

# Why Do Retail Prices Fall During Seasonal Demand Peaks?\*

R. Andrew Butters<sup>†</sup>

Daniel W. Sacks<sup>‡</sup>

Boyoung Seo<sup>§</sup>

March 8, 2020

## Abstract

Examining all widely-sold products in a large, national scanner database, we find that seasonality in demand is large, pervasive across product categories, and heterogeneous in its timing. Yet at seasonal frequencies prices fluctuate little, and typically, countercyclically, falling as demand peaks. A natural explanation for this pattern is that demand becomes more elastic as its level peaks. We find, indeed, that for most seasonal products, demand becomes more elastic when demand levels peak. Quantitatively, the estimated seasonal elasticity changes can roughly rationalize observed countercyclical pricing. A likely mechanism for time-varying elasticities is seasonal shifts in the composition of buyers, as extensive margin responses account for nearly the entirety of seasonal demand shifts.

JEL codes: D40, L11, L81

Key words: seasonality, countercyclical pricing, demand, price elasticity of demand, optimal pricing

---

\*We are grateful to the Marketing Data Center at The University of Chicago Booth School of Business for granting access to the Nielsen Retail Scan and Homescan data. Researcher(s) own analyses calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher(s) and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. Information on availability and access to the data is available at <http://research.chicagobooth.edu/nielsen>. We acknowledge the Indiana University Pervasive Technology Institute for providing HPC resources that have contributed to the research results reported in this paper (<https://pti.iu.edu>). We thank, without implicating, Mike Baye, Joshua Bernstein, Stefano DellaVigna, Rick Harbaugh, Avery Haviv, Rupal Kamdar, Haizhen Lin, John Maxwell, Jeff Prince, Katja Seim, Joel Waldfogel, Matthijs Wildenbeset; seminar participants at Indiana University, the 17th Annual International Industrial Organization Conference, and the 46th Annual European Association for Research in Industrial Economics Conference for helpful comments; and Yaying Zhou for superb research assistance. Additionally, we thank Howard Freedman for helpful discussions on the business practices of retailers.

<sup>†</sup>e-mail: [rabutterm@indiana.edu](mailto:rabutterm@indiana.edu).

<sup>‡</sup>e-mail: [dansacks@indiana.edu](mailto:dansacks@indiana.edu).

<sup>§</sup>e-mail: [seob@indiana.edu](mailto:seob@indiana.edu).

Seasonal fluctuations in quantities can be enormous, so much so that seasonal fluctuations in GDP growth are larger than those associated with business cycles (Barsky and Miron, 1989). Seasonal fluctuations in prices, however, are small and may even be countercyclical; meaning prices appear to fall during periods of peak demand (e.g., Warner and Barsky (1995); MacDonald (2000); Chevalier et al. (2003); Hosken and Reiffen (2004a); Nevo and Hatzitaskos (2005)). As many instances of the seasonal fluctuations in quantities appear demand-driven, systematic countercyclical pricing seasonally would be puzzling. Exploring how microeconomic models can generate countercyclical prices is the focus of many recent studies in industrial organization, which focus on a handful of product categories with obviously seasonal demand such as soup and ice cream (Guler et al., 2014; Bayot and Caminade, 2015; Kwon et al., 2018; Haviv, 2019). Despite a growing number of explanations, no comprehensive assessment of the prevalence of countercyclical pricing exists, as compositional effects—arising from substitution to cheaper products during demand peaks—may account for the measured price declines (Nevo and Hatzitaskos, 2005; Perrone, 2016), and existing evidence focuses on only a small number of product categories. In this paper, we provide the first comprehensive assessment on the extent and path of seasonal fluctuations in demand and prices. We look at a broad set of widely available food products which span dozens of categories, we address compositional effects, and we investigate the ability of simple models of optimal pricing to rationalize the observed seasonal patterns.

Our investigation uses UPC-level data for 1,353 widely available UPCs (hereafter, “products”) in 38 categories, sold in grocery, drug, and mass merchandise stores in 49 states. The 24,500 stores in our sample collectively generate \$191 billion in annual revenue, and therefore constitute a meaningful share of the population of retailers. These data offer several advantages for studying seasonal fluctuations in demand and prices. First, as we will show, seasonality is a central factor in the demand facing these retailers, leading to large quantity fluctuations across a broad array of product categories. Second, the observed seasonal fluctuations in quantities are very likely demand driven. For the products we study, there is little seasonal variability in availability, quality, or new offerings, ruling out many possible supply-side explanations (Pashigian and Bowen, 1991;

Einav, 2007; Cooper and Haltiwanger, 1993). Finally, our scanner data are highly detailed, letting us observe prices at the store-week-product level, so we can quantify price changes that are free of compositional effects that result from aggregation in the time series or cross-section.

We have four results. Our first result is that seasonality in demand is pervasive, substantial, and heterogeneous. We measure seasonality in demand by regressing log quantity on a set of month-of-year dummies, as well as store-year fixed effects and log price (demeaned and interacted with month-of-year dummies). The estimated month-of-year effects therefore measure the month-to-month variation in quantities, adjusting for possible price differences. Our measure of seasonality in demand is the difference between the maximum and minimum quantity month.<sup>1</sup> We find that, for the median product at a fixed price, demand varies by 24 log points from seasonal trough to peak, and by 10 log points at the 10th percentile. These fluctuations are large relative to, for example, quantity fluctuations induced by the business-cycle. Using business-cycle induced fluctuations as a benchmark, we estimate that three-quarters of products are seasonal in the sense that seasonal fluctuations are larger than the fluctuations due to the Great Recession. We find clear seasonality in demand in obvious categories like soup and frozen novelties, but also in less obvious categories like cookies. Not only are seasonal fluctuations in demand large, they are heterogeneous both in their magnitude—cookies are less seasonal than soup—and also in their timing. While many products exhibit demand peaks in January, several peak in June, July, and August.

Our second result is that seasonal fluctuations in prices are small and typically, though not always, countercyclical. We measure seasonality in prices by comparing how prices change from the month of trough demand to the month of peak demand. A negative price change means that prices fall when demand peaks and therefore indicates countercyclical pricing. We find that roughly two-thirds of seasonal products exhibit countercyclical pricing. Our approach looks within store-product-year cells, and therefore is free of compositional effects. In fact, we show that not adjusting for compositional shifts within product

---

<sup>1</sup>We address price endogeneity by using the within-chain Hausman (1996) instruments proposed by DellaVigna and Gentzkow (2019), and we adjust for measurement error in the month-of-year effects with an Empirical Bayes shrinkage procedure.

categories does result in larger measures of countercyclical prices. As further evidence of countercyclical pricing, we find that the peak demand month is particularly likely to be the trough price month, meaning not only do prices fall when demand peaks, but they often fall to their lowest point over the year. Despite clear evidence that countercyclical pricing is common, however, we find that seasonal price changes are small: the median price change (between demand trough and peak) is only about 2 log points. Furthermore, these price changes come from a combination of more frequent promotional prices (i.e., sales) and lower base prices, without any meaningful change in the depth of promotional pricing.

Countercyclical pricing is a puzzle for supply-and-demand models of pricing, but it is easily rationalized if grocers have market power. In that case, prices depend on demand elasticities. And, countercyclical pricing can be rationalized if demand elasticities fall (i.e., become more negative) during seasonal demand peaks. Our third result shows that for a majority of seasonal products, demand indeed becomes more elastic during seasonal demand peaks, and demand elasticities are especially likely to be at their most negative in precisely the months when prices are lowest. Quantitatively, we find that the seasonality in elasticities matches well with the seasonality in prices: looking across product categories, observed prices move roughly one for one with the predicted change in prices, given estimated changes in elasticities. Overall, these results suggest that seasonal pricing patterns—small and typically countercyclical price changes—are roughly optimal.<sup>2</sup> Furthermore, we estimate that the differences between the observed price changes and optimal prices implied by our simple benchmark are economically small, in the sense that retailers are giving up little profit.

Our fourth and final result probes the mechanism underlying the observed countercyclical elasticities: changing consumer composition. Using household panel data, we decompose observed seasonal fluctuations into an extensive margin component (does the household purchase?) and an intensive margin component (given purchase, how much is purchased?). We find the extensive margin is large, explaining nearly the entirety of

---

<sup>2</sup>The optimal prices of our simple benchmark indicates scope for even larger countercyclicity in prices than what we observe.

seasonal fluctuations in demand for the median category. The large extensive margin response means that the identity of the marginal customer during periods of peak demand is potentially quite different than the marginal customer during periods of trough demand. To the extent that the marginal customer in periods of peak demand has a higher price sensitivity than the marginal customer in other times of the year, this would lead to retailers facing more elastic demand during these periods.

These results contribute to three literatures. Most directly, we contribute to a literature in industrial organization that explores seasonal fluctuations in demand and prices. Important early contributions such as Warner and Barsky (1995) and Chevalier et al. (2003) show that prices fall during demand peaks like during Christmas and weekends, but offer differing explanations. Warner and Barsky (1995) argue that falling prices arise from increased search intensity, while Chevalier et al. (2003) conclude that they reflect a “loss-leader” pricing strategy. Generalizing the results of these papers is difficult, however, given their focus on a small number of product categories and retailers in a single market. MacDonald (2000) and Hosken and Reiffen (2004a) examine the price variation across a larger set of product categories, and find evidence that countercyclical pricing is tied to informational advertising. Nevo and Hatzitaskos (2005) call into question the initial findings of this literature, showing that much of the observed countercyclicity in prices paid represents substitution towards cheaper products during periods of peak demand.<sup>3</sup>

More recent work in this literature carefully examines individual product categories with highly seasonal demand (such as ice cream, soup, champagne, or tuna), in an effort to recover aspects of demand that provide incentives for countercyclical pricing (Guler et al., 2014; Bayot and Caminade, 2015; Perrone, 2016; Kwon et al., 2018; Haviv, 2019). Explanations from this work center on the role of heterogeneous consumer segments (Guler et al., 2014; Bayot and Caminade, 2015), stockpiling and endogenous consumption (Kwon et al., 2018), and changing incentives to search (Haviv, 2019). These explanations, as well as many of the earlier studies, typically share the reduced-form implication that the demand facing retailers becomes more elastic during seasonal demand peaks, with some

---

<sup>3</sup>In contemporaneous work, Meza and Sudhir (2006) finds some of the observed countercyclical pricing in the tuna and beer categories result from manufacturer trade promotions.

emphasizing the role of the changing composition of consumers.

We complement and extend this literature in two ways. First, our analysis is much more general, examining the widest variety of products and categories. This breadth reveals that seasonality in demand and the pattern of countercyclical pricing extends beyond the obviously seasonal categories, like soup and ice cream. Second, by adopting an approach to estimating time varying demand elasticities that is applicable to a wide variety of categories, we are able to confirm that explanations leading to the reduced-form implication that demand becomes more elastic during seasonal demand peaks are the most likely to be successful in explaining the majority of countercyclical pricing.

Second, we contribute to an emerging literature exploring uniform pricing puzzles. Several papers have documented that prices appear highly uniform despite large demand variation (Orbach and Einav, 2007; Cho and Rust, 2010; Hitsch et al., 2017; Gagnon and López-Salido, 2019). This literature includes the pioneering work by DellaVigna and Gentzkow (2019), whose empirical approach we follow closely. Many of these papers argue that the observed low variability in pricing is inconsistent with profit maximization (Orbach and Einav, 2007; Cho and Rust, 2010; DellaVigna and Gentzkow, 2019). Our results show that large variations in the level of demand need not be accompanied by large variation in its elasticity, meaning that near-uniform pricing can be optimal even in the face of large demand fluctuations.

Finally, we contribute to a literature in macroeconomics that measures the frequency of price adjustment. The frequency of price adjustment—or the stickiness of prices—has important implications for the welfare consequences of business cycles, optimal monetary policy, and is a crucial feature of a large class of macroeconomic models (e.g., Calvo (1983)). In their seminal work, Bils and Klenow (2004) found a median (monthly) price change frequency of 21%—or an implied duration of about four months, casting doubt on sticky-price models ability to explain inflation behavior. Subsequent work argued that while promotional prices change often, base prices change much less often (Hosken and Reiffen, 2004b; Nakamura and Steinsson, 2008, 2013). Our results show that promotional price changes are in part driven by highly predictable seasonal fluctuations in demand (and its elasticity). Retailers and wholesalers could anticipate and contract on

these price changes months or years in advance. The frequency of promotional price changes, therefore, may not be informative of how retail prices would respond to sudden and unanticipated aggregate shocks such as interest rate changes, and therefore may not be informative for the validity of sticky-price models (Nakamura and Steinsson, 2013).

The remainder of this paper is structured as follow. We describe the data in Section 1. In Section 2, we present a simple model of pricing and use it to motivate our empirical strategy, which we discuss in detail. We document seasonality in demand, prices, and elasticities in Section 3. In Section 4, we compare profits under observed prices to their optimal level. In Section 5, we discuss extensive margin responses using household-level data. Section 6 concludes.

## 1 Data

The primary dataset used for the analysis comes from the unbalanced panel of retail stores reporting revenues and units sold to Nielsen Retail Scanner data (RMS) over the period of 2006–2014. The Nielsen RMS data contain stores from across the country, making it possible to study the seasonal fluctuations in demand and prices for the country as a whole and for most products available in retail stores. The dataset contains units sold and revenue for every week and product, which is defined as a combination of UPC and UPC version, for as many as 41,000 stores. Our baseline sample selection approach follows DellaVigna and Gentzkow (2019) closely.<sup>4</sup>

### 1.1 Sample selection

**Selecting retail formats, chains, and stores** We apply an initial set of screens of chains and stores to obtain our analysis sample. Appendix Table A.1 shows how the number of stores and chains falls as we impose these screens. We start by excluding all convenience stores and gas stations, leaving food, drug, and mass merchandise store types in our sample. Excluding these stores helps avoid the “zero” quantity problem, and also

---

<sup>4</sup>We are grateful to them for sharing their list of products studied. As we follow the same sample selection procedure, we study the identical set of products.

Table 1: Store and product-level summary statistics

	Stores	Chains	States	Total yearly revenue
Food stores	9,834	69	49	\$134bn
Drug stores	11,180	2	49	\$23bn
Mass merchandisers	3,436	6	49	\$34bn

	No. of products	No. of modules	Yearly revenue by store	Weekly average price
Sample of products	1,356	38	\$533K (627K)	\$1.87 (0.51)

Notes: This table reports final sample characteristics. Yearly revenue is aggregate sales divided by 9 years. Yearly revenue by store is total revenue for a given store-product, divided by the number of years the store is in the sample.

lets us focus on a sample of stores in which grocery sales represent a major component of store sales. Next, we impose restrictions at the chain and store level. The primary purpose of these restrictions is to ensure that we correctly measure stores' chain affiliation. This is important because our instrumental variable strategy uses prices charged by other stores in the same chain. We define a chain as a unique combination of parent-code (company) and retailer-code found in the Nielsen dataset. This has the effect of treating stores with different banner names, but the same parent holding company, as different chains. We drop stores that switch chains, and we drop stores that are in the data for only one year. Next, we drop stores which do not show up in Nielsen Consumer Panel (HMS), because we use consumer's store visits to RMS stores to decompose extensive and intensive margin changes in seasonal demand. We then drop a chain if it does not appear in at least eight of the nine years of our panel. Then, we limit the sample to ensure that all chains are valid. Sometimes a given retailer-code is assigned to multiple parent-codes; in such cases, we keep only the modal parent-code-retailer-code combinations, and we drop all parent-code-retailer-code combinations in which the modal combination accounts for less than 80 percent of stores. Finally, to obtain the final sample, we limit to chain-years in

which we are always able to impute missing prices (as in Section 1.2). Our final sample is comprised of stores which collectively generated \$191 billion yearly revenue, about 85% of total revenue in the RMS data. Summary statistics at the store- and product-level are reported in Table 1.

