uniform pricing in us retail chains." *The Quarterly Journal of Economics*, 134(4): 2011–2084.

Dillon, Brian, Joachim De Weerdt, and Ted O'Donoghue. 2021. "Paying More for Less: Why Don't Households in Tanzania Take Advantage of Bulk Discounts?" *The World Bank Economic Review*, 35(1): 148–179.

Erdem, Tülin, Susumu Imai, and Michael P Keane. 2003. "Brand and quantity choice dynamics under price uncertainty." *Quantitative Marketing and Economics*, 1: 5–64.

Fafchamps, Marcel, and Bart Minten. 2001. "Property rights in a flea market economy." *Economic Development and Cultural Change*, 49(2): 229–267.

Fafchamps, Marcel, and Ruth Vargas Hill. 2008. "Price transmission and trader entry in domestic commodity markets." *Economic Development and cultural change*, 56(4): 729–766.

Fehr, Dietmar, Günther Fink, and B Kelsey Jack. 2022. "Poor and rational: Decision-making under scarcity." *Journal of Political Economy*, 130(11): 2862–2897.

Fitzpatrick, Anne. 2017. "Shopping while female: Who pays higher prices and why?" *American Economic Review*, 107(5): 146–149.

Gowrisankaran, Gautam, and Marc Rysman. 2012. "Dynamics of consumer demand for new durable goods." *Journal of Political Economy*, 120(6): 1173–1219.

Griffith, Rachel, Ephraim Leibtag, Andrew Leicester, and Aviv Nevo. 2009. "Consumer shopping behavior: how much do consumers save?" *Journal of Economic Perspectives*, 23(2): 99–120.

Hendel, Igal, and Aviv Nevo. 2006. "Measuring the implications of sales and consumer inventory behavior." *Econometrica*, 74(6): 1637–1673.

Hendel, Igal, and Aviv Nevo. 2013. "Intertemporal price discrimination in storable goods markets." *American Economic Review*, 103(7): 2722–2751.

Karlan, Dean, Margaret McConnell, Sendhil Mullainathan, and Jonathan Zinman. 2016. "Getting to the top of mind: How reminders increase saving." *Management Science*, 62(12): 3393–3411.

Klopack, Ben. 2024. "One Size Fits All? The Value of Standardized Retail Chains." *The RAND Journal of Economics*, 55(1): 55–86.

Kremer, Michael, Jean Lee, Jonathan Robinson, and Olga Rostapshova. 2013. "Behavioral biases and firm behavior: Evidence from Kenyan retail shops." *American Economic Review*, 103(3): 362–368.

Lagakos, David. 2016. "Explaining Cross-Country Productivity Differences in Retail Trade." *Journal of Political Economy*, 124(2): 579–620.

Lal, Rajiv, and Ram Rao. 1997. "Supermarket competition: The case of every day low pricing." *Marketing science*, 16(1): 60–80.

Marshall, Guillermo, and Tiago Pires. 2018. "Measuring the impact of travel costs on grocery shopping." *The Economic Journal*, 128(614): 2538–2557.

Newey, Whitney K, and Kenneth D West. 1987. "Hypothesis testing with efficient method of moments estimation." *International Economic Review*, 777–787.

Orhun, A Yeşim, and Mike Palazzolo. 2019. "Frugality is hard to afford." *Journal of Marketing Research*, 56(1): 1–17.

Pesendorfer, Martin. 2002. "Retail sales: A study of pricing behavior in supermarkets." *The Journal of Business*, 75(1): 33–66.

Rao, Vijayendra. 2000. "Price heterogeneity and "real" inequality: A case study of prices and poverty in rural South India." *Review of Income and Wealth*, 46(2): 201–211.

Szabo, Andrea. 2015. "The value of free water: analyzing South Africa's free basic water policy." *Econometrica*, 83(5): 1913–1961.

Talukdar, Debabrata. 2008. "Cost of being poor: retail price and consumer price search differences across inner-city and suburban neighborhoods." *Journal of Consumer Research*, 35(3): 457–471.

Appendix to "Bundling to save: Analyzing package size choices in South African grocery stores"

Andrea Szabó
Economics Department, University of Houston
E-mail: aszabo2@uh.edu

July 5, 2026

Contents

A.1 Additional data characteristics **2**
A.1.1 Market characteristics . . . 2
A.1.2 Prices . . . 3
A.1.3 Store characteristics . . . 5
A.1.4 Census Data . . . 7

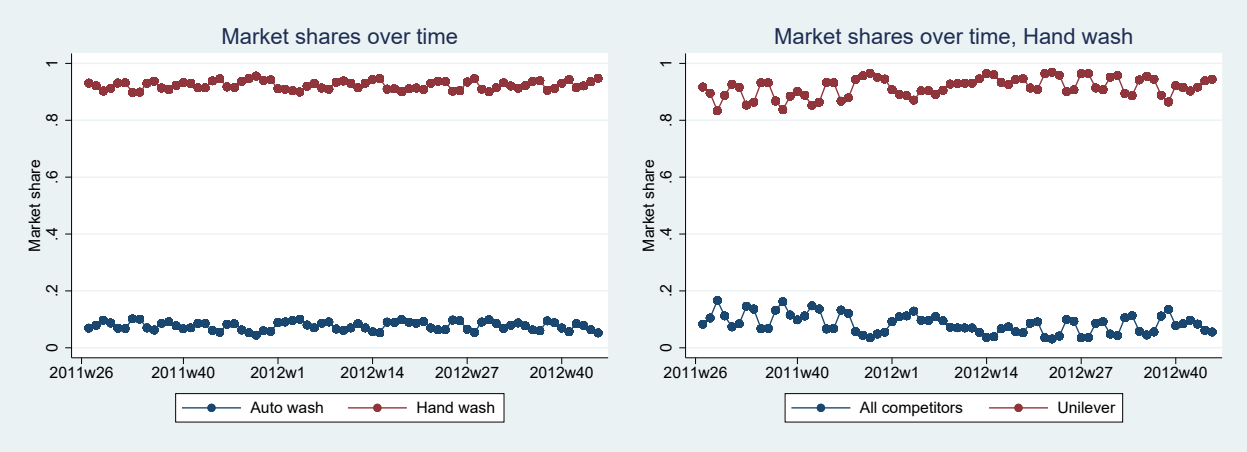
A.2 Estimation details **7**
A.2.1 Details of the static demand estimation . . . 9
A.2.2 Laundry detergent survey and value function approximation details . . . 11
A.2.2.1 Sampling . . . 11
A.2.2.2 Purchase, consumption, and inventory data . . . 12
A.2.2.3 Patterns in the survey data . . . 12
A.2.2.4 Value function approximation . . . 13
A.2.3 Details of the dynamic estimation . . . 14

A.3 Additional results **15**

A.1 Additional data characteristics

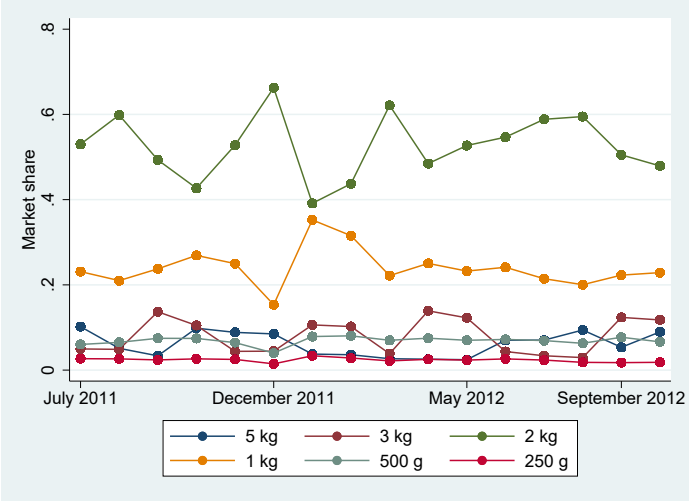
A.1.1 Market characteristics

Figure A.1: Market shares of hand-wash and automatic detergents



Notes: Market shares based on sales value. Left panel: Market share of all hand-wash and automatic detergents. Right panel: Market share of hand-wash detergents only, Unilever and all competitors.

Figure A.2: Market share of Unilever products by package size



Notes: Market share is sales value of the given package size across all stores in a given month divided by the sales value of all 14 Unilever products across all stores in the given month.