**Products** Instead of selecting products which exhibit stark seasonal fluctuations in demand and/or price, we examine all widely available retail products. We define the availability of a product in a given year and format (retail/mass merchandise/food) as the number of store-weeks with positive sales divided by the number of store-weeks with any sales in the category. We define widely available products as ones with greater than 80 percent availability in at least one format and year. We focus on widely available products to avoid the problem that prices are missing when quantity is zero. Alternative solutions to this problem would involve aggregating over heterogeneous products, resulting in possible compositional effects. Our definition of widely available products means that private label (or store brand) products are not included in the analysis. Our final sample comprises of 1,356 products, across 38 categories (Nielsen’s modules).<sup>5</sup> These products on average account for half a million dollars of revenue per store-year, or 5.79% of total revenue.

## 1.2 Quantities, prices, base prices, and promotions

Quantities—units sold—are directly reported in the Nielsen data. In weeks with zero units sold, Nielsen reports a missing value for quantities, which we impute as zero. In these cases, we do not observe price. We have chosen a sample of stores and products to minimize the likelihood of having zero sales, leaving on average 0.32 percent of store-weeks with no reported sales for each product.

Prices are not directly reported by Nielsen, but revenue is. We define price as revenue per unit sold at the store-week-UPC level. Average per unit revenue is likely close but not exactly equal to posted prices. There are two sources of discrepancy, as Einav et al. (2010) and DellaVigna and Gentzkow (2019) explain, and has been documented in other

---

<sup>5</sup>Our analysis of seasonality excludes 131 products for which our seasonal estimates have very high standard errors, as described below.

scanner datasets by Eichenbaum et al. (2014). First, some shoppers use loyalty cards to get a discount. Second, the Nielsen week is Sunday to Saturday, but some stores change prices midweek, so average per unit revenue is a quantity-weighted average of the two posted prices occurring during that seven-day period. As we do not observe prices for store-weeks with zero sales, we impute prices in those weeks as the median price of other stores in the same chain-state-week. This imputation is likely accurate given the near-uniform pricing within stores of the same chain (DellaVigna and Gentzkow (2019)). This imputation fills in prices for nearly all zero quantity weeks.<sup>6</sup>

In some of the analysis, we distinguish between “base prices” and promotional prices. Unlike other commonly used sources for price information (e.g., the microdata underlying the CPI), Nielsen does not identify promotional prices. To overcome this issue, we developed an algorithm to identify base and promotional prices.<sup>7</sup> The basic idea is that, given a store’s time series of observed prices, the base price should be common and high, whereas promotional prices should be infrequent and low. Thus, in each week the algorithm looks for the base price as the highest price in prior and subsequent weeks. To avoid sensitivity to measurement error in prices, we apply a filter to the time series of prices that resembles a mini-max program. We select the width of the window of the filter to ensure a low false negative rate for base price changes. Given a measure of a store’s base price for every week from the algorithm, we define promotion periods as ones when actual prices are at least 10 percent below the base price.

The algorithm yields sensible results in our data. Table B.1 in the appendix show

---

<sup>6</sup>In most specifications, we work with the log of quantity, which means the zero quantity weeks are dropped. We use the imputed prices in these weeks when estimating seasonality in prices and in specifications that work with the inverse hyperbolic sine of quantity instead of the logarithm.

<sup>7</sup>We believe this algorithm is an improvement over prior approaches, which often look at whether prices are below a researcher-specified threshold such as \$1 (e.g., Hendel and Nevo (2003, 2013)), or the modal value for the year (e.g., Hosken and Reiffen (2004b)). Such procedures may work well for a few stores or products, but they would be infeasible for hundreds of products and thousands of stores—and typically require less transparent ad hoc decisions. Our approach offers an approach to calculating base prices that can be implemented efficiently, and can accommodate base prices that change over time and at unknown frequencies for a given store. Hitsch et al. (2017) also report developing an algorithm for identifying base prices and promotion periods, and provide figures like Figure B.1 in the appendix. We took inspiration from their approach, but we differ from them in three respects. First, we impute missing prices as the median price of stores in the same chain-state-week based on the uniform pricing within them, before implementing the base price algorithm. Second, we provide validation of the algorithm in the form of a simulation study and a verification with actual promotion data. Third, our approach is simple, which allows for the efficient estimation of base prices for many stores.

that the median store-year has four base prices in the data, which means it changes base prices three times a year. Base prices therefore change relatively infrequently, consistent with other findings in the literature (e.g., Nakamura and Steinsson (2013)). Additionally, products are on promotion on average 22% of the time with an average discount being 24%.<sup>8</sup> Full details can be found in Appendix B.

## 2 Empirical approach to measuring seasonality

### 2.1 A benchmark model

We begin by considering a simple model of seasonal pricing, for a firm with some market power. The model guides our empirical analysis, and is therefore intended to be applicable to a wide variety of products. We deliberately abstract from some salient features of the retailer’s pricing decision (e.g., storability). We do this in the hope of attaining a simple model that still captures the key aspects of seasonal price changes for many product categories.

We assume the residual demand curve for a given product facing a single-product firm at time  $t$  takes the following form

$$\ln q_t(p_t) = \alpha_t + \eta_t \ln p_t.$$

The quantity  $\alpha_t - \alpha_{t'}$  measures seasonal shifts in the level of demand between  $t$  and  $t'$ , and  $\eta_t - \eta_{t'}$  measures seasonal shifts in the elasticity of demand. Given marginal cost  $c_t$ , the profit-maximizing price is

$$\ln p_t^* = \ln c_t + \ln \left( \frac{\eta_t}{\eta_t + 1} \right).$$

Seasonal fluctuations in optimal prices reflect changes in marginal cost  $c_t$  or demand elas-

---

<sup>8</sup>We checked the performance of the algorithm, using the Dominick’s scanner data which reports promotional activities. We show under conservative assumptions that the algorithm has less than a 10 percent rate of both Type I and Type II errors. We have also simulated many time series for prices which are roughly realistic with those found in our sample, but with known base prices and promotions. We verify that the algorithm accurately recovers base prices and promotions.

ticities  $\eta_t$ , but not variations in the level of demand at a fixed price,  $\alpha_t$ .

The change in prices between any two periods  $t$  and  $t'$  is

$$\Delta_{t,t'} \ln p^* = [\ln c_t - \ln c_{t'}] + \left[ \ln \left( \frac{\eta_t}{\eta_t + 1} \right) - \ln \left( \frac{\eta_{t'}}{\eta_{t'} + 1} \right) \right].$$

The change in prices can stem from changes in demand elasticities or changes in marginal costs (because of time-varying or non-constant marginal costs). Our analysis maintains the assumption that marginal costs are constant across seasons (and do not change with quantity). The assumption that marginal costs are constant in quantity and over the year is standard in the retail pricing literature (e.g., Guler et al. (2014); Bayot and Caminade (2015); Kwon et al. (2018); Haviv (2019)), and is consistent with evidence that wholesale prices and input prices do not systematically vary over the year (MacDonald, 2000; Chevalier et al., 2003; Hosken and Reiffen, 2004b; Anderson et al., 2017). Moreover, we view it as plausible in our context, because the majority of products we study are consumer packaged goods, which are mass produced and likely do not depend on seasonal inputs.<sup>9</sup> One concern is that, plausibly, the economic marginal cost of a good does vary over the year, as the opportunity cost of shelf space changes with aggregate shopping intensity. However, this cost is common to all products, and we will show that the timing of demand peaks varies from one product to the next. Therefore, it is difficult to argue that the opportunity cost of shelf space could jointly explain the seasonal fluctuations in quantities across all our products. The assumption that marginal costs do not depend on quantity is more ambiguous, but we would expect, if anything, that marginal costs increase with quantity. This would imply that prices rise with seasonal demand, however, and we typically observe the opposite.

The benchmark model implies that to understand seasonal pricing and how it relates to demand, we should estimate seasonal shifts and rotations of the demand curves, the  $\alpha_t$  and  $\eta_t$  terms, and relate the term  $\ln \left( \frac{\eta_t}{\eta_t + 1} \right)$  to the seasonal shifts in prices. This is exactly the approach we take in our empirical analysis, which we describe next.

---

<sup>9</sup>An important exception is fresh fruit and vegetables. We show below that this category also exhibits unusual seasonal patterns in our analysis below, and our results are stronger if we exclude it.

## 2.2 Seasonality in demand

To measure seasonality in demand, we estimate the empirical analog of our demand equation, separately for each product  $i$  in the data. Letting  $j$  index stores and  $t$  weeks, our approach to measuring seasonality in demand starts with the following regression:

$$\begin{aligned} \ln q_{ijt} = & \sum_{m=1}^{12} \alpha_{im} 1\{\text{Month}_t = m\} \\ & + \eta_i (\ln p_{ijt} - \overline{\ln p_{ijy}}) + \sum_{m=1}^{12} \eta_{im} 1\{\text{Month}_t = m\} (\ln p_{ijt} - \overline{\ln p_{ijy}}) \\ & + \gamma_i \text{unemployment}_{jt} + \theta_{ijy} + \varepsilon_{ijt}. \end{aligned} \quad (1)$$

This regression includes a set of month-of-year dummy variables, log price  $\ln p_{ijt}$  (de-meaned relative to the store-product-year average,  $\overline{\ln p_{ijy}}$ ), a set of interactions between month-of-year dummies and log price, and the local unemployment rate, as well as a store-year fixed effect. All coefficients are product specific. We estimate product-by-product and cluster standard errors at the chain level. Below we discuss the product-month specific elasticities and explain why we control for the local unemployment rate.

In estimating this equation, we use two normalizations: we set  $\sum_m \alpha_{im} = 0$  and  $\sum_m \eta_{im} = 0$ . These normalizations imply that the  $\alpha_{im}$  terms can be interpreted as the monthly deviation from store-year average demand, and the  $\eta_{im}$  terms as the monthly deviation from store-year average elasticities. Both normalizations are natural in the context of seasonality (e.g., Suits (1984)), and also facilitate our application of the Empirical Bayes procedure (discussed further below).

We interpret Equation 1 as the residual demand curve facing a given retailer, the empirical analog of the demand curve in our benchmark model. This interpretation requires that seasonal variability in quantities reflects changing demand and not changing costs. Costs can only affect quantities, however, through their effect on prices. As Equation 1 adjusts for prices, the remaining variation in quantities reflects demand rather than cost side factors.

Given the demand curve, a natural measure of the seasonality in demand for product

$i$  is

$$\text{Amplitude}_i = \max_m \alpha_{im} - \min_m \alpha_{im}.$$

This is the difference in log quantities between the month with maximal sales and the month with minimal sales, for a given store-year-product, holding prices fixed at their store-year-product mean. We can also use the estimates to obtain peak and trough demand months:

$$\bar{m}_i = \operatorname{argmax}_m \{ \alpha_{im} \} \tag{2}$$

$$\underline{m}_i = \operatorname{argmin}_m \{ \alpha_{im} \}. \tag{3}$$

These are the months when level of demand achieves its minimum or maximum.

This approach has several appealing features. First, like MacDonald (2000), we let the data indicate the timing of seasonal peaks and troughs, rather than impose a functional form or assume that seasonal fluctuations in demand are tied to temperature or holidays, as Chevalier et al. (2003) do. This data-driven approach is especially valuable when studying seasonality across a wide array of products, like we do, because it would be difficult a priori to specify the seasonal patterns for over a thousand products. A potential limitation of a data-driven approach, however, is that it risks conflating price-induced shifts along the demand curve with seasonality-induced shifts in demand. Our specification addresses this possibility by explicitly adjusting for prices shifts, with month-specific coefficients. Finally, because we estimate Equation 1 product-by-product, we obtain granular estimates of seasonality, which is especially valuable for measuring changes in prices that are free of compositional effects, as we explain below.

The estimation of Equation 1 must overcome two challenges: price endogeneity and estimation error. The price endogeneity challenge is the usual concern that differences in prices could themselves be driven by demand shocks that are unobserved to the econometrician. Although our specification is already tightly controlled—store-year fixed effects control for a great deal of demand heterogeneity—we address price endogeneity using an instrumental variables approach. We instrument for the price of a product in given store-week using the prices of that product in the same chain and week but in stores in other

markets, specifically all other DMAs. These are within-chain versions of the Hausman (1996) instruments.<sup>10</sup> The justification for this instrument is that for a given product and week, prices are highly correlated within a chain, and so a given chain’s out-of-market prices are informative for local prices, but are unlikely to be correlated with local demand conditions (DellaVigna and Gentzkow, 2019; Hitsch et al., 2017). Past work using the instrument has found it to be strong and to produce sensible demand elasticities (DellaVigna and Gentzkow, 2019; Allcott et al., 2019), a finding which we confirm in our sample as well.

This instrument corrects for two other potential endogeneity problems that relate to the measurement error in prices that might result from misaligned store-weeks and loyalty cards. The misalignment is that Nielsen weeks are defined as Saturday-Saturday, but some stores change prices midweek. If people buy more during the low price week, there is a mechanical price-quantity correlation. The loyalty card problem is that we do not observe discounts for using loyalty cards, and greater loyalty card use generates both greater quantities and lower prices, again creating a mechanical price-quantity correlation. Our instrument solves both these problems by using other stores’ prices, breaking the mechanical link between prices and quantities.

The second challenge our approach faces is estimation error, which can lead to bias in estimates of seasonal amplitudes. To see why, note that the natural plug-in estimator for the seasonal amplitude is:

$$\text{Amplitude}_i = \max_m \hat{\alpha}_{im} - \min_m \hat{\alpha}_{im},$$

where  $\hat{\alpha}_{mi}$  is the estimated value of the month-of-year effect. Even in the absence of any seasonal effects, so that  $\alpha_{mi} = 0$  for every month, estimation error will generally mean that some  $\alpha_{mi}$  are positive and some are negative, producing a positive estimate of seasonality. More generally, estimation error increases the expected difference between estimated  $\alpha$ ’s, biasing seasonal amplitudes upward.

To address this challenge, we use Empirical Bayes shrinkage. We briefly summarize

---

<sup>10</sup>Specifically, we instrument for the price-month interactions with a series of interactions between within chain out-of-market prices and month-of-year dummies.

our approach here, and provide more details in Appendix C. This process yields shrunk coefficients  $\tilde{\alpha}_{im}$ , defined as weighted averages of  $\hat{\alpha}_{im}$  and 0 (the value of  $\alpha_{im}$  under the null of no seasonality). The weights we use depend on the variability of the coefficients (across months) and the estimated variance of the estimates (i.e., the standard errors); the noisier are the individual estimates, the more weight is put on the null of zero. We then estimate seasonal amplitudes as

$$\text{Shrunk Amplitude}_i = \max_m \tilde{\alpha}_{im} - \min_m \tilde{\alpha}_{im}. \quad (4)$$

In the extreme case when all variability is due to estimation error, the shrinkage procedure shrinks all the estimates to their mean, and so we obtain an amplitude of zero. In general, this procedure tightens the range of seasonal coefficients, adjusting for the fact that selecting the maximum and minimum of a set of noisy estimates will generally result in too large a range. Throughout our empirical analysis, we focus on results based on the shrunk estimates. However, we also present as a robustness check the set of results which use the estimates before shrinkage as well. These results show more seasonality in quantities, more variable elasticities, and almost no change in the cyclicity of prices. Thus, the shrunk estimates are conservative and more precise.

**Using unemployment effects as a benchmark** Ideally, we would have a benchmark to assess whether observed seasonal fluctuations are large. We obtain such a benchmark by including as an additional regressor the local (county-specific) unemployment rate, a natural measure of the stage of the business cycle.<sup>11</sup> Using the local unemployment rate as a measure of the business cycle has been used many other settings including those examining scanner data (e.g., Coibion et al. (2015)). Including this regressor does not substantially affect any of our estimates, but the size of its effects provides a useful benchmark, letting us compare the size of fluctuations driven by seasonal shifts in demand to those induced by business cycle considerations.

**Abstracting from holiday effects** In estimation, we exclude all weeks that include a major holiday. The major holidays include New Years, the Superbowl, Memorial Day, July

---

<sup>11</sup>Given its role as a benchmark to seasonal fluctuations, we use the seasonally adjusted measure of the unemployment rate.

4, Labor Day, Thanksgiving, and Christmas. We abstract from them for several reasons. First, for certain holidays and products (such as ketchup and mustard during July 4, or Turkeys during Thanksgiving), we are concerned about the possibility of “loss leaders”—temporary deep sales on a particular product in an effort to pull in people to a retailer (e.g., Lal and Matutes (1994); Chevalier et al. (2003); DeGraba (2006)). Such effects would be difficult to detect empirically (because they would rely on cross-category demand elasticities), but they are likely not representative of seasonal pricing, because (as we show) many products exhibit seasonal peaks in demand, and it is unlikely that every product could be focal enough to act as a loss leader. A second reason to abstract from holidays is that fairness constraints may make it difficult for retailers to suddenly and sharply raise prices (Kahneman et al., 1986). We expect that such constraints do not, however, prevent retailers from seasonally adjusting their pricing strategies over the course of the year. Thus, the price response to holidays need not be representative of the general pricing patterns over the year. A final reason to abstract from holiday effects is that it is difficult to identify holiday-specific demand elasticities for every product. This is because our empirical strategy uses within product-chain-year price variation, essentially looking at chain-wide discounts; not all products are discounted on holidays, making it impossible to identify holiday shifts in price elasticities.

### 2.3 Seasonality in prices

We are interested in how seasonal price fluctuations compare to seasonal fluctuations in demand. In particular, we want to measure the sign and size of price changes between peak and trough demand. To do so, we estimate

$$\ln p_{ijt} = \sum_{m=1}^{12} \beta_{im} 1\{\text{Month}_t = m\} + \tau_i \text{unemployment}_{jt} + \mu_{ijy} + \epsilon_{ijt}. \quad (5)$$

This is a regression of log price on a set of month-of-year dummies (with  $\sum_m \beta_{im}$  normalized to zero), plus a store-year fixed effect. Again, we estimate product-by-product and cluster standard errors at the chain level. As with our quantity estimates, we do

not work directly with the estimated month-of-year effects, but instead work with their shrunk values,  $\tilde{\beta}_{im}$ , to avoid estimation error leading us to overstate the extent of seasonal fluctuations.

We use Equation 5 to assess the cyclical in prices, and specifically, how it compares to demand. Therefore, we use the peak and trough months of demand, and compare the prices observed during these periods. In particular, we calculate

$$\Delta p_i \equiv \tilde{\beta}_{i\bar{m}_i} - \tilde{\beta}_{i\underline{m}_i}. \quad (6)$$

This is the difference in (shrunk) month-of-year coefficients from the price equation, but comparing the months when *demand* peaks and troughs (see Equations 2 and 3). For example, if we estimate for product  $i$  that price-adjusted quantities peak in December and trough in July, then we would take the difference  $\beta_{12i} - \beta_{7i}$ . If prices are countercyclical, we would expect this difference to be negative. A second use of these estimates is to calculate the peak and trough price months, as we do with quantities. Countercyclical pricing would imply that a given product's peak demand month is its trough price month.

We emphasize that we implement this approach product-by-product, and obtain product-specific measures of pricing patterns. We therefore avoid the potential problem of confounding the composition effect highlighted by Nevo and Hatzitaskos (2005). The composition effect could arise if during peak demand periods, people substitute towards low-price varieties. Here, we hold fixed the product, and so avoid compositional shifts affecting our measures of countercyclical pricing.

## 2.4 Seasonality in elasticities

Our benchmark model implies that how optimal prices fluctuate seasonally depends crucially on if there are seasonal changes in elasticity of demand. Therefore, we also investigate how demand elasticities vary over the year, and in particular how they might vary between peak and trough demand. Our baseline specification, Equation 1, yields an estimate of a product-level elasticity  $\hat{\eta}_i$ , as well as product-month specific deviations,  $\hat{\eta}_{im}$ . As with our other estimates, we focus on the shrunk values of the monthly deviation,  $\tilde{\eta}_{im}$ ,

to avoid spuriously inflating seasonal fluctuations.

Given these shrunk values, we can find the seasonal difference in elasticities as

$$\Delta\eta_i = \tilde{\eta}_{i\bar{m}_i} - \tilde{\eta}_{i\underline{m}_i}, \quad (7)$$

or the difference in elasticities between peak and trough demand months (with peak and troughs defined by price-adjusted demand levels).  $\Delta\eta_i$  is interesting because it is informative about the optimality of seasonal pricing.

The change in elasticities is informative about the sign of optimal price changes, but not directly their magnitude. We therefore also consider another object, the change between trough and peak demand months in the optimal price implied by our benchmark model:

$$\Delta p_i^* = \ln \frac{\hat{\eta}_i + \tilde{\eta}_{i\bar{m}_i}}{\hat{\eta}_i + \tilde{\eta}_{i\bar{m}_i} + 1} - \ln \frac{\hat{\eta}_i + \tilde{\eta}_{i\underline{m}_i}}{\hat{\eta}_i + \tilde{\eta}_{i\underline{m}_i} + 1}, \quad (8)$$

$\Delta p_i^*$  is the (log) change in optimal price for a single-product monopolist experiencing an all-else-equal elasticity change from  $\hat{\eta}_i + \tilde{\eta}_{i\bar{m}_i}$  to  $\hat{\eta}_i + \tilde{\eta}_{i\underline{m}_i}$ . It is therefore a useful benchmark not only for the sign but also for how large a price change we should expect to see for product  $i$ . Of course, richer models would imply that other factors (such as cross-price, or other products' own price elasticities) also influence the optimal price. We view this as a useful benchmark, however, rather than a definitive pricing rule.

## 2.5 Identification and robustness

We interpret  $\alpha_{im}$  and  $\eta_{im}$  as measuring the level and elasticity of the residual demand curve facing a retailer in a given month. This interpretation is justified under the assumption that prices in other stores (in a given chain and conditional on product-store-year and product-month fixed effects) are uncorrelated with unobserved shocks to demand. This assumption is plausible because price variation over time within a chain-product largely reflects the promotional pricing strategies that are pre-negotiated between manufacturers and retailers, and not temporary changes in local demand conditions.

Thus, we interpret our estimates as measuring the level and slope of the residual (log)

demand curve facing a given retailer. A limitation of our estimates is that we assume a single elasticity (at a point in time). In practice, elasticities may differ for short- and long-lasting price changes (for example because of storability) and they may differ for single-product price changes relative to category-wide price changes. Assuming away these rich complications lets us estimate elasticities month-by-month for over 1,000 products.

However, a reasonable concern is that the confounding influences of other prices—either past prices or competitor’s current price—might affect our estimates. There are three potential concerns. First, we estimate Equation 1 product by product and do not account for other products’ prices. To the extent that price changes are coordinated, we may underestimate the own-price elasticity. We address this concern in two robustness specifications that control for the average of other (log) prices of products in the same module, and the minimum log price of other products in the same category. (DellaVigna and Gentzkow (2019) and Strulov-Shlain (2019) also use this strategy to address this concern.) A second concern is that we do not adjust for the price that competing retailers are charging. To address this concern, in a robustness test, we include a control for the average log price of other chains in the same market (DMA).<sup>12</sup> Finally, a third concern is that some products are storable, in which case past prices influence current demand. We address this concern in a final robustness test by including a control for “weeks since last promotion,” following Hendel and Nevo (2003) and Strulov-Shlain (2019).

### 3 Seasonality in demand, prices, and elasticities

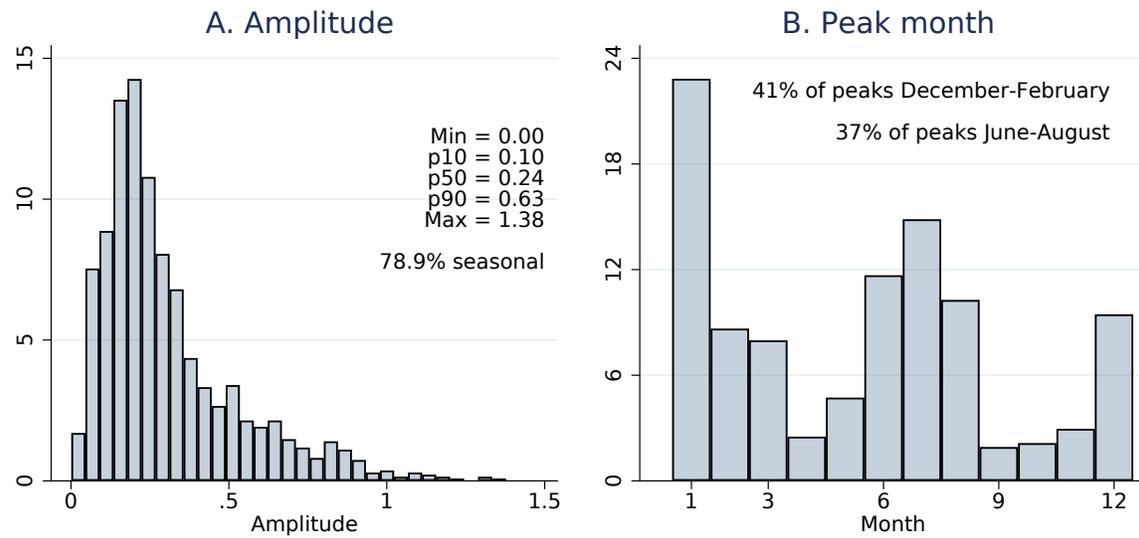
#### 3.1 Seasonality in demand

We begin by describing seasonal fluctuations in demand. For these and all subsequent results, we exclude products with extremely large standard errors. For these products, it is impossible for us to reliably measure seasonality. We exclude products if, for any month, the standard error of its demand level ( $\alpha_{mi}$ ), demand elasticity ( $\eta_{mi}$ ), or price level ( $\beta_{mi}$ ) is larger than two. Applying this filter removes 178 products, bringing our

---

<sup>12</sup>We take a store-weighted average price, to avoid compositional biases.

Figure 1: Seasonality in demand



Notes: Panel A plots the distribution of estimated amplitudes, across products, defined as the difference in shrunk log quantity (adjusted for price) between the peak and trough (of demand) month. Panel B plots the distribution of peak demand months.

final sample down to 1,353. Inspection of the excluded products reveal they are usually products sold in relatively few chains, and with limited price variation.<sup>13</sup>

We present the distribution of peak-to-trough amplitudes for price-adjusted log quantities in Panel A of Figure 1. This figure shows the magnitude of the seasonal fluctuations in quantities for the final 1,353 products across 38 modules, and therefore gives a comprehensive sense of the size and scope of seasonal fluctuations of demand. Seasonal fluctuations in demand are large for most products. For example, the 10th percentile is 10 log points, the median is 24 log points, and the most seasonal products have fluctuations above 63 log points. The breadth of the seasonal fluctuations in demand that we are findings are further underscored by the design of our product selection. Unlike past work, our sample selection was not pre-disposed towards seasonal or non-seasonal prod-

<sup>13</sup>In particular, we have determined that high standard errors arise for two reasons. First some products are widely available but only in drug or mass merchandise stores. This results in high standard errors because there are few drug and mass merchandise chains in the Nielsen data, and we cluster our standard errors on chain. Second, some products have limited within chain-year price variation, making it impossible to precisely estimate time varying elasticities.

ucts, and so we view these results as providing the first comprehensive assessment of the overall size of seasonal demand fluctuations facing retailers.

These seasonal fluctuations are large relative to other cyclical variations, especially variations attributable to business cycles. We estimate, for example, that movements in the local unemployment rate has a relatively small effect on the sales of most products. The median effect (across products) of a 1-percentage point increase in unemployment on quantities is -2 log points.<sup>14</sup> As unemployment rates rose by about 4.5 percentage points during the Great Recession in the average county, our estimates imply the median product's demand declined by about 9 log points during the Great Recession—small relative to the typical seasonal fluctuation. As a benchmark, if a product's seasonal (price adjusted) quantity amplitude change is greater than 4.5 times its unemployment effect, we define the product as having seasonal demand. By this measure, about three quarters of the products in our sample are seasonal. The prevalence of seasonal demand products, as measured by this benchmark, partly reflects the fact that grocery shopping in general is sensitive to seasonal factors (MacDonald, 2000; Chevalier et al., 2003), and less sensitive to business cycles (Coibion et al., 2015). Nonetheless, our results indicate substantial and heterogeneous seasonality in demand.

Significant seasonal fluctuations in demand is not limited to select set of product categories. We report several dimensions of the seasonality in demand by product module (Nielsen's category measure) in Appendix Table A.2. The product categories with the most seasonal demand are canned soup and frozen novelties, as might be expected given their obvious connection to the outside temperature. Even a not-so-obvious seasonal category such as cookies, however, exhibits fairly high seasonality: 80 percent of products within the category are seasonal, with an average amplitude of 28 log points. Of course, cookies are much less seasonal than canned soup. It is not the case, however, that our method simply finds large seasonal amplitudes for all products. For example, we find that dog food, cat food, and pain remedies all have very modest seasonal amplitudes in demand—a feature one might have expected.<sup>15</sup>

---

<sup>14</sup>We report in Appendix Figure A.1 the distribution of unemployment coefficients.

<sup>15</sup>These categories also have little business cycle variation, so some of their products do still have larger seasonal components than business cycle components, and are thus deemed seasonal.