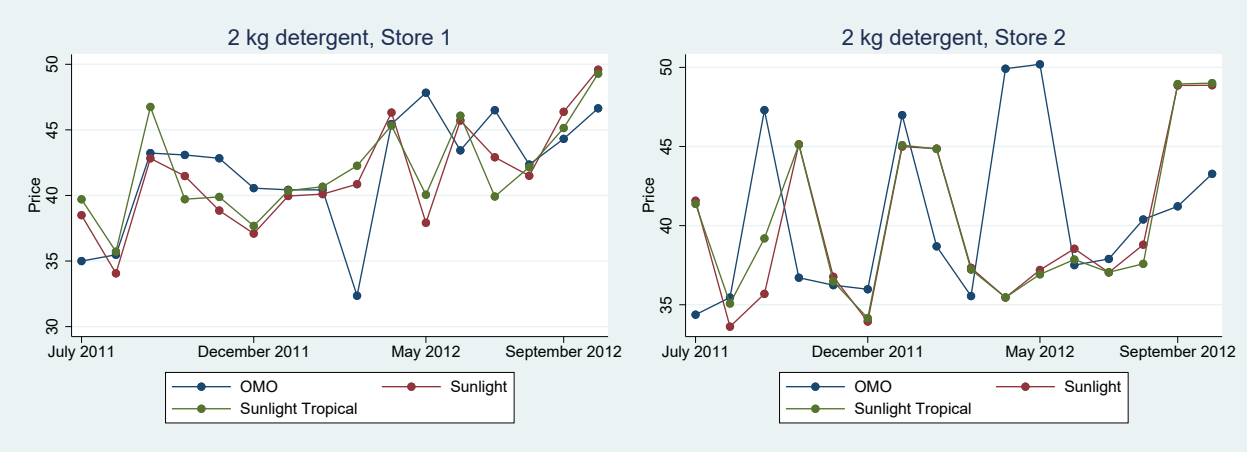
Table A.3: Availability of various brands and sizes

	Sunlight	Sunlight Tropical	OMO
250g	0.94	0.81	0.95
500g	1	0.77	0.99
1kg	1	0.96	1
2kg	1	0.99	1
3kg	0.99		
5kg	0.98		

Notes: Fraction of all markets (months × stores) in the data where the given product is available. N = 5255.

A.1.2 Prices

Figure A.4: Prices of 2 kg packages in selected stores



Notes: Prices of 2 kg packages in Rand in two selected stores.

Figure A.5 shows the evolution of prices for two distinct products over time in two selected stores. The first product (Sunlight 250 g, left panel) had only 1 week of promotion during the study period and the second (Sunlight 1 kg, right panel) had the most frequent promotions. In a given store, the price of the product with less promotion stays steady over a longer period. The price can be more than 10 percent higher in a different store. In both stores, we see an increasing price trend over the 16 month period, although the timing of the price increases is different.

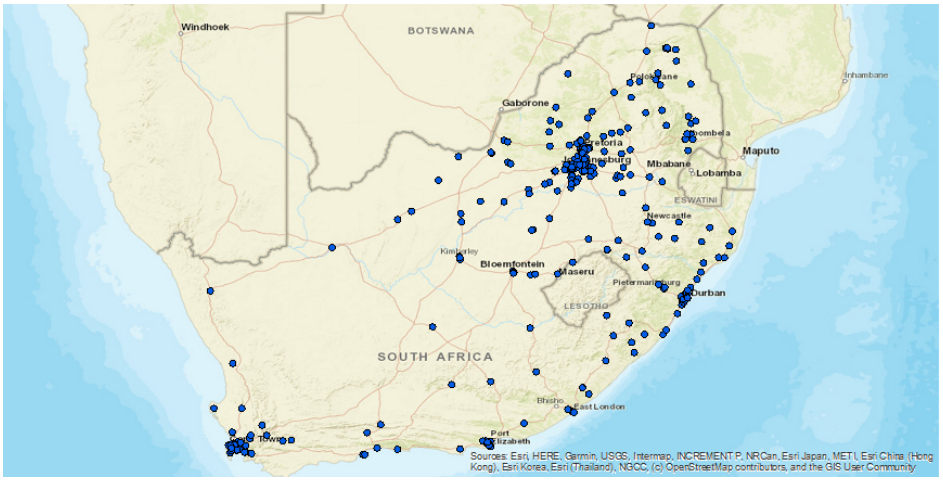
Figure A.5: Evolution of prices with and without frequent promotions



Notes: Prices in Rand. Left panel shows Sunlight 250 g. Right panel shows Sunlight 1 kg.

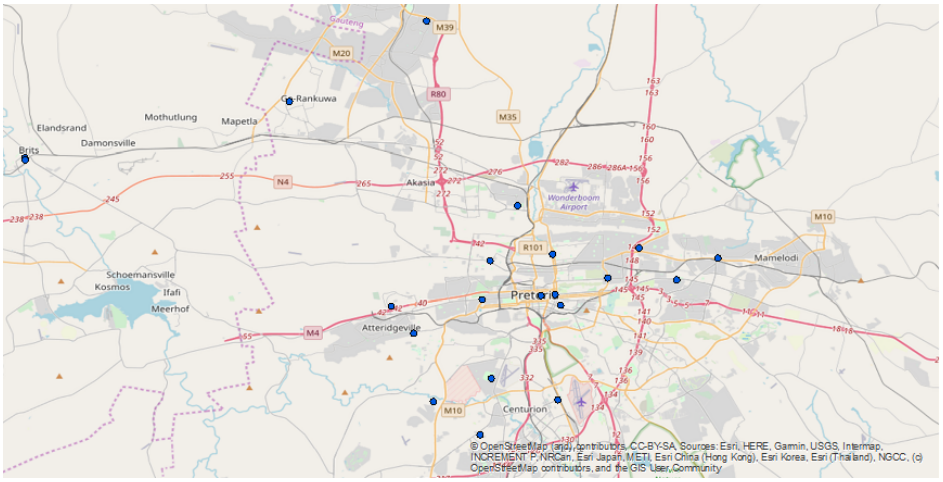
A.1.3 Store characteristics

Figure A.6: Location of stores in the sample



Notes: Locations based on GPS coordinates of the stores, collected from www.shoprite.co.za

Figure A.7: Location of the stores in the sample around Pretoria



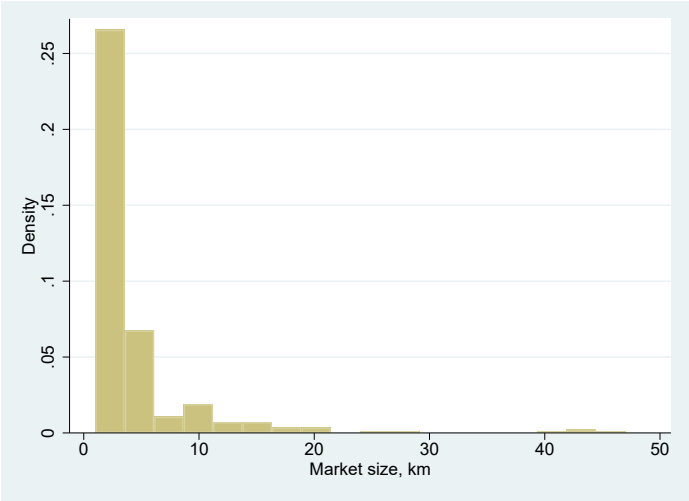
Notes: Locations based on GPS coordinates of the stores, collected from www.shoprite.co.za

Table A.8: Other store characteristics

	Percent
Low-income area	17.27
Middle-income area	43.64
High-income area	39.09
In a shopping mall	8.79
In city centre	24.85
Sunday closing time	
13:00	18.18
14:00	26.67
15:00	13.33
15:30	1.82
16:00	3.33
17:00	16.06
18:00	2.42
19:00	3.33
20:00	12.42
21:00	0.30
Closed on Sunday	2.12

Notes: Income area categories are from Unilever. Other characteristics based on store locator information at www.shoprite.co.za

Figure A.9: Distribution of the market radius



Notes: Distribution of market radius corresponding to each store. See the main text for the market size definition.

A.1.4 Census Data

The paper uses dataset "4.6. Household goods" from the "Census 2011: Community Profiles" CD. The data is accessed using SuperCross, a software provided by the South African Census.

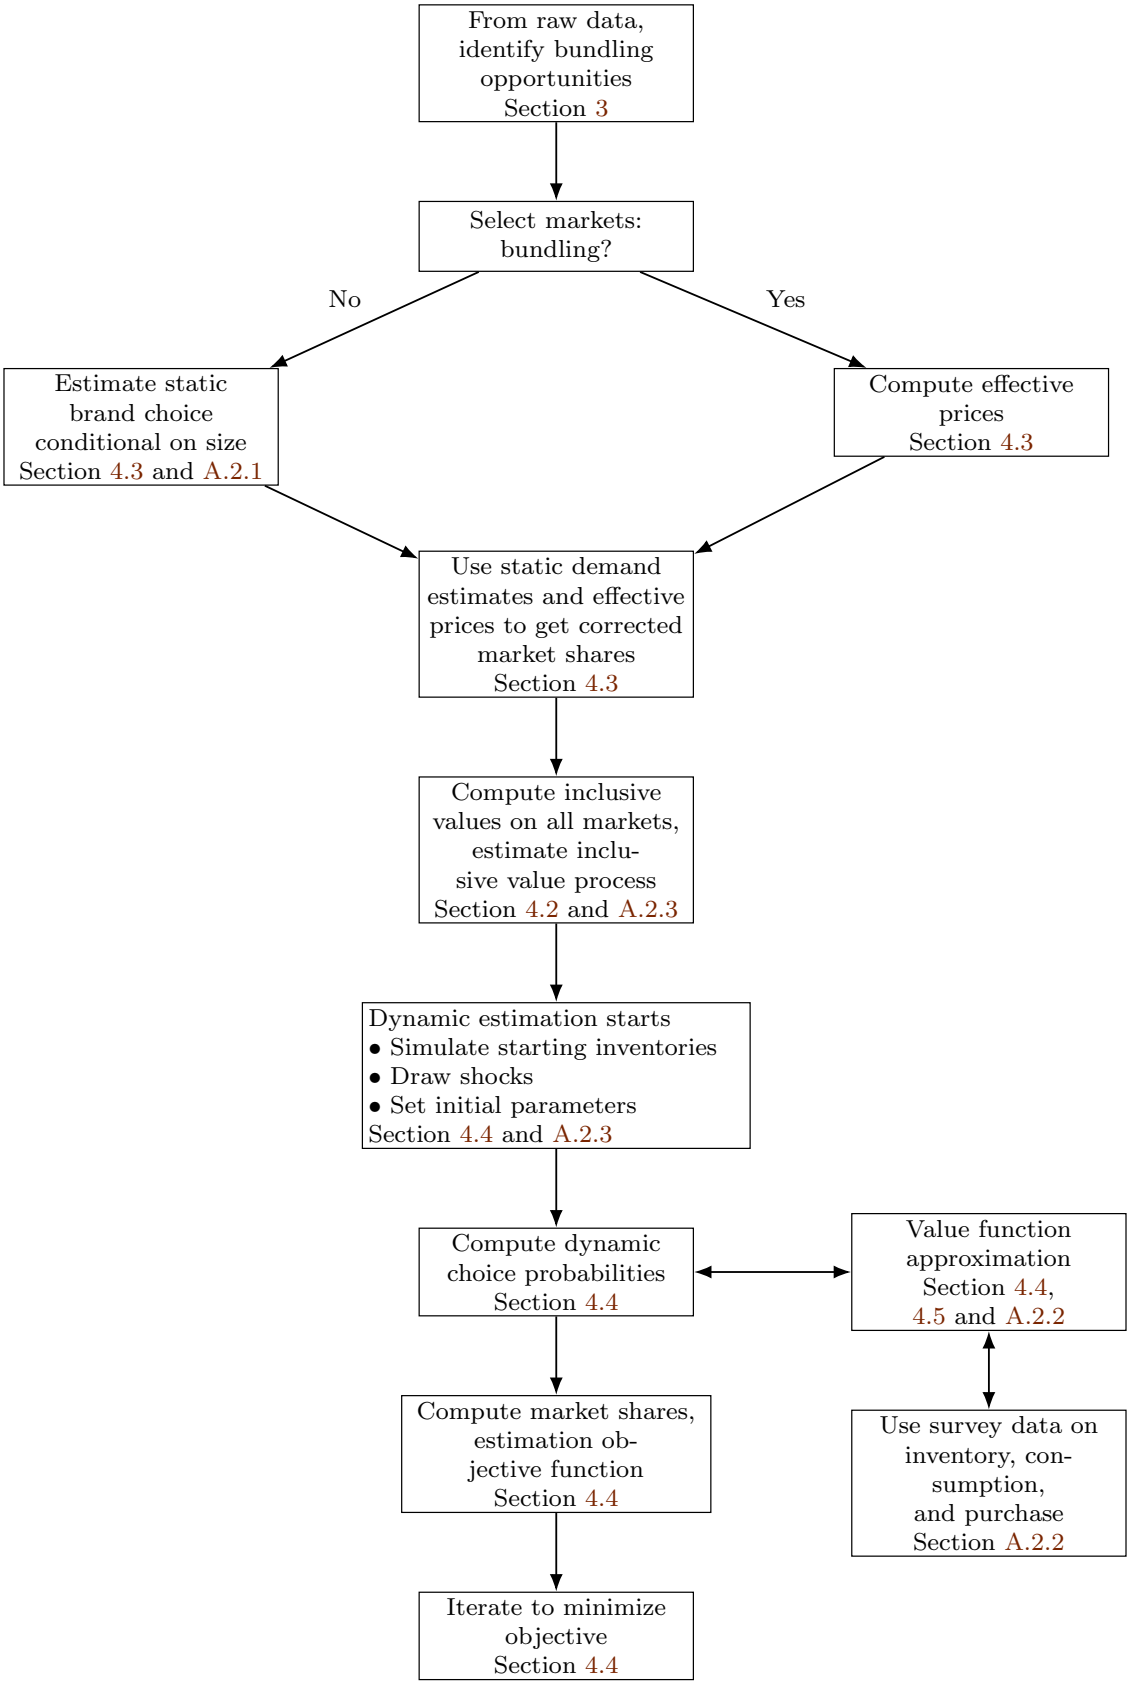
Table A.10: Household demographics

	South Africa	Markets in the sample
Monthly household income	9092.11	10410.27
Urban area	68.03	84.48
Male household head	57.61	59.89
Owns car and washm	22.40	28.73
Owns no car or washm	55.81	46.44
Owns washm only	12.02	14.79
Owns car only	9.77	10.04
Population group of household head		
Black African	77.71	69.7
White	10.96	13.49
Other	11.33	16.8
Type of dwelling		
House	66.97	67.41
Flat/apartment	12.45	18.3
Other (Informal dwelling, shack in backyard)	20.58	14.28
N	10,261,921	3,012,142

Notes: Based on 2011 South African Census. Households in sample markets are identified based on their distance from a store. See Section 2.2.4 for details. Income is monthly household income in Rand. Other variables are percentages of total.

A.2 Estimation details

Figure A.11: Estimation flow chart



A.2.1 Details of the static demand estimation

Estimation of the static demand parameters follows the Generalized Method of Moments (GMM) algorithm proposed by Berry, Levinsohn and Pakes (1995). Detailed treatments of the procedure can be found in Berry, Levinsohn and Pakes (1995) and Nevo (2001). Briefly, consider a dataset with information on product and market characteristics $\mathbf{a}_{jxt}$ and actual product shares $S_{jt}^{(x)}$. Berry et al. (1995) show that, for given $\mu_h(.)$, it is possible to numerically solve for $\boldsymbol{\xi}_{xt}$ from the equations $s_{jxt}(\mathbf{a}_{jxt}, \boldsymbol{\beta}_x, \boldsymbol{\xi}_{xt}, \mu_h(\mathbf{a}_{jxt})) = S_{jt}^{(x)}$, i.e., equating the model-predicted market shares to those observed in the data. Identification relies on moment conditions $E[\xi_{jxt}|\mathbf{Z}_{jxt}] = 0$ where the $\mathbf{Z}_{jxt}$ are suitable instruments, and estimation is via GMM.

Linear variables. The linear part of the utility, $\boldsymbol{\beta}_x \mathbf{a}_{jxt}$, includes price, brand dummies, store characteristics and average demographics. Store characteristics are included as they are likely to be correlated with consumer choices as well as prices and/or the availability of a given package size. For example, if stores located in city centers are more likely to offer smaller package sizes, then we need to control for the location of the stores. Similarly, I control for whether the store is in a larger shopping mall, opening hours on Sunday, and the distance to the nearest store to control for popularity and accessibility. For example, stores closed on Sunday and located in the city center might be used by people returning home from work and looking for smaller package sizes.

The linear part of the utility also includes a set of market-level average household demographics that exhibit little to no variation within markets. This lack of variation in some demographic variables is due to features of the South African setting. In this application, perhaps uniquely among similar discrete choice demand applications, I see very segregated markets. This arises from the segregation of South African neighborhoods, and it implies that there is little to no variation in, e.g., ethnicity within markets. In addition, as described in the paper, markets are defined as relatively small areas around each store, and consequently characteristics such as "urban" or "rural" do not change across individuals within markets. The full set of variables included is shown in Table A.18 under "Linear parameters."

Nonlinear variables. To model individual level heterogeneity among simulated consumers, I use demographic characteristics that vary within market and are relevant in the current context. I use household income, a binary variable indicating whether the household head is male, and four categories based on whether the household owns a washing machine and/or a car. Table A.10 illustrates the variation in these variables. For example, there is a meaningful portion of the population that owns a car but does not own a washing machine. Since I am able to draw individuals from the Census, I do not need to estimate the covariance of the different demographics and can instead use their empirical distribution.

The nonlinear part of the utility, $\mu_h(\mathbf{a}_{jxt})$, interacts these household demographics with price and with a constant. Interactions with the constant capture heterogeneity in the valuation of the outside good among individuals. The interpretation of the outside good in each BLP estimation is that the consumer decides to purchase no product (or a product not modeled here), or a different size. For example, the fact that a household has a car is likely to be an important determinant of which size they purchase.[1] For the full set of nonlinear variables included, see Table A.18 under "Nonlinear parameters."

Instruments. To identify the model, instruments are needed for two reasons: to control for price endogeneity, and to identify the nonlinear parameters. Instruments based on product characteristics, such as the differentiation IVs of Gandhi and Houde (2019), are not feasible in the current context. There are no product characteristics in $\mathbf{a}_{jxt}$ other than prices and brand dummies. None of the product characteristics change across markets (either across months or across stores) and there is little variation in the availability of different brands conditional on size across stores.