The heterogeneity in seasonal demand appears not only in the size of fluctuations but also in their timing. We measure the timing of seasonality by the month in which demand peaks, and we plot the distribution of peak demand months in Panel B of Figure 1. The modal peak demand month is January, which is broadly consistent with the findings in MacDonald (2000), although our estimates control for price changes. There is a clear second mode in the summer, however, when demand peaks for seasonal products like ice cream, hot dogs, beer, and juice and water.<sup>16</sup>

The heterogeneity in the timing of demand peaks is important for the rest of our analysis for at least two reasons. First, the variation in demand peaks across products indicates it is unlikely that any variation in marginal cost that is common to all products (e.g., cost of shelf space) could be the primary driver for the quantity fluctuations we are observing. This supports the view that our empirical approach is picking up seasonal fluctuations in demand, and not movements along a demand curve. Second, the variation in demand peaks indicates that it is also unlikely that a common (or aggregate) shock to demand for all retail products (e.g., increased shopping intensity around Christmas) can explain the overall size and breadth of seasonal demand fluctuations across products.<sup>17</sup>

Thus, the results reveal that seasonal fluctuations in demand are large, common, and heterogeneous in magnitude and timing. This finding is important because it indicates that seasonal fluctuations in demand are a central part of the retail sector, and a likely economically important dimension in the effectiveness of their pricing decisions. The results also reveal the importance of a broad examination; studies focusing on individual product categories, even obviously seasonal ones, are likely to miss much of the seasonality in demand facing retailers.

## 3.2 Countercyclical prices

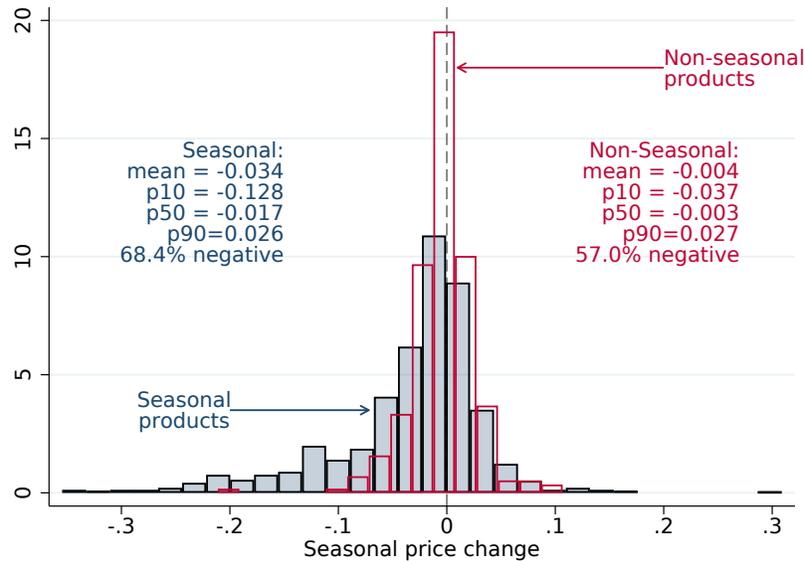
How do movements in prices relate to these large seasonal fluctuations in quantities? We measure the seasonality in prices by looking at how prices change between the peak

---

<sup>16</sup>We report, for each module, the mode peak month (when it is unique) in Appendix Table A.2.

<sup>17</sup>To this point, it is also important to note that our empirical estimates of the seasonal fluctuations in demand abstract from holiday weeks.

Figure 2: Distribution of seasonal price changes



Notes: This figure plots the distribution of observed price changes between peak and trough demand months, separately for seasonal and non-seasonal products. Seasonal products are those with a peak-trough difference in demand greater than the implied effect of a 4.5 percentage point increase in unemployment.

and trough of demand. We plot the distribution of these price changes in Figure 2, looking separately at seasonal and non-seasonal products. For the typical product—seasonal or not—prices move very little at seasonal frequencies, less than 2 log points. For the 1,067 seasonal products, between peak and trough demand months the average price change is negative and small, about -3 log points. This negative price change indicates that for the average seasonal product, prices are countercyclical, in that they fall when demand peaks. Indeed, prices are countercyclical for a majority of seasonal products, about two-thirds. There is, however, a heavy left-tail of products with larger countercyclical price movements. At the 10th percentile of seasonal products, for example, prices decline by 13 log points from trough to peak demand months. Even this, however, is a relatively small price change when judged against the benchmark of quantity changes, where the median change is 23 log points. Countercyclical pricing, limited though it may be, is a common property of products with seasonal demand. This conclusion is further strengthened by the tight and symmetric distribution of seasonal price changes around zero for the prod-

ucts without seasonal demand.

Prior research has found that countercyclical movements in the average price paid largely reflect changing composition of the set of products purchased (Nevo and Hatzitaskos, 2005; Perrone, 2016). Our estimates of price changes are free from this problem because we look at price changes on a UPC-by-UPC basis. This finding does not mean, however, that compositional changes are unimportant. To illustrate the importance of compositional changes, we re-estimate our demand and price equations, but using data aggregated to the category (i.e., module) level, and look at cyclicity in the sales-weighted average price. We find that category-level prices are much more countercyclical than are product-level prices. For the average module, category-level prices fall by 2.5 log points between peak and troughs of demand, whereas the fall in product-level prices is only 1.4 log points. Looking across modules, Appendix Figure A.3 shows the scatter plot of category-level price cyclicity against product-level cyclicity (averaging up to a given product category). Although there is a strong association, points are typically below the 45-degree, especially for products with the most countercyclical prices, reflecting that category-level prices are more countercyclical than product-level prices (see, e.g., Nevo and Hatzitaskos (2005)).

Thus, prices often fall when demand peaks, although rarely by a large amount. This is one indication that prices are countercyclical. An alternative indication is: how often are prices at their trough when demand peaks? We answer that question in Figure 3. In this figure, we divide up the seasonal products according to the month when demand peaks. Then, for each demand peak month, we plot the distribution of what month prices for those products reached their trough. For example, the top left panel of the figure is limited to the 257 products in which demand peaks in January. For those products, we plot the distribution of the price trough month. If prices are strongly countercyclical, we should see that they tend to trough in exactly the month when demand peaks. That is the pattern we find, with some exceptions. For products with a demand peak in a particular month, we find that there is a clearly elevated probability of having a price trough in that month. To show this more clearly, in each panel, we also plot the distribution of trough price months for products with peak quantities in all other months. To quantify the patterns

Figure 3: Timing of price troughs, given timing of demand peaks



Notes: This figure plots the distribution of observed trough price months, among seasonal products, separately by the timing of peak demand months (e.g., in the top left, we plot the distribution of price trough months among products whose demand peaks in January). The hollow bars plot the distribution of timings for other products (e.g., those that do not peak in January). Seasonal products are those with a peak-trough difference in demand greater than the implied effect of a 4.5 percentage point increase in unemployment.

in the figures, we make a product-month data set, and regress an indicator for “trough price month” on an indicator for “peak demand month”, and obtain a coefficient of 0.10 (with a standard error of 0.01, clustering on product), meaning there is a 10 percentage point increase in the likelihood that a given month is the trough price month if it is also the peak demand month. Since the base rate is 1/12, this is about a 100 percent increase.

The observed seasonal price changes could be because of falling base prices, deeper sales, or greater frequency of promotional prices. We decompose the sources of countercyclical pricing in Table 2. The table reports the distribution of changes in log price, log base price, frequency of promotional price discounts, and depth of promotional price discounts, between trough and peak demand months. On average base price changes account for about half of the countercyclical price change, although they account for less than a quarter of the median product’s price change. The balance of countercyclical pricing is accounted for by an increased frequency of promotional prices. The average seasonal product has promotional prices 4.6 percentage points more often during peak quantity periods, relative to trough periods, and for the median product the difference is 2.9 percentage points. The depth of promotional price discounts varies little at seasonal

Table 2: Sources of seasonal price changes

Seasonal change in...	Mean	10 <sup>th</sup> percentile	Median	90 <sup>th</sup> percentile
Price	-0.034	-0.128	-0.017	0.026
Base price	-0.016	-0.067	-0.004	0.020
Has promotional price	0.046	-0.040	0.029	0.157
Depth of sale	-0.004	-0.044	-0.000	0.033

Notes: This table reports statistics on the distribution (across products) of changes in the indicated price component, between peak and trough demand months, for seasonal products. Seasonal products are those with a peak-trough difference in demand greater than the implied effect of a 4.5 percentage point increase in unemployment.

frequencies.

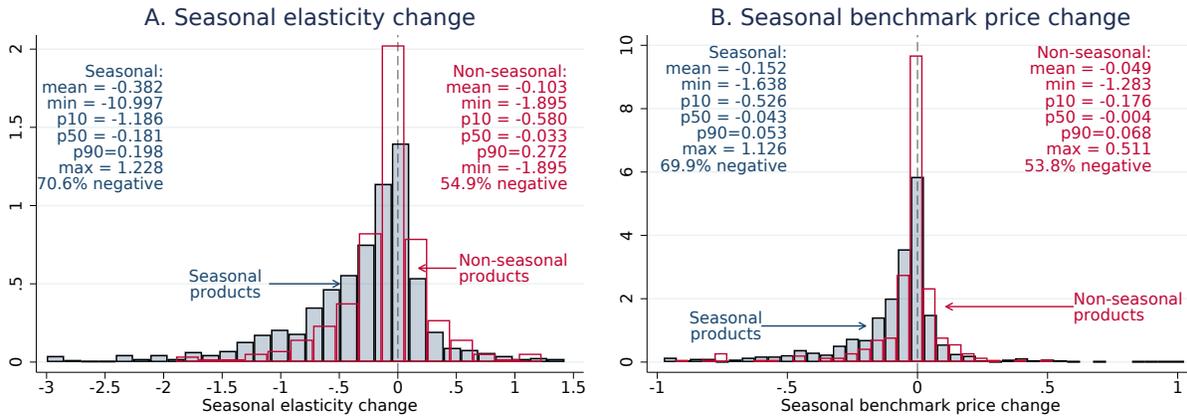
Overall, we find that when demand peaks, prices typically fall, in equal parts because of more frequent pricing promotions and lower base prices. Furthermore, prices are especially likely to reach their trough in months of peak demand.<sup>18</sup> However, the price changes are fairly small for the typical product, and there are some seasonal products that do not show countercyclical pricing. Thus, case studies of individual product categories could easily reveal pro- or countercyclical pricing, and any assessment of whether prices are countercyclical requires a broad perspective.

### 3.3 Demand elasticities are often countercyclical

The results so far show that many products exhibit large seasonal fluctuations in demand, with little corresponding price changes. What price changes we do observe are often countercyclical, in the sense that prices often fall when demand peaks. This finding is counterintuitive when viewed through a perfectly competitive, simple (upward sloping) supply and demand framework, which would imply that prices should remain steady or rise with demand. However, a perfectly competitive model likely does not adequately describe the retail sector, as retailers likely have considerable market power. For a firm with market power, the optimal prices depends crucially on the price elasticity of

<sup>18</sup>We emphasize that our measure of peak month of demand is price adjusted, so it is not caused by the price changes we describe here.

Figure 4: Distribution of seasonal elasticity and benchmark price changes



Notes: Panel A of the figure plots the distribution (across products) of elasticity changes between peak and trough demand months, separately for seasonal and non-seasonal products. Seasonal products are those with a peak-trough difference in demand greater than the implied effect of a 4.5 percentage point increase in unemployment. Panel B plots the distribution of the implied change in the benchmark price, given by Equation 8. For visual clarity, the histograms plot truncated distributions, but we report statistics from the untruncated distributions.

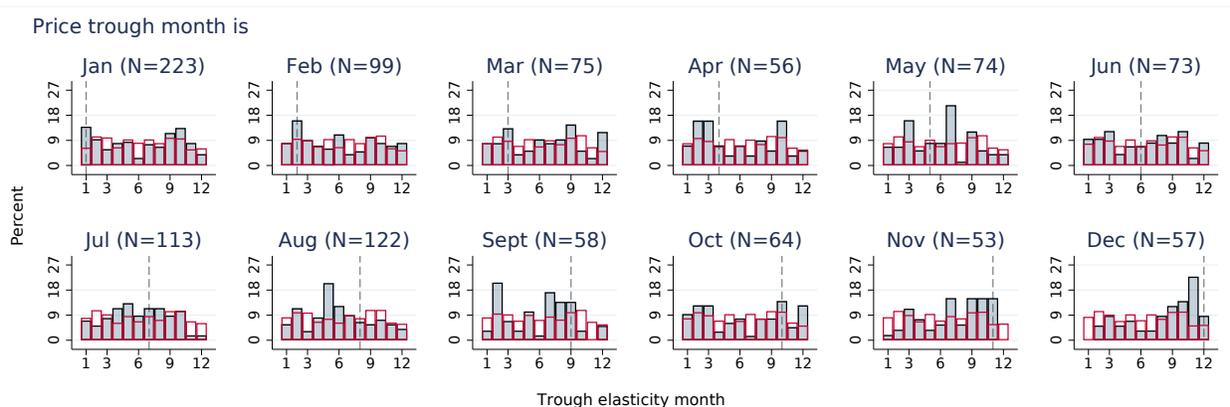
demand. We therefore investigate how these demand elasticities vary with the level of demand across the seasons.

Our baseline specification, Equation 1, yields demand elasticities that vary by month-of-year.<sup>19</sup> We calculate how these (shrunk) estimates change between trough and peak quantity months, and we report the distribution of changes in elasticities in Panel A of Figure 4. We look separately at seasonal and non-seasonal products. For seasonal products, demand elasticities are typically countercyclical, meaning that demand usually becomes more elastic when demand peaks. About 70 percent of seasonal products exhibit countercyclical elasticities, and for the average seasonal product, the elasticity falls by 0.4 between peak and trough demand months. For non-seasonal products, the seasonal changes in elasticities are smaller and nearly symmetric around zero.

These estimated demand elasticities broadly rationalize the observed price patterns of the previous section. Demand becomes more price sensitive during demand peaks,

<sup>19</sup>We plot the distribution of estimated elasticities (across products; not the monthly deviations) in Appendix Figure A.2. The elasticities appear sensible. The mean and median are -2.64 and -2.51. Every one of our 1,353 products has a negative elasticity, and all but 8 are elastic (i.e. less than -1).