To instrument for price, similar to Nevo (2001), I construct instruments using the prices of the same product in different markets. Specifically, I rank a store's neighbors based on the magnitude of the correlation between the prices in the two stores. To determine the number of neighboring prices used as instrument, I successively include neighbors until there is no further increase in the F statistic.

To identify the nonlinear parameters, I use the market-level average of the demographic characteristics that enter the nonlinear part of the utility. This includes household income, male household head, and the fraction of households owning cars/washing machines. In some specifications, I found it useful to add the interaction of some of these instruments and the most relevant price instrument.

Other considerations. Demand specifications are the same for all sizes. For each market, I simulate 400 individuals, and the same individuals are used in estimating the demand for each size. Standard errors reported below are clustered at the store-month level, to allow for both heteroskedasticity and correlation of the shocks ξ_{jxt} across products within a market.

[1] I do not include interactions between brand dummies and household demographics. Including these would only be useful if there was variation in either the availability of brands across markets, or in market shares, correlated with the demographics included in the model. Here, there is little variation in the availability of brands across markets. In addition, all the products I consider are powdered hand-washing detergents, and the biggest differences between them (other than the brand) are likely to be factors such as scent. Preference for these factors is unlikely to be explained by the demographic variables that are typically available in Census data.

A.2.2 Laundry detergent survey and value function approximation details

The survey data is used for estimating the value function of the households, which in turn is an input for computing the dynamic choice probabilities in estimating the dynamic parameters of the model. Estimation of the value function is via value function approximation (see Section 4.5 in the paper), which requires describing the relationship between consumption, inventory, and purchase. Previous studies do not have data on consumption and inventory as these are not available in scanner data, which poses a challenge for identification. Even though my survey was conducted in a different time period than the scanner data, it helps alleviate these identification challenges as long as households with a given inventory and consumption choose similar sizes. The level or distribution of households' inventory, consumption, or purchase patterns (market shares) do not need to be the same at the time of the survey and the scanner data.

Out of the 300 respondents in my survey, 91.3% typically buy powdered detergents only.[2] This is very close to the market share from 2012 (see Figure A.1).

A.2.2.1 Sampling

The survey was entirely funded by the University of Houston. It was approved by the Human Subject Committee of the University of Houston, and was conducted in accordance with the standards of that institution regarding the ethical treatment of human subjects (Protocol number: 2626). Participation in the survey was voluntary and respondents could stop participating in the survey at any time. Only adults between the ages of 18 and 65 were asked to participate.

Surveys were collected from 300 households in December 2020 and again in March 2022. For logistical reasons, sampling had to be restricted to a single metropolitan area. I chose the area around Pretoria because of the diverse socio-economic characteristics of its population.

I first took all the stores in my dataset located within 20 miles from Pretoria (25 stores). I then extended this area 5 miles to the north to include more rural areas, resulting in a total of 27 stores. Of these 27 stores, 4 are located in low-income areas, 15 stores in middle-income areas and 8 stores in high-income areas. One of these stores was closed at the time of the survey due to damage from a tornado. Of the remaining 26 stores, I randomly selected a store from each of the three income areas. I selected the sample of households to be surveyed around each of these 3 stores as follows.

[2]Only 4 respondents stated that they typically buy liquid detergents and only 1 said that they typically use bar soap instead of detergent. 21 additional respondents buy a combination of powdered detergent and either liquid or bar soap.

For each store, I randomly selected 5 of the 10 closest small area layers of the 2011 South African Census. Surveyors were provided maps of each of these 5×3 areas. From each map, they selected a street intersection, and starting from there interviewed 5 households in each direction. Specifically, surveyors visited every 5th house in each direction, subject to the constraint that the final sample had to be stratified based on dwelling type recorded in the Census ("house," "flat," and "informal/other"). Surveyors recorded the GPS coordinate and a detailed description of the selected houses. Based on this information, surveyors visited the same houses during the second round of the survey.

A.2.2.2 Purchase, consumption, and inventory data

Recall that inventory levels at the start and end of a period, consumption and purchase during the period satisfy $i_{t+1} = i_t - c_t + x_t$. For each household, I define the period $[t, t+1]$ as consisting of the one month just before the survey visit. To infer the household's purchase x_t and inventory i_t from the survey, I proceed as follows. I compute the sum of the observed stock and consumption, call this SC. If SC is smaller than the size of the package kept by the household, X, then I assume that the package was purchased more than a month ago. In this case, I assign $x_t = 0$ and $i_t = SC$. If $SC \geq X$, then I assign $x_t = X$ and $i_t = SC - X$. I drop 2 observations where $i_t = 0$, 26 observations where the package was unopened, as well as 5 outliers whose inventory i_t is above 4 kg (over 3 standard deviations above the average of 1173 g). This results in 542 survey observations used in the value function approximation.

Based on the data, 32.87% of the households did not purchase detergent during the current month.

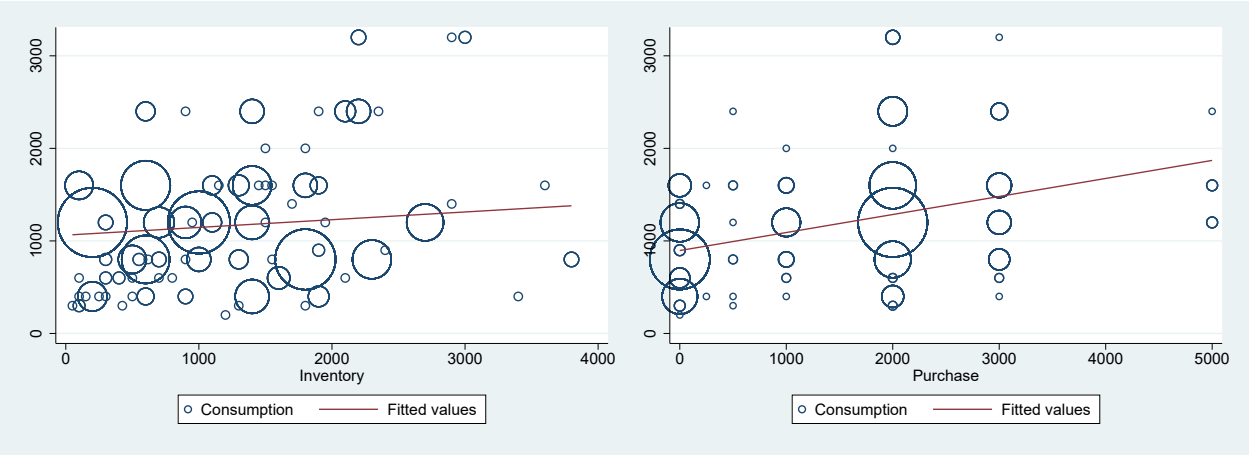
A.2.2.3 Patterns in the survey data

Three patterns are visible in the survey data. First, reported consumption is not correlated with inventory at home (correlation = 0.12, left panel of Figure A.12). This makes sense since the households are unlikely to use more detergent just because they have a new package at home.

Second, there is a stronger relationship between consumption and purchase size (correlation = 0.42, right panel of Figure A.12). Households who tend to buy larger packages consume more on average.

Third, average inventory during the first and the second round of the survey, conducted 16 months apart, is not statistically different. This is true both for all households and for households within income groups (Figure A.13). If average inventory is the same at the start and at the end of the period, this implies that average consumption equals average quantity

Figure A.12: Consumption as a function of inventory and purchase in the survey



Notes: Scatterplots and linear fits of consumption as a function of inventory and purchase, all in grams. N = 542.

purchased over this period period. (Note that 16 months is also the difference between the first and last period observed in the scanner data.)

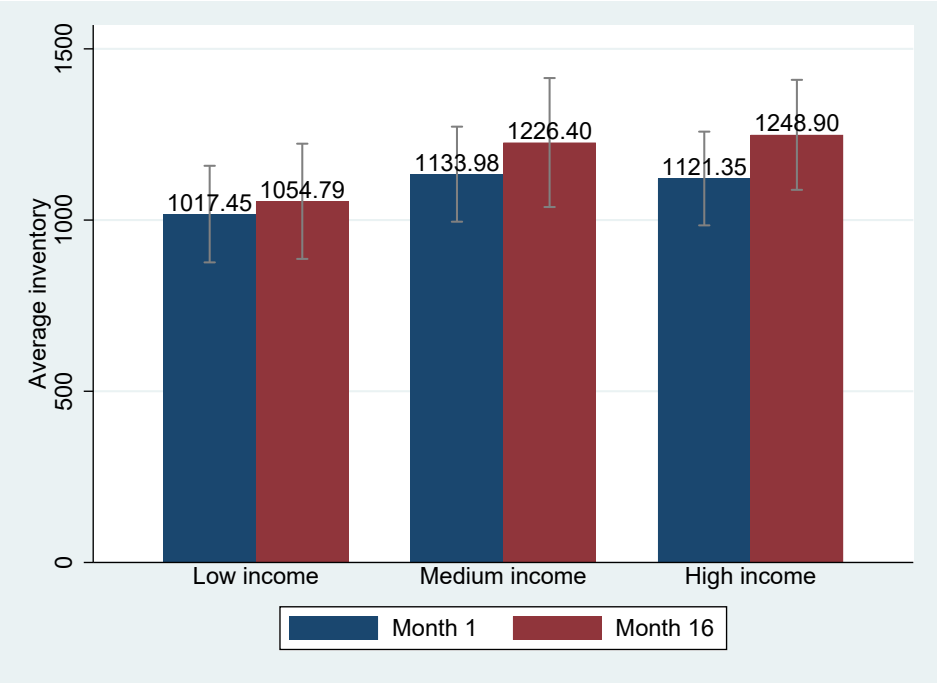
A.2.2.4 Value function approximation

For the value function approximation (equations (5) and (6) in the paper), I draw 400 markets from the relevant income groups (low only, medium only, high only, or all), and draw 50 individuals on each market, with corresponding inclusive values, from the BLP exercise. For each simulated individual, I draw a purchase x_t based on the market shares. The survey data provides information on the joint distribution of inventory i_t and consumption c_t conditional on purchase, and I draw inventory and consumption pairs for each simulated individual from this distribution.

Future inventory i_{t+1} is computed as $i_t + x_t - c_t$. Future values for inclusive values are predicted from the Markov process described in Section A.2.3 below. Based on the patterns described in the previous section, households' consumption needs are assumed to be a time-invariant constant, plus an idiosyncratic shock drawn from a log-normal distribution parametrized as in Hendel and Nevo (2006).

With these ingredients, for a given set of dynamic parameters, the value function can be approximated as described in Section 4.5 of the paper.

Figure A.13: Average inventory of households in the survey



Notes: Inventory in grams. Month 1 and 16 refer to the two surveys. p-values of Month 1 vs. Month 16 differences: 0.7520, 0.5018, 0.1353, respectively.

A.2.3 Details of the dynamic estimation

Markets and simulations. The dynamic estimation uses both markets with and without bundling opportunities. I drop only markets that do not have all brands and sizes, to keep the consumer's choice set constant (430 out of 5255 markets). I also drop markets where I do not observe at least 16 consecutive time periods (this can happen because some stores opened during the period and/or the store did not carry all sizes during the 16 month period). This results in 528, 1248 and 992 markets, respectively, for each income area.

From the static part of the estimation, I have 400 simulated consumers on each of these markets. To reduce the computational complexity of the dynamic estimation, I draw 50 of these consumers on each market. For each individual, the model predicts individual choice probabilities $Pr(x_{ht} = x)$ for each possible package size. I compute each of these choice probabilities for 10 draws of the consumption shock ν_{ht} and take the average.

Starting inventory. To obtain households' starting inventory, I follow Erdem, Imai and Keane (2003) and Hendel and Nevo (2006), who take as the starting inventory the inventory levels obtained from simulating the model for a number of periods from an inventory of 0 for all households. Specifically, for each set of dynamic parameters, I use the inclusive values and market shares to simulate the model from a starting inventory of zero for all households for

16 periods. I then take the resulting inventory levels as the starting inventory, and simulate another 16 periods which are used for estimation. The maximum potential inventory is set to twice the largest package size in the data.

Consumption. To compute consumption c_{ht} for the dynamic estimation, I first draw a sequence of 16 monthly purchased quantities based the observed market shares and take consumption to be the average of each sequence, c_h. This is based on the patterns described above, which suggest that consumption is correlated with purchase (but not inventory) and that, over a 16-month period, consumption is roughly equal to average purchase.

When simulating the model, I impose the constraint that current consumption (including the consumption shock) cannot be larger than current inventory, i.e., $c_{ht} + \nu_{ht} = min\{c_h + \nu_{ht}, i_t\}$.

Inclusive value process. Once the static demand parameters are estimated, I compute inclusive values using equation (3). I assume a first-order Markov process for the inclusive values of each package size. Since the demand parameters contain individual level heterogeneity, the inclusive values vary at the individual level as well. To reduce the number of processes to be estimated, I group individuals, and estimate either a single process for each size (Table A.19), or separate processes for each of the three income areas (top panel of Table A.20). I also show results where the grouping is done based on the ownership of washing machines and cars (middle panel of Table A.20). These are the variables which capture most of the individual level heterogeneity in the demand estimation, therefore I use the same variables to capture heterogeneity in the evolution of inclusive values. I assess the first-order Markov process assumption by including additional lags of the inclusive values of different sizes, as in Hendel and Nevo (2006) (bottom panel of Table A.20).

A.3 Additional results

Table A.14: Markets with no bundling opportunities

	Unique stores	Unique month	Markets (store × month)
250g	326	16	1064
500g	326	16	1088
1kg	326	16	1088
2kg	326	16	1088
3kg	324	16	1049
5kg	325	16	1029
All	330	16	5255

Notes: Total number of stores is 330, total number of months is 16. Total number of markets is 5255.

Table A.15: Correlation between bundling opportunities and market characteristics

	Low income	Middle income	Mall	City center	Sunday	% Black	% Flat
250 g	0.048	-0.005	-0.036	0.035	0.014	0.153	-0.022
	(0.024)	(0.015)	(0.018)	(0.018)	(0.015)	(0.032)	(0.036)
	0.25	0.25	0.25	0.25	0.25	0.26	0.25
500 g	0.053	-0.006	-0.035	0.029	0.012	0.147	-0.040
	(0.024)	(0.015)	(0.018)	(0.019)	(0.016)	(0.031)	(0.037)
	0.23	0.23	0.23	0.23	0.23	0.24	0.23
1 kg	0.004	-0.002	0.015	-0.008	0.000	-0.043	0.077
	(0.012)	(0.009)	(0.010)	(0.010)	(0.012)	(0.019)	(0.020)
	0.38	0.38	0.38	0.38	0.38	0.38	0.38
2 kg	0.020	-0.012	-0.012	-0.010	0.001	0.014	0.065
	(0.014)	(0.011)	(0.013)	(0.013)	(0.019)	(0.021)	(0.027)
	0.32	0.32	0.32	0.32	0.32	0.32	0.32
3 kg	0.018	-0.011	-0.018	-0.002	-0.006	0.033	0.044
	(0.014)	(0.010)	(0.013)	(0.012)	(0.013)	(0.018)	(0.026)
	0.40	0.40	0.40	0.40	0.40	0.40	0.40
5 kg	0.022	-0.012	-0.019	-0.002	0.002	0.057	-0.004
	(0.013)	(0.009)	(0.010)	(0.010)	(0.012)	(0.018)	(0.024)
	0.49	0.49	0.49	0.49	0.49	0.49	0.49
	% House	% Male HH head	% Urban	% White	No car or washm	Washm only	Car only
250 g	0.009	0.113	-0.081	-0.038	0.139	-0.412	0.889
	(0.039)	(0.103)	(0.037)	(0.033)	(0.042)	(0.096)	(0.250)
	0.25	0.25	0.25	0.25	0.25	0.25	0.25
500 g	0.011	0.143	-0.079	-0.033	0.135	-0.433	0.911
	(0.039)	(0.107)	(0.035)	(0.035)	(0.041)	(0.094)	(0.254)
	0.23	0.23	0.23	0.23	0.23	0.24	0.23
1 kg	-0.086	-0.085	-0.012	0.006	-0.040	-0.017	0.041
	(0.026)	(0.067)	(0.022)	(0.021)	(0.029)	(0.058)	(0.150)
	0.38	0.38	0.38	0.38	0.38	0.38	0.38
2 kg	-0.084	0.023	-0.050	-0.004	0.029	-0.149	0.291
	(0.032)	(0.080)	(0.022)	(0.024)	(0.033)	(0.065)	(0.165)
	0.32	0.32	0.32	0.32	0.32	0.32	0.32
3 kg	-0.072	0.059	-0.056	0.000	0.048	-0.158	0.309
	(0.030)	(0.076)	(0.021)	(0.024)	(0.028)	(0.063)	(0.158)
	0.40	0.40	0.40	0.40	0.40	0.40	0.40
5 kg	-0.019	0.063	-0.061	-0.020	0.081	-0.159	0.333
	(0.026)	(0.071)	(0.019)	(0.024)	(0.026)	(0.053)	(0.147)
	0.49	0.49	0.49	0.49	0.49	0.49	0.49