Figure 5: Timing of elasticity troughs, given trough price month



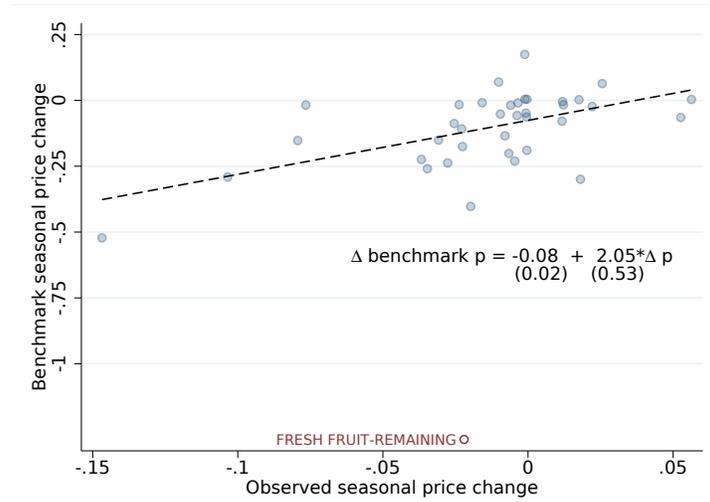
Notes: This figure plots the distribution of observed trough elasticity months, among seasonal products, separately by the timing of trough price months (e.g., in the top left, we plot the distribution of elasticity trough months among products whose prices trough peak in January). The hollow bars plot the distribution of timings for other products (e.g., those that do not trough in January). Seasonal products are those with a peak-trough difference in demand greater than the implied effect of a 4.5 percentage point increase in unemployment.

providing an incentive for retailers to decrease their prices. About 70 of products show countercyclical prices, and about 70 percent show countercyclical elasticities. However, the results so far do not necessarily indicate that the observed magnitude of price changes is consistent with elasticity changes, because (in a our benchmark model) the optimal price is a nonlinear function of the elasticity. We therefore report in Panel B of Figure 4 the distribution of implied changes in the benchmark price, defined as  $\ln \frac{\eta}{\eta+1}$ , where  $\eta$  is the shrunk elasticity estimate.<sup>20</sup> For seasonal products, the mean change in benchmark price is -0.15 and the median change is -0.04. Although the mean change in  $\ln p^*$  is larger than the observed change in prices, the median is close. For most products, the change in the benchmark price is both small and not too different from what is observed.

We present two further pieces of evidence that low and typically countercyclical price changes are consistent with optimal pricing. First, we show that the elasticities are particularly likely to reach their trough when prices do as well, so the timing of changes

<sup>20</sup>This is undefined for products with elasticities greater than -1. Following DellaVigna and Gentzkow (2019), we use winsorized elasticities to calculate optimal prices, winsorizing at -7 and -1.2. This winsorizing explains why there are fewer negative benchmark price changes than elasticity changes (because it sometimes results in exactly zero change).

Figure 6: Benchmark and observed price changes, category-level



Notes: This figure plots the (category-level average) benchmark price change against (category-level average) observed price change, between trough and peak demand months. The best fit line (black dashed line), estimated at the category-level, excludes the fresh fruit category.

in prices and elasticities line up. We show in Figure 5 the distribution of trough (most elastic) elasticity months, among products with a given trough price month, analogous to our exercise for the timing of prices and demand levels. The figure reveals more noise in the timing of elasticity changes than what we observed for prices. Nonetheless, there is a clear pattern for elasticities to trough when prices trough, and indeed we find in a regression framework that the trough price month is 4.1 percentage points more likely to be the trough elasticity month (standard error, clustered on product, of 0.01), a 50 percent increase over the base rate.

Our second piece of evidence that price changes are consistent with optimal pricing comes from comparing the magnitude of price changes across modules to the implied benchmark price change. We aggregate to the module level because the noisiness in the product-level elasticity estimates makes it difficult to assess their correlation with product-level price changes. To aggregate, we take the simple average, across products in a module, of the estimated seasonal price change and benchmark price change. We plot the benchmark price change against the observed price change in Figure 6. Our

benchmark models implies that these price changes should be perfectly correlated with a 1:1 relationship. Measurement error (especially in the elasticity estimates) will attenuate the correlation. Nonetheless, the figure shows a clear connection between observed price changes and benchmark price changes, with one important exception: fresh fruit-remaining. Our simple benchmark model assumes that demand fluctuations are the only reason for price changes, but this assumption is likely to be a poor one for fresh fruit, where harvest conditions are likely to be just as seasonal (or more) as demand. Excluding fresh-fruit, we estimate a strong relationship between observed price changes and benchmark seasonal price changes, with a slope coefficient of 2.05. This magnitude means that benchmark prices are strongly correlated with seasonal prices, and on average move more than one for one. While this might indicate that seasonal prices do not move enough, relative to an optimal benchmark, it is only marginally significantly different from 1 ( $p$ -value= 0.06), i.e., meaning we cannot necessarily reject the hypothesis that prices move one for one with those implied by our benchmark. We view this evidence as showing that countercyclical elasticities can roughly quantitatively rationalize countercyclical prices. If anything, observed prices are not as countercyclical as our elasticity estimates say they should be.

### **3.4 Robustness**

Our approach to measuring seasonality assumes a particular form for demand. Our basic findings—high seasonality in quantities, limited and often countercyclical price variation, and countercyclical elasticities that align with observed price changes—are robust to alternative approaches. We consider several extensions and modifications of our basic model. To allow for competitive effects, we consider specifications that allow demand to depend on rival retailers' prices, the average price of other products in the same category, or the minimum price of other products in the same category. To assure that storability considerations are not driving the results, we estimate models in which demand also depends on time since the last sale. Finally, although our base specifications use instrumental variables to adjust for price changes, we also consider specifications that

use OLS, or forgo the price adjustment entirely.<sup>21</sup>

We report the distribution of demand amplitudes for each model in Appendix Table A.3. Allowing for competitive effects has essentially no effect on our estimates, and neither does allowing for storability. Alternative functional forms—using inverse hyperbolic sine rather than logs—raises our estimated seasonal amplitudes, although it does not change the seasonal share. Using OLS instead of IV to estimate the price parameters also makes very little difference. The price adjustments themselves are important, however. When we do not adjust for price, we find more seasonal products and larger seasonal amplitudes, because prices are countercyclical, and therefore amplify seasonal demand shifts. We consider the same set of robustness checks for our measure of seasonal price change. These measures are potentially sensitive to the demand specification because they depend on the timing of demand peaks and troughs. We present distributions of seasonal price changes from these alternative specifications in Appendix Table A.4. Across specifications, we find very similar results: the typical seasonal product exhibits countercyclical pricing, but many seasonal products do not, and the seasonal fluctuation in prices is typically small relative to seasonal fluctuations in demand.

We also examine the robustness of our elasticity and benchmark price estimates, again using the same set of alternative demand specifications. Adjusting for competitor's price has little impact on the estimated seasonality in elasticities or the share of products in which elasticities fall from trough to peak quantity months. Likewise, adjusting for the time since last sale makes little difference.<sup>22</sup>

Finally, we show the robustness of our observed price-benchmark price association to alternative demand systems in Table A.6.<sup>23</sup> Across our alternative specifications, we find

---

<sup>21</sup>We continue to exclude products with standard errors greater than 2. As the standard errors vary from specification to specification, the sample of products in each specification varies modestly as well.

<sup>22</sup>This might seem surprising, as the previous literature has found that the level of elasticity estimates for storable products can be sensitive to the handling of sales (e.g., Hendel and Nevo (2013); Haviv (2019)). There are two explanations. First, we find that the *difference* in elasticities between seasonal peaks and troughs is robust to controlling for time-since-sale. Second, many of our products are not storable and some are perishable. We show in Appendix Table A.7 that quantities are much more sensitive to time-since-sale for storable products than for perishable ones, confirming the validity of time-since-sale as a proxy for storage incentives.

<sup>23</sup>We do not consider robustness of the elasticity or benchmark price results to the inverse hyperbolic sine specification, because this specification yields an elasticity which depends on quantity and, hence the price level itself.

a large slope coefficient around 2, which is always statistically significantly different from zero and sometimes significantly different from one.

Finally, we show how our estimates change when we do not use the Empirical Bayes shrinkage procedure. Without shrinkage, we classify more products as seasonal, we find larger amplitudes and more varied elasticity changes, consistent with expectations that shrinkage reduces bias and noise in our estimates. The effect of shrinkage on the price estimates are smaller because these estimates are more precisely estimated.

## 4 Are seasonal price changes optimal?

We have shown that across seasons prices vary little, and for products with seasonal demand the price changes are typically countercyclical. These countercyclical price changes are qualitatively consistent with elasticity fluctuations, but observed prices vary less than optimal prices from a benchmark model. How meaningful are these deviations from the benchmark model? We answer this question by looking at how much profit retailers appear to be giving up by pricing at the observed level rather than our benchmark optimum. Using the parsimonious framework of Section 2, lets us calculate lost profits for our large number of products at the individual product level, providing a comprehensive assessment of the overall incentive to adjust prices seasonally, and for countercyclical pricing.

Our approach, which follows DellaVigna and Gentzkow (2019) closely and is described in more detail in Appendix Section D, proceeds in two steps. First, we back out an estimate of marginal cost at the product-store-year level assuming that the average price over the year is set optimally. Second, we calculate profits with optimally time-varying prices, for the observed price path, and also for constant prices. Profits under constant prices provide a compelling benchmark because we can ask whether observed countercyclical prices increase profit relative to an inactive pricing strategy. This gives a product-specific estimate of lost profits (relative to the benchmark optimal price). We aggregate across all products in our sample to obtain an in-sample estimate of total lost profit. We extrapolate to all products by assuming that lost profits are a constant share of revenue. Additionally, to highlight the scope that time varying elasticities can rationalize counter-

Table 3: How much profit could retailers be losing relative to benchmark prices?

	(1) Constant Price		(2) Observed Price		(3) Peak-Trough Gain	
Panel A: Median Chain Year–Sample Products						
Baseline	\$264,784	(0.005)	\$262,617	(0.005)	\$973	(0.001)
DMA competition	\$263,754	(0.005)	\$265,695	(0.005)	\$855	(0.001)
Other product competition - avg	\$243,596	(0.005)	\$245,591	(0.005)	\$11	(0.000)
Other product competition - min	\$261,221	(0.005)	\$260,496	(0.005)	\$964	(0.001)
Panel B: Median Chain Year–All Products						
Baseline	\$3,287,698	(0.005)	\$3,304,234	(0.005)	.	.
DMA competition	\$3,364,634	(0.005)	\$3,392,337	(0.005)	.	.
Other product competition - avg	\$3,025,780	(0.005)	\$3,100,350	(0.005)	.	.
Other product competition - min	\$3,250,099	(0.005)	\$3,282,576	(0.005)	.	.

Notes: This table reports the results of the lost profit calculations described in Appendix Section D, in annual dollars and as a share of revenue in parentheses, for the median chain-year under constant prices and observed prices in our sample of food stores. Additionally, profit gains from observed prices relative to constant prices exclusively for peak and trough months of demand is also reported. In the top panel, the lost profits are reported for our set of sample products. In the bottom panel, the revenue shares of our sample products are used to construct an estimate of the aggregate lost profits across all the products within a chain. Within each panel, the lost profits are reported using our baseline (shrunk) estimates of the price elasticity of demand as well as each of our alternative specifications. See notes to Table A.3 for an explanation of the different specifications.

cyclical pricing, we also report the profits gained exclusively between peak and trough months between observed prices and uniform prices across those two months. Each of these steps are repeated using the seasonal fluctuations in elasticities across several of our demand specifications.<sup>24</sup>

Table 3 reports the lost profits from constant (column 1) and observed (column 2) pricing, relative to our benchmark of optimal pricing, for the median chain-year. First, aggregating across all the products in our sample (Panel A), the estimated lost profits for the median chain-year are just about \$250,000, or about 0.5 percent of revenue. Scaling up to the aggregate chain level (Panel B), the median chain loses about \$3 million. Additionally,

<sup>24</sup>In what follows, we ignore the issue of pricing around holidays, as our time varying elasticity estimates do not pertain to holiday periods. Furthermore, in this section we limit the analysis to food stores to provide a more consistent means of extrapolating the amount of lost profits for our product modules and their share of yearly revenue for the chains in our sample.

these lost profit estimates are robust to the multi-product and competitive considerations we explored in our alternative demand specifications.

The lost profits arising from both the observed seasonal price paths of retailers as well as a perfectly constant price path are remarkably similar. This partly reflects the fact that broadly speaking elasticities vary little over the year. To shed further light on how this relates to the pattern of countercyclical pricing, we also calculate the profits (gained) of the observed relative to constant path, specifically over the peak and trough months of demand (column 3). Focusing exclusively on the pricing behavior between peak and trough demand months, the median retailer increases profits by about 0.1 percent of revenue with the observed price changes relative to uniform pricing. This calculation of course confirms the product category level results discussed in the previous section.

The overall magnitude of lost profits for both the observed and constant is small in multiple senses. First, it is small relative to revenue, which totals about \$1 billion per year for the typical chain in our data. Second, it works out to about than \$190 per product (i.e.  $\$250,000/1321$ ), meaning any product-level costs of reoptimizing would likely swamp the gains. Even this per product number is likely an overestimate given that our sample of products focuses on widely available products. Finally, it is small relative to lost profit from other forms of apparent mispricing. For example, DellaVigna and Gentzkow (2019) argue that the median chain loses about \$16 million from charging uniform prices across stores. On the other hand, retail profit margins tend to be quite small, because of large fixed costs. On a profit margin of just a few percent, an increase of 0.5 percent of revenue could be large. Overall, although observed prices do not perfectly coincide with optimal prices in our simple benchmark model, for the median chain the returns to addressing this issue are likely low, and are not in conflict with the countercyclical pricing we observe.

While admittedly stylized, the lost profit estimates are generally supportive of approximate profit maximization in the context of seasonal demand fluctuations—and that countercyclical price elasticities can partially rationalize the observed countercyclical pricing patterns in retail. The approach used to evaluate this claim is in contrast with the standard approach, especially in industrial organization, which typically assumes profit maximization and works out its implications (e.g., implied marginal costs or merger effects).

Interestingly, our results indicate only modest departures from this common assumption, which has been questioned with the emerging literature that has documented departures from it (Hortaçsu and Puller (2008); Goldfarb and Xiao (2011); Cho and Rust (2010); Goldfarb and Xiao (2018); Ellison et al. (2018); DellaVigna and Gentzkow (2019)).

We caution, however, that this conclusion, comes at the cost of several strong assumptions including constant marginal costs, and largely abstracting from multi-product and/or competitive effect considerations. We have discussed the credibility and robustness of our estimates associated with each of these assumptions above, and broadly view them as unlikely to strongly alter the results reported here. An assumption unique to the lost profit calculation is that average prices over the year are set optimally. This assumption is of course weaker than what is commonly made in the literature. We have no way to verify the optimal average price assumption and we view it as the most critical and perhaps the most suspect. Testing this assumption would likely require, at a minimum, data on marginal cost, although such data might be difficult to obtain given that the marginal cost to a retailer includes the opportunity cost of retail space. Evidence on the optimality of the average level of prices would be an important question for future research.

## **5 Seasonality on the intensive and extensive margin**

We find that most products exhibit countercyclical price fluctuations, and these fluctuations are broadly consistent with a simple benchmark pricing model, because the demand curve facing retailers becomes more elastic during periods of demand peaks. What explains this time-varying elasticity? A natural candidate is that the identity of shoppers in a given category changes over the year. For example, the average soup buyer in July may be quite different than the average soup buyer in February. If people who only purchase during peak demand periods have more elastic demand, than the aggregate demand curve facing a retailer will become more elastic during peak demand periods, even if at an individual level the demand curve has a time-invariant elasticity. Several researchers have proposed and investigated versions of this explanation for countercyclical pricing for particular products (e.g. Guler et al. (2014); Bayot and Caminade (2015); Per-

rone (2016); Kwon et al. (2018); Haviv (2019)). Here, we test an implication of this idea, using a wide range of products. Specifically, for this compositional explanation to be correct, it must be true that categories with larger seasonal shifts in demand experience large *extensive margin* shifts, i.e. large changes in the fraction of people purchasing any products in the category at all.