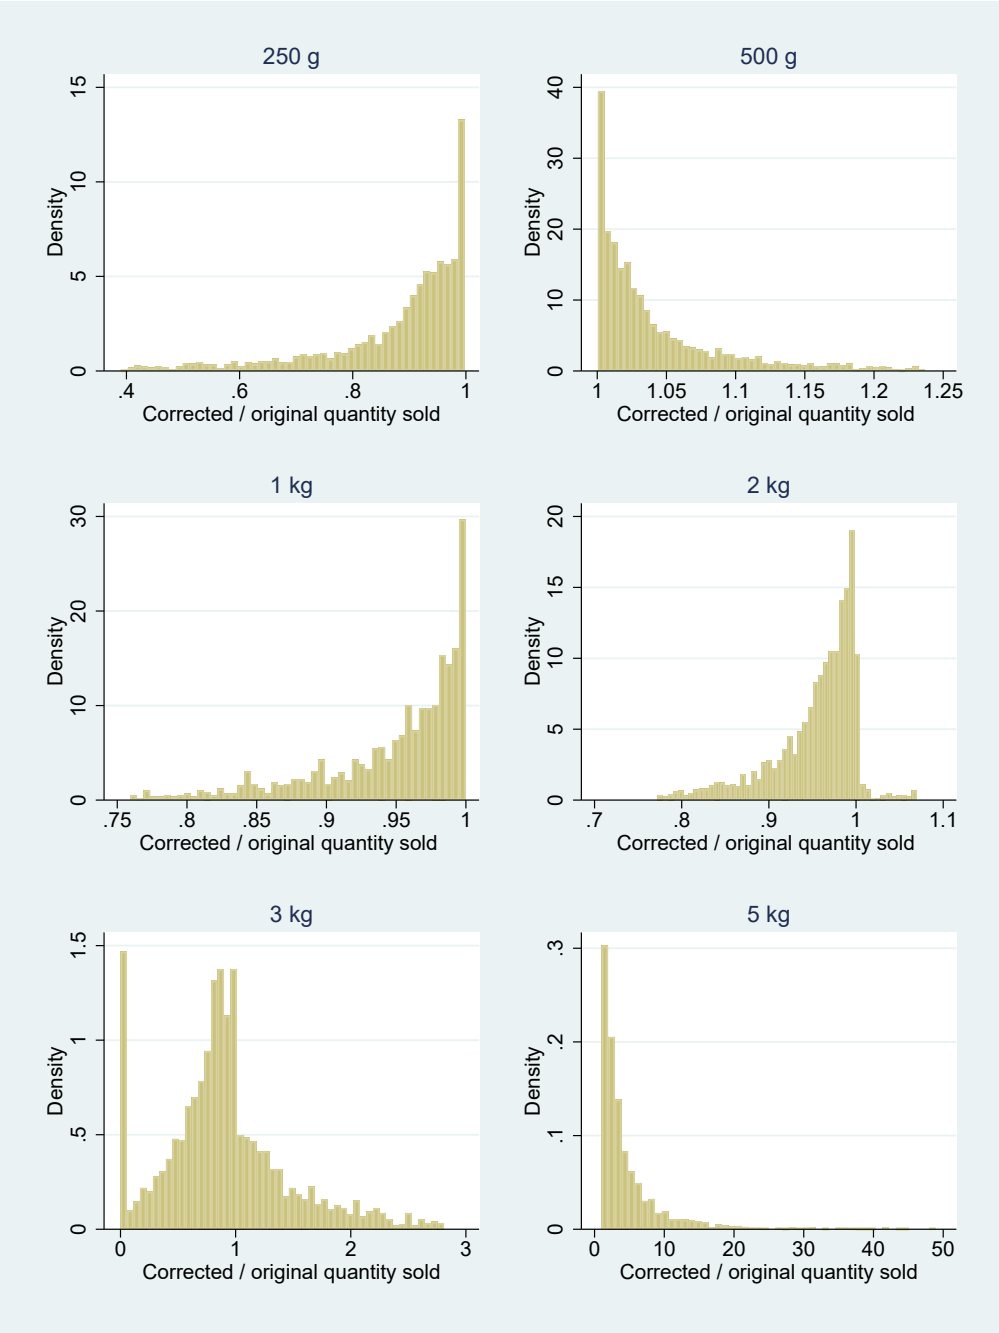
Notes: Regressions of an indicator for the presence of bundling opportunities on market characteristics. Observations are brands of the size listed in the row heading across all markets. Each cell is from a different regression. All regressions control for month and province fixed effects, market size, distance to closest store and number of other stores in the market. Number of observations for each size are shown in Table A.16.

Table A.16: Correlation between bundling opportunities and market characteristics

	250 g	500 g	1 kg	2 kg	3 kg	5 kg
Low income area	0.037	0.045	0.008	0.016	0.009	0.008
	(0.028)	(0.028)	(0.016)	(0.017)	(0.017)	(0.015)
Middle income area	0.018	0.021	0.003	-0.002	-0.002	-0.006
	(0.016)	(0.016)	(0.011)	(0.012)	(0.011)	(0.009)
Mall	-0.017	-0.016	0.006	-0.015	-0.013	-0.015
	(0.021)	(0.022)	(0.012)	(0.015)	(0.015)	(0.012)
City center	0.025	0.019	-0.006	-0.012	-0.002	-0.005
	(0.016)	(0.017)	(0.010)	(0.013)	(0.013)	(0.010)
Sunday hours	0.010	0.009	-0.001	0.001	-0.006	-0.000
	(0.016)	(0.017)	(0.012)	(0.020)	(0.014)	(0.013)
HH share Black	0.297	0.271	-0.081	0.007	0.051	0.059
	(0.071)	(0.070)	(0.036)	(0.037)	(0.039)	(0.035)
HH share White	-0.069	-0.068	-0.117	-0.074	0.004	0.015
	(0.086)	(0.081)	(0.040)	(0.056)	(0.055)	(0.054)
HH share flat	0.263	0.177	-0.059	0.067	0.035	0.083
	(0.088)	(0.082)	(0.076)	(0.099)	(0.072)	(0.061)
HH share house	0.163	0.101	-0.123	-0.030	-0.058	0.040
	(0.086)	(0.080)	(0.079)	(0.099)	(0.073)	(0.057)
HH share male HH head	0.512	0.506	-0.169	0.152	0.233	0.286
	(0.120)	(0.125)	(0.074)	(0.094)	(0.091)	(0.081)
HH share urban	-0.049	-0.044	-0.014	-0.053	-0.070	-0.059
	(0.040)	(0.039)	(0.026)	(0.027)	(0.029)	(0.025)
HH share no car or washm	-0.193	-0.205	-0.105	-0.043	-0.018	0.050
	(0.133)	(0.127)	(0.067)	(0.080)	(0.074)	(0.062)
HH share washm only	-0.391	-0.382	-0.198	-0.168	-0.020	-0.008
	(0.175)	(0.169)	(0.100)	(0.129)	(0.123)	(0.111)
HH share car only	-0.428	-0.294	0.017	-0.103	-0.024	-0.042
	(0.337)	(0.338)	(0.182)	(0.238)	(0.224)	(0.209)
Adj. R^2	0.26	0.24	0.38	0.33	0.40	0.49
Adj. R^2 controls only	0.25	0.23	0.38	0.32	0.40	0.49
N	14,189	14,483	15,548	15,696	5,199	5,167

Notes: Regressions of an indicator for the presence of bundling opportunities on market characteristics. Observations are brands of the size listed in the column heading across all markets. All regressions control for month and province fixed effects, market size, distance to closest store and number of other stores in the market.

Figure A.17: Ratio of corrected and original quantities sold



Notes: Distribution of the ratio of corrected and observed quantity sold across markets. Ratios above 1 indicate that some consumers purchased the given quantity by creating bundles of smaller packages; ratios below 1 indicate that some consumers used the given package size to create bundles of larger total quantities. 5-95 percentile range shown by package size. N = 2709, 2752, 1352, 3057, 2816, 2040, respectively.

Table A.18: Parameter estimates: static demand

	250g	500g	1kg	2kg	3kg	5kg
Linear parameters						
Price / 100	-97.290	-53.601	-36.182	-15.916	-10.113	-15.526
	(13.923)	(17.745)	(6.917)	(1.402)	(1.171)	(5.507)
Omo	0.492	0.371	0.614	1.109		
	(0.183)	(0.247)	(0.349)	(0.172)		
Sun	0.313	0.543	0.358	0.377	-3.877	2.630
	(0.151)	(0.222)	(0.256)	(0.138)	(5.294)	(2.335)
Constant	-3.798	-0.917	-2.136	0.917		
	(4.222)	(2.119)	(4.850)	(1.189)		
Low-income area	-0.148	0.164	-0.106	-0.108	-0.269	0.246
	(0.114)	(0.140)	(0.188)	(0.112)	(0.110)	(0.105)
Middle-income area	0.029	0.105	-0.062	-0.053	-0.116	0.163
	(0.086)	(0.102)	(0.135)	(0.069)	(0.070)	(0.081)
Mall	-0.235	-0.107	-0.044	-0.196	0.089	-0.297
	(0.094)	(0.108)	(0.179)	(0.094)	(0.117)	(0.115)
Center	-0.222	-0.147	-0.140	0.073	-0.088	-0.113
	(0.051)	(0.052)	(0.085)	(0.070)	(0.088)	(0.074)
Open Sunday	-0.052	-0.229	-0.012	0.122	0.059	0.096
	(0.058)	(0.048)	(0.089)	(0.056)	(0.066)	(0.080)
Neighboring stores	-4.808	-1.073	0.398	2.041	2.480	1.867
	(1.159)	(0.805)	(1.742)	(1.216)	(0.965)	(1.026)
Dist. to nearest store	-0.749	0.001	0.478	0.988	0.674	0.819
	(0.229)	(0.203)	(0.381)	(0.282)	(0.239)	(0.260)
HH share flat	-0.193	0.061	-0.745	-1.648	-2.172	-0.919
	(0.407)	(0.390)	(0.556)	(0.519)	(0.504)	(1.196)
HH share house	-0.306	0.094	0.004	-1.117	-1.195	-1.806
	(0.299)	(0.347)	(0.605)	(0.455)	(0.344)	(0.601)
HH share White	1.601	1.467	2.555	0.877	0.656	0.970
	(0.475)	(0.379)	(0.970)	(0.315)	(0.332)	(0.416)
HH share Black	-0.758	-0.419	-0.472	-0.213	-0.008	0.229
	(0.188)	(0.208)	(0.372)	(0.204)	(0.214)	(0.228)
HH share urban	0.320	0.263	0.113	-0.415	-0.693	-0.423
	(0.199)	(0.195)	(0.222)	(0.278)	(0.198)	(0.193)
Omo × HH share Black	0.110	0.679	0.204	0.201		
	(0.118)	(0.175)	(0.171)	(0.087)		
Sunlight x HH share Black	0.647	1.113	1.014	1.012		
	(0.107)	(0.176)	(0.155)	(0.074)		
Omo × HH share urban	-0.302	-0.222	-0.208	-0.746		
	(0.151)	(0.212)	(0.325)	(0.163)		
Sun × HH share urban	-0.320	-0.066	0.022	-0.258		
	(0.132)	(0.189)	(0.234)	(0.127)		
Omo × Low-inc area	-0.042	-0.350	-0.083	-0.273		
	(0.118)	(0.148)	(0.196)	(0.113)		
Sun × Low-inc area	0.098	-0.313	0.071	-0.112		
	(0.098)	(0.139)	(0.139)	(0.086)		
Omo × Mid-inc area	-0.077	-0.140	-0.113	0.009		
	(0.081)	(0.101)	(0.118)	(0.066)		
Sun × Mid-inc area	-0.043	-0.132	-0.096	0.014		
	(0.073)	(0.095)	(0.102)	(0.054)		

Cont'd on next page

Table A.18 cont'd

	250g	500g	1kg	2kg	3kg	5kg
Non-linear parameters						
Price × male	19.839	22.998	8.600	0.388	-0.446	6.251
	(8.907)	(15.893)	(2.759)	(1.261)	(0.973)	(5.303)
Constant × income	6.094	1.818	1.636	1.575	5.871	1.558
	(4.059)	(1.736)	(0.962)	(1.165)	(4.681)	(2.502)
Constant × no car or washm	4.237	2.175	4.919	-1.799	-5.297	-1.570
	(1.619)	(0.870)	(2.673)	(1.093)	(256.806)	(4.780)
Constant × washm only	4.027	2.701	6.896	1.831	2.682	1.538
	(1.346)	(0.793)	(2.651)	(0.675)	(1.688)	(1.202)
Constant × car only	-2.050	-6.024	5.490	2.276	2.925	1.396
	(61.514)	(693.013)	(2.613)	(0.778)	(1.198)	(0.817)
J	1.305	1.613	0.095	1.156	4.137	0.501
p-value	0.521	0.446	0.758	0.561	0.530	0.779
Newey-West D	48.441	51.073	52.501	20.960	29.922	33.820
p-value	0.000	0.000	0.000	0.001	0.000	0.000
N (market × products)	2798	2888	3193	3230	1049	1029
Unique markets	1064	1088	1088	1088	1049	1029
Unique months	16	16	16	16	16	15
Unique stores	326	326	326	326	324	325

Notes: Parameter estimates from the BLP model. Standard errors robust to heteroskedasticity and intra-market correlation in parentheses. All specifications contain province and quarter fixed effects. J is the overidentification test statistic with corresponding p-value. Newey-West D is a likelihood ratio test for the null hypothesis that the nonlinear parameters are jointly 0 with the corresponding p-value.

Table A.19: Inclusive values

	250 g	500 g	1 kg	2 kg	3 kg	5 kg
ω_{t-1}^{250}	0.964	0.020	0.012	-0.144	0.103	0.143
	(0.005)	(0.002)	(0.002)	(0.007)	(0.008)	(0.007)
ω_{t-1}^{500}	0.020	0.985	0.006	0.021	-0.044	-0.060
	(0.003)	(0.001)	(0.001)	(0.004)	(0.004)	(0.003)
ω_{t-1}^{1}	0.017	-0.006	0.976	0.134	-0.072	-0.058
	(0.002)	(0.001)	(0.002)	(0.005)	(0.005)	(0.004)
ω_{t-1}^{2}	-0.031	0.013	0.049	0.524	0.254	0.130
	(0.004)	(0.003)	(0.004)	(0.019)	(0.019)	(0.011)
ω_{t-1}^{3}	0.009	-0.002	-0.007	0.127	0.883	0.095
	(0.002)	(0.001)	(0.002)	(0.007)	(0.007)	(0.005)
ω_{t-1}^{5}	0.009	-0.011	-0.014	0.087	-0.022	0.705
	(0.004)	(0.002)	(0.003)	(0.009)	(0.012)	(0.011)
Constant	-0.030	-0.033	-0.029	0.236	-0.322	-0.524
	(0.007)	(0.005)	(0.007)	(0.027)	(0.026)	(0.027)
Adj. R^2	0.98	0.99	0.98	0.78	0.94	0.86
N	1,794,000	1,814,800	1,814,800	1,815,200	1,750,000	1,795,200

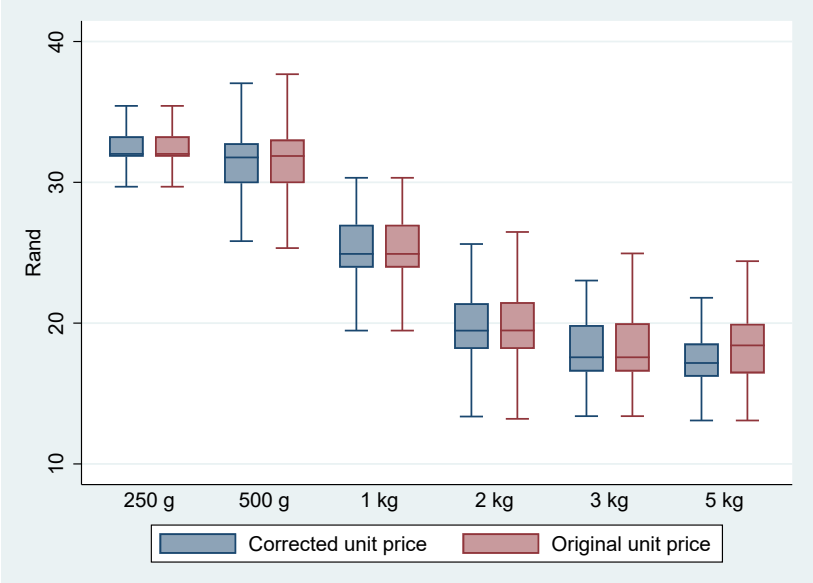
Notes: Estimates of the inclusive value process. The explanatory variables are lagged values of the inclusive value of every package size.

Table A.20: Fit of different inclusive value specifications

	250 g	500 g	1 kg	2 kg	3 kg	5 kg
Income area						
Low-income area	0.97	0.99	0.98	0.76	0.94	0.88
Middle-income area	0.98	0.99	0.98	0.79	0.95	0.87
High-income area	0.98	0.99	0.98	0.78	0.93	0.85
Ownership status						
No car or washm	0.97	0.98	0.94	0.54	0.80	0.78
Washm only	0.97	0.98	0.94	0.63	0.78	0.77
Car only	0.97	0.98	0.94	0.51	0.77	0.77
Car and washm	0.97	0.98	0.93	0.58	0.78	0.78
Alternative specifications						
Sum of five additional lags	0.98	0.99	0.98	0.84	0.96	0.90
Second lag	0.98	0.99	0.98	0.80	0.95	0.87

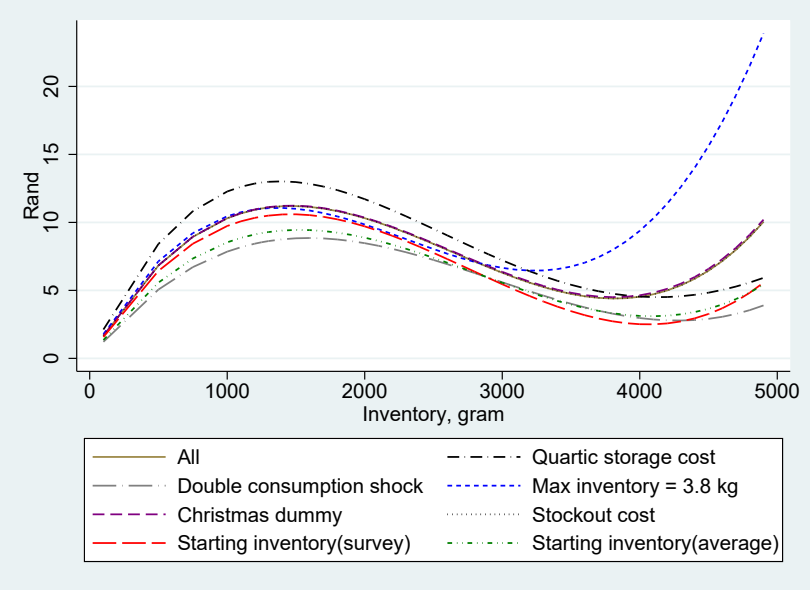
Notes: Adjusted R^2 values from different specifications of the inclusive value process. The top panel estimates separate processes by income area. The middle panel estimates separate processes by car/washing machine ownership. Both of these include the first lag of all sizes. On the bottom panel, the first row also includes the sum of five additional lags (t-2 to t-6) of all sizes, and the last row includes one additional lag (t-2) of all sizes.