We use a simple decomposition to show the importance of the extensive margin for overall seasonality. Specifically, the change in average quantity purchased (at the individual level) between two months  $m$  and  $m'$ , can be written as

$$E[q|m] - E[q|m'] = (1/2) (E[q|q > 0, m] + E[q|q > 0, m']) \underbrace{(Pr(q > 0|m) - Pr(q > 0|m'))}_{\equiv \Delta Pr} + (1/2) (Pr(q > 0|m) + Pr(q > 0|m')) \underbrace{(E[q|q > 0, m] - E[q|q > 0, m'])}_{\equiv \Delta ECOP}.$$

This is the difference in probability of any purchase (which we call  $\Delta Pr$ ), scaled by the expected amount purchased conditional on making a purchase, plus the change in the expected amount purchased conditional on making a purchase (which we call  $\Delta ECOP$ ). We therefore define the extensive margin share as

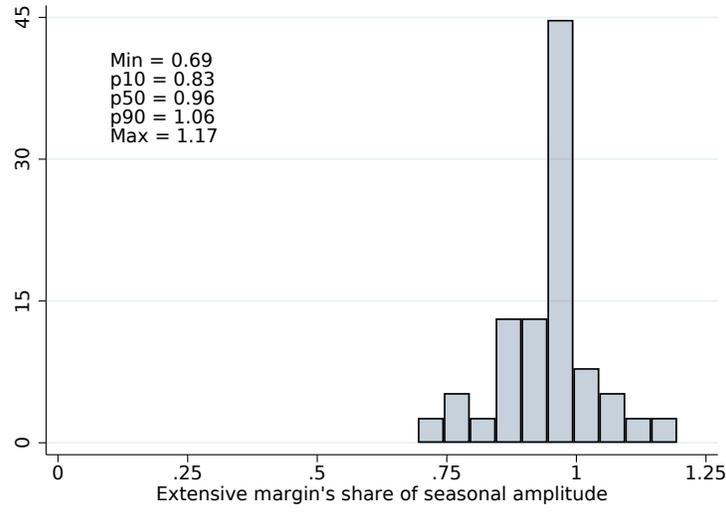
$$\text{Extensive share}_{m,m'} = \frac{\Delta Pr(1/2) (E[q|q > 0, m] + E[q_{m'}|q > 0, m'])}{E[q|m] - E[q|m']}. \quad (9)$$

If changes in the composition of consumers drive changes in the store-level elasticity, then we should observe that the extensive share accounts for a large share of seasonal changes in quantities, i.e. a large share of the peak-trough difference  $E[q_{\bar{m}}] - E[q_m]$ .

To estimate the extensive share, we use Nielsen's household level dataset, the Consumer Panel. We aggregate to the household-week-category (i.e. module) level and record total purchases, whether a household has any purchases, and the amount purchased if positive. Then for each month we calculate the empirical analogs of  $E[q|Month = m]$ ,  $Pr(q > 0|Month = m)$ , and  $E[q|q > 0, month = m]$ , and we use these to seasonal amplitudes and the extensive margin share.<sup>25</sup>

<sup>25</sup>For consistency with our main analysis, we exclude holiday weeks. However, in the household data we do not adjust for prices or other factors because there is not a straightforward way to measure the price

Figure 7: The extensive margin accounts for most of the seasonal variation in demand



Notes: This figure reports the distribution, across modules, of the extensive margin's contribution to seasonal amplitudes. We estimate the extensive margin share using Equation 9 and Nielsen's Consumer Panel.

We show the distribution of extensive margin shares across modules in Figure 7. For all our modules, the extensive margin accounts for a majority of seasonal fluctuations (including modules in which these fluctuations are small). For the median module, the extensive margin accounts for 96 percent of the seasonal fluctuation, and for some modules, the extensive margin accounts for greater than 100 percent of seasonality. This is possible because, for these modules, the intensive margin (the amount purchased conditional on a purchase) actually declines when demand peaks. These intensive margin declines are consistent with both large extensive margin shifts and considerable consumer heterogeneity, with the marginal customers who purchase during only peak periods having much lower demand than the consumers who purchase during demand troughs. Thus seasonal shifts in demand are accompanied by large shifts in the composition of buyers. As these buyers have less attachment to the product category, it is likely that they are more price sensitive and hence, much more elastic. These extensive margin response therefore help explain why we find that retailers face time-varying elasticities.

a customer faces when she buys no product. We show in Figure A.4 in appendix, however, that seasonality in the household data is very similar to (price-adjusted) seasonality in the retail scanner data.

## 6 Conclusions

Seasonality in demand is large, pervasive, and heterogeneous. The typical product experiences demand swings of 25 log points or more, and most products are seasonal in the sense that their seasonal demand fluctuations are larger than their business cycle fluctuations. The timing of these fluctuations is heterogeneous: for a plurality of products, demand peaks in the winter, but many products show a summer peak. Despite these large demand fluctuations, prices vary little at seasonal frequencies, and prices are typically countercyclical. For the average seasonal products, prices fall by a few percent from demand trough to demand peak, and peak demand months are usually trough price months. Roughly half of the countercyclical price changes come from greater frequency of promotional prices. These price changes are for a fixed product and therefore avoid compositional bias, which we show results in even greater measured countercyclicality.

We show that the simplest model of optimal pricing can roughly account for countercyclical prices. The model implies that prices depend on demand elasticities rather than levels, and we find that, indeed, demand elasticities typically become more negative when the level of demand peaks. Elasticity and price movements are temporally closely aligned; we find that demand is often most elastic in the month when prices are most negative. Quantitatively, we find a roughly one-for-one relationship between seasonal price changes and seasonal changes in the theoretical benchmark price, looking across product categories. Our benchmark model in fact implies that prices should be somewhat more countercyclical than they are. Although observed seasonal price changes are not exactly equal to our benchmark, the lost profits from mispricing relative to our benchmark are a small share of revenue. Thus despite abstracting from many important features such as varying marginal cost, intertemporal considerations, or competitive effects, the simple pricing model can rationalize seasonal pricing patterns. Key to this rationalization is the fact that demand becomes more elastic during seasonal peaks. We argue that the mechanism behind this countercyclical elasticity is a changing composition of consumers. People who buy soup in summer likely have a strong attachment to soup and low price sensitivity, while people who buy soup only in winter are likely more price sensitive.

In support of this mechanism, we show using household-level data that nearly the entirety of seasonal fluctuations in quantities comes from the extensive margin, shifts in the fraction of people buying any product in a given month and category. Thus, we have provided new evidence on the extent of seasonality that is substantially broader than previous work, and shown that simple models can largely rationalize apparent pricing puzzles.

## References

- Allcott, H., B. B. Lockwood, and D. Taubinsky (2019, 05). Regressive Sin Taxes, with an Application to the Optimal Soda Tax\*. *The Quarterly Journal of Economics* 134(3), 1557–1626.
- Anderson, E., B. A. Malin, E. Nakamura, D. Simester, and J. Steinsson (2017). Informational rigidities and the stickiness of temporary sales. *Journal of Monetary Economics* 90(C), 64–83.
- Angrist, J. D., P. D. Hull, P. A. Pathak, and C. R. Walters (2017). Leveraging lotteries for school value-added: Testing and estimation. *Quarterly Journal of Economics* 132, 871–919.
- Barsky, R. B. and J. A. Miron (1989). The seasonal cycle and the business cycle. *Journal of Political Economy* 97(3), 503–534.
- Bayot, D. and J. Caminade (2015). Popping the cork: The role of the extensive margin in countercyclical pricing.
- Bils, M. and P. Klenow (2004). Some evidence on the importance of sticky prices. *Journal of Political Economy* 112(5), 947–985.
- Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12(3), 383–398.
- Chevalier, J. A., A. K. Kashyap, and P. E. Rossi (2003). Why don't prices rise during periods of peak demand? Evidence from scanner data. *The American Economic Review* 93(1), 15–37.
- Cho, S. and J. Rust (2010). The flat rental puzzle. *The Review of Economic Studies* 77(2), 560–594.
- Coibion, O., Y. Gorodnichenko, and G. H. Hong (2015). The Cyclicity of sales, regular and effective prices: Business cycle and policy implications. *The American Economic Review* 105(3), 993–1029.
- Cooper, R. and J. Haltiwanger (1993). Automobiles and the national industrial recovery act: Evidence on industry complementarities. *The Quarterly Journal of Economics* 108(4), 1043–1071.
- DeGraba, P. (2006). The loss leader is a turkey: Targeted discounts from multi-product competitors. *International Journal of Industrial Organization* 24(3), 613 – 628.
- DellaVigna, S. and M. Gentzkow (2019, 06). Uniform Pricing in U.S. Retail Chains\*. *The Quarterly Journal of Economics* 134(4), 2011–2084.
- Eichenbaum, M., N. Jaimovich, S. Rebelo, and J. Smith (2014). How frequent are small price changes? *American Economic Journal: Macroeconomics* 6(2), 137–155.

- Einav, L. (2007). Seasonality in the U.S. motion picture industry. *The RAND Journal of Economics* 38(1), 127–145.
- Einav, L., E. Leibtag, and A. Nevo (2010, June). Recording discrepancies in Nielsen Homescan data: Are they present and do they matter? *Quantitative Marketing and Economics* 8(2), 207–239.
- Ellison, S. F., C. Snyder, and H. Zhang (2018). Costs of managerial attention and activity as a source of sticky prices: Structural estimates from an online market. Working Paper 24680, National Bureau of Economic Research.
- Gagnon, E. and D. López-Salido (2019, 02). Small Price Responses to Large Demand Shocks. *Journal of the European Economic Association*. jvz002.
- Goldfarb, A. and M. Xiao (2011). Who thinks about the competition: Managerial ability and strategic entry in US local telephone markets. *American Economic Review* 101, 3130–3161.
- Goldfarb, A. and M. Xiao (2018). Transitory shocks, limited attention, and a firm’s decision to exit. Working paper.
- Guler, A. U., K. Misra, and N. Vilcassim (2014). Countercyclical pricing: A consumer heterogeneity explanation. *Economics Letters* 122(2), 343–347.
- Hausman, J. A. (1996, January). *Valuation of New Goods under Perfect and Imperfect Competition*, pp. 207–248. University of Chicago Press.
- Haviv, A. (2019). Consumer search, price promotions, and counter-cyclic pricing. Working Paper FR. 15-19, University of Rochester, Simon Business School.
- Hendel, I. and A. Nevo (2003). The post-promotion dip puzzle: What do the data have to say? *Quantitative Marketing and Economics* 1(4), 409–424.
- Hendel, I. and A. Nevo (2013). Intertemporal price discrimination in storable goods markets. *The American Economic Review* 103(7), 2722–2751.
- Hitsch, G. J., A. Hortaçsu, and X. Lin (2017). Prices and promotions in U.S. retail markets: Evidence from big data. Working Paper 17-18, Chicago Booth Research Paper.
- Hortaçsu, A. and S. L. Puller (2008). Understanding strategic bidding in multi-unit auctions: A case study of the Texas electricity spot market. *The RAND Journal of Economics* 39(1), 86–114.
- Hosken, D. and D. Reiffen (2004a). How retailers determine which products should go on sale: Evidence from store-level data. *Journal of Consumer Policy* 27(2), 141–177.
- Hosken, D. and D. Reiffen (2004b). Patterns of retail price variation. *The RAND Journal of Economics* 35(1), 128–146.

- Jacob, B. A. and L. Lefregren (2008). Can principals identify effective teachers? evidence on subjective performance evaluation in education. *Journal of Labor Economics* 26, 101–136.
- Kahneman, D., J. L. Knetsch, and R. Thaler (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *The American Economic Review* 76(4), 728–741.
- Kane, T. J. and D. O. Staiger (2008). Estimating teacher impacts on student achievement: An experimental evaluation. NBER Working Paper No. 14607.
- Kwon, M., T. Erdem, and M. Ishihara (2018). Counter-cyclical price promotion: Capturing seasonal category expansion under endogenous consumption. Working paper, New York University, Stern School of Business.
- Lal, R. and C. Matutes (1994). Retail pricing and advertising strategies. *The Journal of Business* 67(3), 345–370.
- MacDonald, J. M. (2000). Demand, information, and competition: Why do food prices fall at seasonal demand peaks? *The Journal of Industrial Economics* 48(1), 27–45.
- Meza, S. and K. Sudhir (2006). Pass-through timing. *Quantitative Marketing and Economics* 4(4), 351–382.
- Nakamura, E. and J. Steinsson (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics* 123(4), 1415–1464.
- Nakamura, E. and J. Steinsson (2013). Price rigidity: Microeconomic evidence and macroeconomic implications. *Annual Review of Economics* 5(1), 133–163.
- Nevo, A. and K. Hatzitaskos (2005). Why does the average price of tuna fall during Lent? Working Paper 11572, National Bureau of Economic Research.
- Orbach, B. Y. and L. Einav (2007). Uniform prices for differentiated goods: The case of the movie-theater industry. *International Review of Law and Economics* 27(2), 129–153.
- Pashigian, B. P. and B. Bowen (1991). Why are products sold on sale?: Explanations of pricing regularities. *The Quarterly Journal of Economics* 106(4), 1015–1038.
- Perrone, H. (2016). Consumers' quality choices during demand peaks. *International Journal of Industrial Organization* 44, 154–162.
- Strulov-Shlain, A. (2019). More than a penny's worth: Left-digit bias and firm pricing. Working Paper 19-22, Chicago Booth Research Paper.
- Suits, D. B. (1984). Dummy variables: Mechanics v. interpretation. *The Review of Economics and Statistics* 66(1), 177–180.
- Warner, E. J. and R. B. Barsky (1995). The timing and magnitude of retail store markdowns: Evidence from weekends and holidays. *The Quarterly Journal of Economics* 110(2), 321–352.

# For Online Publication

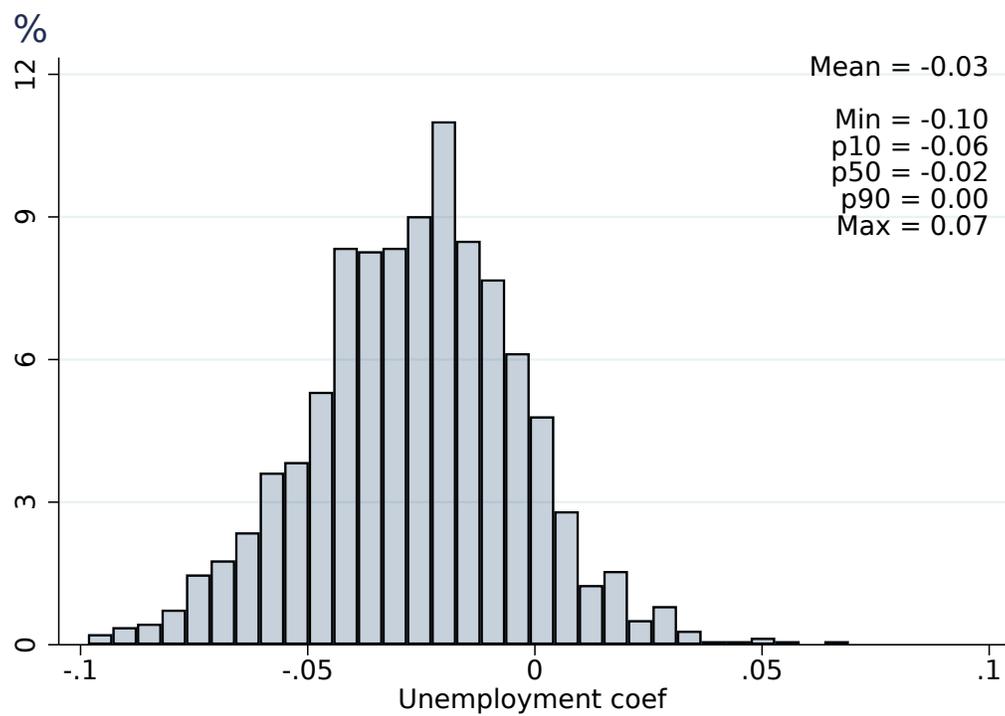
## A Additional results

Table A.1: Selecting chains and stores

Restriction imposed	No. of Stores	No. of Chains	No. of States	Total yearly revenue
Initial sample of stores	40,070	398	49	\$225bn
Food, drug, and mass merchandise only	38,052	390	49	\$222bn
Stores do not switch chain	34,519	133	49	\$208bn
Stores $\geq 2$ years	33,595	121	49	\$208bn
Stores in Household panel	33,595	121	49	\$208bn
Chain present $\geq 8$ years	24,597	79	49	\$194bn
Valid chain	24,471	77	49	\$191bn
Base price algorithm (baseline sample)	24,450	77	49	\$191bn

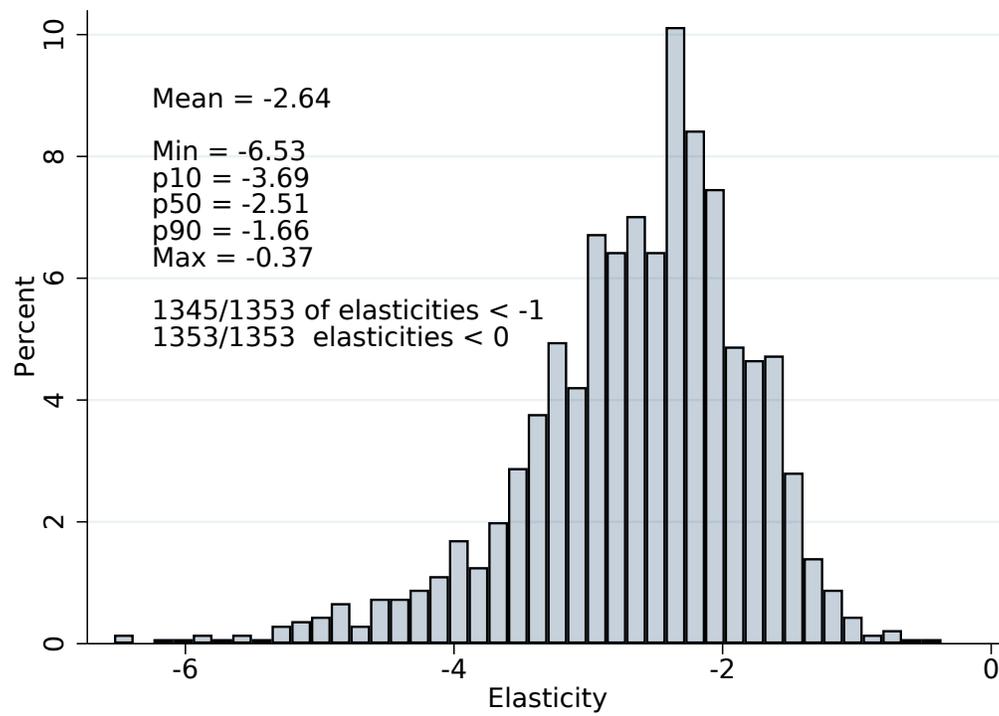
Notes: The table shows the breakdown of stores and chains as we refine our samples. This table follows Table 1 of DellaVigna and Gentzkow (2019) except the last panel aggregates over all types of stores, not just food stores. Total yearly revenue is the aggregate sales divided by 9 years.

Figure A.1: Unemployment coefficients



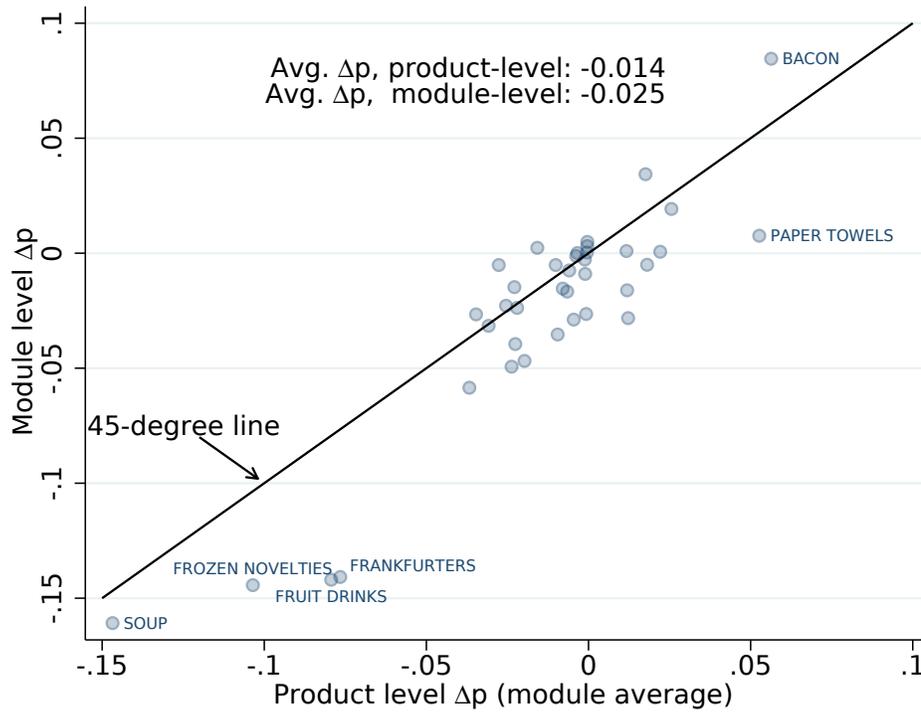
Notes: Figure plots the distribution (across products) of estimated coefficients on the local unemployment rate (measured on a 0-100 scale), from Equation 1.

Figure A.2: Elasticity distribution, across products



Notes: Figure plots the distribution of estimated elasticities, across products, from Equation 1.

Figure A.3: Module-level seasonal aggregate price changes vs. product-level seasonal price changes



Notes: Figure illustrates how composition bias can lead to greater countercyclicality in module level average prices than in the individual prices. To do so, we estimate Equation 1 on data aggregated to the module-store-week level, yielding an estimate of peak and trough quantity months. We then estimate Equation 5 using module-store-week log average price as the dependent variable (using a revenue weighted average).  $\Delta p$  at the module level is the difference in log average price between peak and trough months. We plot that against the module-level average of the product level  $\Delta p$  (averaging these using equal weights). Composition bias makes the module level price change more negative than the product level price change.

Table A.2: Seasonality by module

Module	Products	% Seasonal	Quantity			Price changes		
			Amp	Mode peak month	% Jan peak	ln p	ln base price	promo rate
SOUP-CANNED	79	1	.71	1	.68	-.15	-.06	.19
FROZEN NOVELTIES	22	1	.63	7	0	-.1	-.06	.11
BATTERIES	5	1	.53	12	0	0	0	.02
FRUIT DRINKS-OTHER CONTAINER	101	.91	.44	7	.04	-.08	-.06	.05
FRESH FRUIT-REMAINING	1	1	.43	5	0	-.02	-.02	.06
FRANKFURTERS-REFRIGERATED	18	.83	.39	6	0	-.08	-.04	.09
WATER-BOTTLED	30	.83	.38	7	.17	-.02	-.01	.03
YOGURT-REFRIGERATED	90	.96	.38	1	.3	-.02	0	.08
CANDY-NON-CHOCOLATE	27	.93	.37	6	0	-.02	-.02	0
CANDY-CHOCOLATE	87	.95	.37	12	.05	-.03	-.02	.02
BAKERY-CAKES-FRESH	9	1	.35	5	.22	-.03	0	.09
BEER	6	1	.34	6	0	0	0	0
ICE CREAM - BULK	17	.94	.29	7	0	-.04	-.01	.08
PAPER TOWELS	10	.6	.28	7	0	.05	.03	-.06
LIGHT BEER (LOW CALORIE/ALCOHOL)	6	.83	.28	7	0	0	0	0
COOKIES	85	.82	.28	2	.13	-.03	-.01	.06
SNACKS - TORTILLA CHIPS	53	.51	.27	1	.36	-.01	0	.01
CEREAL - READY TO EAT	194	.8	.27	1	.3	-.02	-.01	.04
CHEESE - SHREDDED	7	1	.26	1	.71	.02	0	-.01
BACON-REFRIGERATED	8	.25	.24	8	.13	.06	.03	-.05
BLEACH - LIQUID/GEL	11	.91	.23	8	0	-.01	0	.02
LUNCHMEAT-DELI POUCHES-REFRIGERATED	21	.81	.23	7	0	-.01	0	.03
DETERGENTS - HEAVY DUTY - LIQUID	21	.52	.22	1	.38	0	0	.02
SNACKS - POTATO CHIPS	79	.52	.22	6	.05	-.01	0	.03
SOFT DRINKS - LOW CALORIE	66	.76	.22	1	.27	.01	0	0
EGGS-FRESH	2	.5	.22	12	0	.03	.01	-.03
ENTREES - ITALIAN - 1 FOOD - FROZEN	33	.94	.21	1	.3	-.03	0	.07
GROUND AND WHOLE BEAN COFFEE	17	.53	.21	2	.06	.01	.01	.01
TOILET TISSUE	27	.7	.2	12	.3	.01	.01	-.01
FRUIT JUICE - ORANGE - OTHER CONTAINER	39	.33	.2	1	.44	0	0	.01
LUNCHMEAT-SLICED-REFRIGERATED	13	.69	.19		.31	0	0	.01
ENTREES-REFRIGERATED	3	.33	.18		0	.02	.02	.03
PIZZA-FROZEN	25	1	.18	1	.6	-.02	-.01	.03
SOFT DRINKS - CARBONATED	73	.66	.18	1	.22	0	-.01	-.01
PAIN REMEDIES - HEADACHE	2	.5	.14	12	0	0	0	-.01
DOG FOOD - DRY TYPE	8	.88	.14		.25	.02	.01	-.05
CAT FOOD - WET TYPE	54	.7	.08		.22	0	0	.01
DAIRY-MILK-REFRIGERATED	4	0	.06	1	1	-.01	-.01	0

Notes: Sorted by amplitude. Table reports, for each module, the number of products, the mean seasonal amplitude (“amp”), the modal peak quantity month (if unique, otherwise missing), and the percent of products whose peak month is January, as well as the seasonal change in log price, log base price, promotional price rate, and benchmark optimal price.

Table A.3: Robustness table for seasonal quantity amplitudes

Specification	# Products	% Seasonal	Seasonal amplitude							
			Mean	p1	p10	p25	p50	p75	p90	p99
Baseline	1353	0.79	0.35	0.05	0.13	0.19	0.28	0.46	0.69	1.10
DMA competition	1350	0.77	0.35	0.05	0.14	0.19	0.29	0.47	0.70	1.08
Other product competition - average	1351	0.79	0.36	0.05	0.13	0.19	0.29	0.48	0.70	1.15
Other product competition - minimum	1353	0.79	0.35	0.05	0.13	0.19	0.28	0.46	0.69	1.10
Time since last sale	1334	0.78	0.36	0.05	0.13	0.19	0.29	0.47	0.70	1.12
Inverse hyperbolic sine	1265	0.76	0.50	0.06	0.18	0.26	0.41	0.63	0.93	1.83
OLS with price adjustment	1531	0.80	0.37	0.05	0.13	0.20	0.30	0.48	0.75	1.15
OLS without price adjustment	1531	0.88	0.46	0.06	0.16	0.24	0.36	0.57	0.92	1.50
No shrinkage	1353	0.83	0.37	0.06	0.15	0.21	0.30	0.48	0.71	1.15

Notes: Each row is a different specification. For each specification, we exclude products with too-large standard errors, and then report the number of products, the percent seasonal, and statistics on the distribution of seasonal amplitudes in quantities. “DMA competition” controls for the average log price charged by other chains in the same DMA. “Other product competition - average” controls for the average log price of other products in the same module and store. “Other product - minimum” instead controls for the minimum log price of other products in the same module and store. “Time since sale” controls for weeks since the last promotional pricing period. Inverse hyperbolic sine uses the inverse hyperbolic sine,  $\ln \sqrt{q^2 + 1}$ , as the dependent variable. OLS with price adjustment is identical to our baseline specification, except we estimate by OLS rather than 2SLS. OLS without price adjustment drops all price controls, and includes only product-month dummies, product-store-year fixed effects, and product-unemployment controls.

Table A.4: Robustness table for seasonal price changes

Specification	# Seasonal	Price change, peak q month minus trough q month								
		% negative	Mean	p1	p10	p25	p50	p75	p90	p99
Baseline	1067	0.68	-0.03	-0.27	-0.13	-0.06	-0.02	0.00	0.03	0.11
DMA competition	1041	0.70	-0.04	-0.27	-0.13	-0.06	-0.02	0.00	0.03	0.09
Other product competition - average	1073	0.71	-0.04	-0.27	-0.13	-0.06	-0.02	0.00	0.02	0.08
Other product competition - minimum	1072	0.68	-0.03	-0.27	-0.13	-0.05	-0.02	0.00	0.03	0.11
Time since last sale	1043	0.68	-0.03	-0.25	-0.13	-0.06	-0.02	0.00	0.03	0.10
Inverse hyperbolic sine	966	0.69	-0.04	-0.25	-0.13	-0.06	-0.02	0.00	0.02	0.09
No shrinkage	1129	0.68	-0.03	-0.25	-0.12	-0.05	-0.02	0.01	0.03	0.10

Notes: Each row is a different specification. For each specification, we exclude products with too-large standard errors, and then report the number of seasonal products, the percent with negative price changes between seasonal trough and peak, and the distribution across products of seasonal price changes. See notes to Table A.3 for an explanation of the different specifications.