Figure A.21: Average unit price by package size



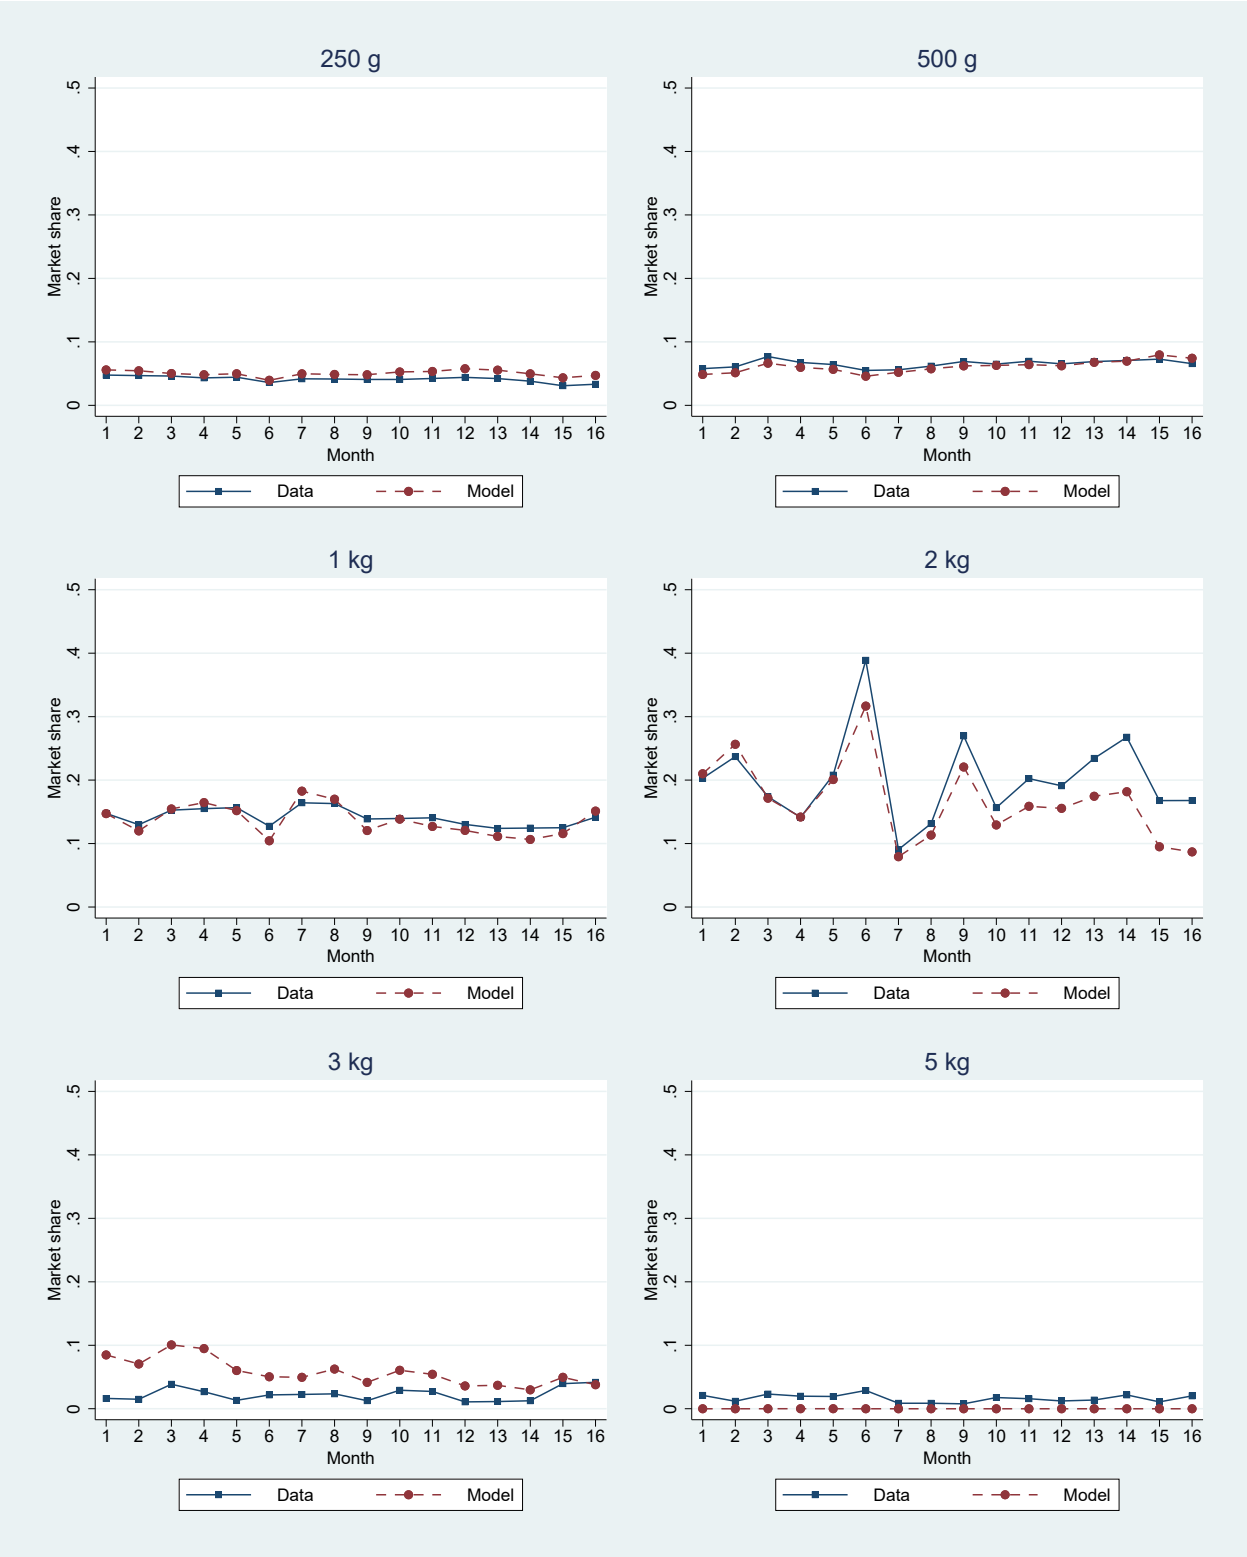
Notes: The plots show the interquartile range, median, min and max values of price per 1kg of detergent. Original unit prices refer to prices observed in the data. Corrected refers to prices corrected for bundling opportunities as described in the paper.

Figure A.22: Inventory cost estimates from alternative specifications



Notes: Inventory cost estimates under the different specification discussed in Section 5.3 in the paper. Estimates are for all income areas. "All" shows the main specification (from column 1 of Table 6).

Figure A.23: Model fit



Notes: Market shares of different sizes observed in the data and predicted by the model, based on column 1 in Table 6.

Table A.24: Household inventory and consumption when reducing transportation cost

	Low-income area		High-income area	
	Consumption	Inventory	Consumption	Inventory
Baseline	63.25	157.90	62.75	216.82
Access to car	69.75	302.77	63.18	290.91

Notes: Average consumption and inventory at baseline and a counterfactual where all households have access to a car. Consumption and inventory are measured in 10 g.

References

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4): 841–890. 9

Erdem, Tülin, Susumu Imai, and Michael P Keane. 2003. "Brand and quantity choice dynamics under price uncertainty." *Quantitative Marketing and Economics*, 1: 5–64. 14

Gandhi, Amit, and Jean-François Houde. 2019. "Measuring substitution patterns in differentiated-products industries." *NBER Working paper w26375*. 10

Hendel, Igal, and Aviv Nevo. 2006. "Measuring the implications of sales and consumer inventory behavior." *Econometrica*, 74(6): 1637–1673. 13, 14, 15

Nevo, Aviv. 2001. "Measuring market power in the ready-to-eat cereal industry." *Econometrica*, 69(2): 307–342. 9, 10