Table A.5: Robustness table for seasonal elasticity changes

Specification	# Seasonal	Elasticity change, peak q month minus trough q month									
		% negative	Mean	p1	p10	p25	p50	p75	p90	p99	
Baseline	1067	0.71	-0.38	-3.50	-1.19	-0.58	-0.18	0.00	0.20	1.00	
DMA competition	1041	0.69	-0.37	-3.40	-1.13	-0.58	-0.17	0.02	0.20	0.98	
Other product competition - average	1073	0.71	-0.36	-3.38	-1.09	-0.57	-0.16	0.01	0.18	0.97	
Other product competition - minimum	1072	0.70	-0.38	-3.35	-1.16	-0.58	-0.19	0.00	0.18	0.98	
Time since last sale	1043	0.70	-0.39	-3.67	-1.11	-0.59	-0.19	0.00	0.19	0.98	
No shrinkage	1129	0.73	-0.58	-4.34	-1.72	-0.98	-0.40	0.04	0.44	1.53	

Notes: Each row is a different specification. For each specification, we exclude products with too-large standard errors, and then report the number of seasonal products, the percent with negative elasticity changes between seasonal trough and peak, and the distribution (across products) of seasonal elasticity changes. See notes to Table A.3 for an explanation of the different specifications.

Table A.6: Robustness table for price-benchmark price association

Demand specification	# Modules	Constant	(SE)	Slope	(SE)
Baseline	37	-0.08	(0.02)	2.05	(0.53)
DMA competition	37	-0.05	(0.03)	2.22	(0.54)
Other product competition - average	37	-0.06	(0.02)	1.78	(0.44)
Other product competition - minimum	37	-0.07	(0.02)	2.10	(0.51)
Time since last sale	37	-0.13	(0.05)	1.83	(0.68)

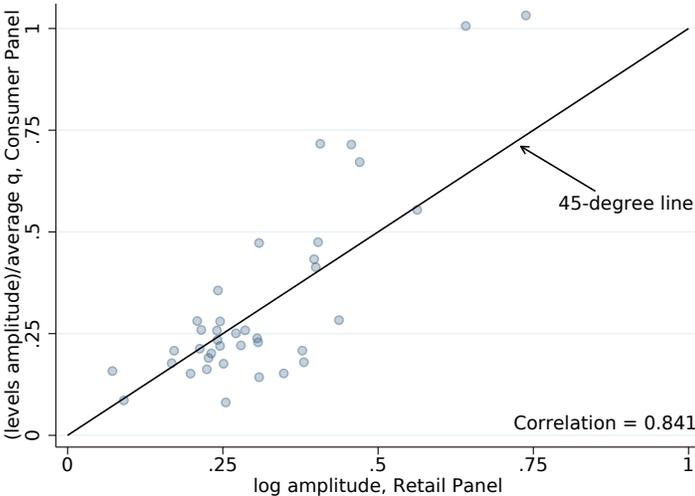
Notes: Table reports the results of a module level regression of the (module average) change in benchmark price against the (module average) observed change in price, among seasonal products and excluding fresh fruit. Each row uses a different specification to determine the change in benchmark price and benchmark price. See notes to Table A.3 for an explanation of the different specifications.

Table A.7: Time-since-sale matters much more for storable products than for perishable products

Product grouping	Average $\Delta\eta$ , Baseline	Average $\Delta\eta$ controlling for weeks-since-sale	Average coefficient on weeks-since-sale
Perishable	-0.033	-0.026	0.0004
Storable	-0.407	-0.431	0.0012

Notes: Table shows the validity of the weeks-since-sale control as a proxy for storability incentives. While the estimated change in elasticities is robust to controlling for weeks-since-sale, we see that weeks-since-sale has a much stronger effect on quantities for storable products than for perishable ones, as we would expect from a model of storability. Table reports the average  $\Delta\eta$  (across products) in the baseline specification and in the robustness specification, and the average coefficient on week-since-sale. The storable products are defined as canned soup, toilet tissue, soft drinks (carbonated or low calorie), and detergent. The non-storable products are orange juice, yogurt, and milk.

Figure A.4: Seasonal amplitudes are highly correlated between the Retail Panel and Consumer Panel



Notes: Figure plots seasonal amplitudes in quantities estimated in the household data (at the module level) against seasonal amplitudes estimated in the retail panel (aggregated to the module level). The estimation of the retail panel amplitudes is described in the text; it is based on shrunken monthly differences in log quantities. To estimate seasonal amplitudes in the household data, we look at the difference in the level of weekly purchases between the peak and trough month, scaled by the average purchase amount across those months (so that the scale is comparable to the retail panel scaling).

## B Base Price Algorithm

### B.1 Overview

We suggest a simple and straightforward algorithm to construct “base” prices allowing for time-varying base prices. Traditionally (e.g., Hendel and Nevo (2003, 2013)), ad-hoc procedures have been used to establish these base prices (e.g., single base price over a predetermined time period), and thus subsequently the promotional periods that are further inferred from price deviations (e.g., beyond a threshold) from the base price. For example, Hendel and Nevo (2003) establish the base price at a particular period to be the modal price observed for that product over the two years, and define a “promo” as any period where the price is 5 percent below this level. Our algorithm is different from the literature in three ways. First, we do not predetermine the number of base prices over a fixed length of time. Second, the algorithm is robust to various depth of either price dips or jumps (i.e., we do not need to take a stance on the minimum/maximum price difference between base and non-base prices). Third, selection of the parameters for the algorithm can be guided by documented practices of grocery retailers, as well as the price path and promotion activity for stores with information on both.

The algorithm is robust to instances of retailers in the Nielsen retail scanner dataset which have pricing schedules that do not align with the Nielsen reporting week. For those retailers and weeks with multiple prices our base price algorithm is likely to be less prone to incorrectly inferring changes in the base price, because of the algorithms built in accommodation for measurement error in prices.

As an example, Figure B.1 plots observed prices and the imputed base prices for canned soup for the store with the most canned soup sales in our data. Inspection of the figure reveals that each year has one or two obvious base prices (except 2009), and the algorithm correctly identifies these as such. In some years price is a base price but in other years it is a promotional price, and the algorithm successfully distinguishes between these cases. Our algorithm does well at avoiding false negatives, and successfully distinguishes between promotions and base prices that occur at the same dollar value. On the other hand, the algorithm has trouble with long-lasting spells of promotional prices; for example, in 2006 and 2007, the algorithm detects a base price drop where there is (likely) none. Implementing the algorithm requires that we specify hyperparameters that let us trade off these false negatives and false positives. In selecting these hyperparameters, we have erred on the side of caution, with a low false negative rate for base price changes.

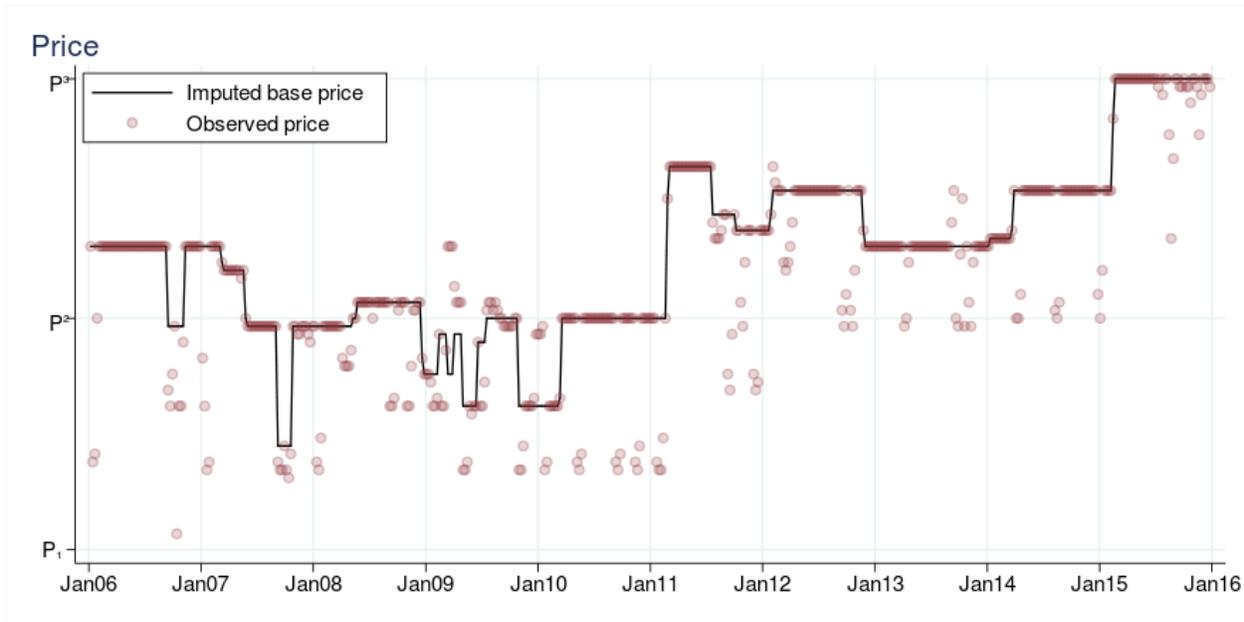
### B.2 Base price algorithm in detail

From a time series of observed weekly prices,  $(p_t)_{t=\tau}^{t=T}$ , for each product and each store, the algorithm constructs a time series of “base” prices,  $(b_t)_{t=\tau}^{t=T}$ , over the same set of weeks.<sup>26</sup> The algorithm requires two parameters—the maximum *consecutive* weeks al-

---

<sup>26</sup> $(p_t)_{t=\tau}^{t=T}$  would not have any missing values as they are already imputed by the median chain-state-week prices.

Figure B.1: Example of base and actual prices



Notes: The figure plots the weekly time series of the observed price per unit of a 10.75 oz can of Campbell's Chicken Noodle Soup, and the imputed base price, for the largest store in our data, given  $\gamma_1 = 5$  and  $\gamma_2 = 7$ . Our data use agreement prohibits us from disclosing the price of individual UPCs, so we cannot label the y-axis of this figure.

lowed for a promo period ( $\gamma_1 - 1$ ) and the minimum length of weeks required for non-promo period ( $\gamma_2$ ).

To present in more detail the specifics of the algorithm, an example time series of prices ( $p_t$ ) is presented in the second column of table B.2. The price path exhibited in table B.2 indicates that a typical price dip lasts one week. From  $p_t$ , it is evident that there are likely two sets of base prices. The first base price is 1.0 during weeks 1-3 and 8-11. The second base price is 0.8 during weeks 4-7. Among these base prices, a promotional price of 0.4 is run during week 4, and 0.6 during week 6. Akin to what would be experienced in the price paths of the Nielsen dataset, it appears that a modest deviation from the base price is experienced during week 9 when the observed price *increases* to \$1.1 for just one week. Ideally, the base price algorithm would be able to distinguish between the two base prices of a 1.0 and 0.8, without incorrectly attributing the other prices observed in the time series to base prices.

The first step of the algorithm detects smooth base prices during price dips. We use the hyperparameter ( $\gamma_1$ ) and construct the lower envelope of both a "forward" and "backward"  $\gamma_1$ -period cumulative maximum of the observed time series of prices to construct the candidate time series of base prices. In the example, we let  $\gamma_1 = 2$  based on the observation that price dips last one week typically. The "forward" 2-period cumulative maximum ( $p_t^f$ ) is reported in the third column, while the "backward" 2-period cumulative

Table B.1: Base Price Summary Statistics

	Mean	25th	Median	75th
% missing	0.0032	0	0	0
No. of base price	4.46	3	4	6
% p > base price	0.07	0	0.06	0.12
% promo	0.22	0.08	0.21	0.35
Promo depth (%)	-0.25	-0.3	-0.24	-0.18

The table reports the summary statistics of base price and promotion per store-year given  $\gamma_1 = 5$  and  $\gamma_2 = 7$ .

maximum ( $p_t^b$ ) is reported in the fourth column. To construct the lower envelope of these two price series, we take the period-wise minimum of both of these series to construct the first step’s candidate base price time series. We call this step of the algorithm the minimax filter-smoother step, and report the resulting time series in the fifth column ( $\tilde{b}_t$ ).

In case of price paths where promotions never lasted longer than  $\gamma_1 - 1$  periods, the minimax filter-smoother step could almost perfectly identify base prices from promotional prices. However, in the event that small week-to-week deviations in prices occur for likely idiosyncratic reasons tied to measurement or reporting issues, this single step will still yield a modest amount of noise in the construction of the base price. Given the minimax formulation of the first step, positive measurement errors, like the one experienced in week nine, are likely to be the most influential to the base price construction. For example, at the culmination of the minimax step of the algorithm, week nine posits that the base price for that period is 1.1, a price that only occurs once in the time series, and is much more likely to be the consequence of measurement error.

To ameliorate the algorithm’s sensitivity to these sorts of measurement errors, the second step of the algorithm detects smooth base prices during price jumps or short dips. In particular, the maximum length of the autocorrelation of measurement errors will be controlled by the hyperparameter,  $\gamma_2$ . Another way to view this hyperparameter is the minimum uninterrupted sequence length a base price must have over the duration of its spell.

We let the upper envelope of both a “forward” and “backward”  $\gamma_2$ -period cumulative minimum of the time series of candidate base prices resulting from the first step ( $\tilde{b}_t$ ) to construct the final candidate time series of base prices, ( $b_t$ ). In our example, data shows that the typical plateau periods last at least two weeks, and therefore we set  $\gamma_2 = 2$ . In table B.2, the “forward” 2-period cumulative minimum of the first step’s base price time series is reported in the sixth column ( $\tilde{b}_t^f$ ). The “backward” 2-period cumulative minimum of the observed time series is reported in the seventh column ( $\tilde{b}_t^b$ ). To construct the upper envelope of these two price series, we take the period-wise maximum of both of these series to construct the second step’s final candidate base price time series,  $b_t$ . We call this step of the algorithm the maximin filter-smoother step. As was desired, this second step has the effect of smoothing over the short-lived idiosyncratic price increases that are less likely to be base prices and more likely to be the result of measurement error.

Table B.2: An Example Price Path: Minimax/Maximin filter-smoother

$t$	$p_t$	$p_t^f$	$p_t^b$	$\tilde{b}_t$	$\tilde{b}_t^f$	$\tilde{b}_t^b$	$b_t$
1	1.0	-	1.0	1.0	-	1.0	1.0
2	1.0	1.0	1.0	1.0	1.0	1.0	1.0
3	1.0	1.0	1.0	1.0	1.0	0.8	1.0
4	0.4	1.0	0.8	0.8	0.8	0.8	0.8
5	0.8	0.8	0.8	0.8	0.8	0.8	0.8
6	0.6	0.8	0.8	0.8	0.8	0.8	0.8
7	0.8	0.8	1.0	0.8	0.8	0.8	0.8
8	1.0	1.0	1.1	1.0	0.8	1.0	1.0
9	1.1	1.1	1.1	1.1	1.0	1.0	1.0
10	1.0	1.1	1.0	1.0	1.0	1.0	1.0
11	1.0	1.0	-	1.0	1.0	-	1.0

Notes: The table provides an example of the minimax (first) filter-smoother step of the base price algorithm, as well as the maximin (second) filter-smoother step for a hypothetical observed time series path of prices. The second column presents the observed price series ( $p_t$ ). The third column ( $p_t^f$ ) presents the “forward”  $\gamma_1 = 2$ -period cumulative maximum, while the fourth column ( $p_t^b$ ) presents the “backward”  $\gamma_1 = 2$ -period cumulative maximum. The fifth column ( $\tilde{b}_t$ ) presents the lower envelope of the two cumulative maximum series by taking the period-wise minimum. The sixth column presents the resulting price series from the first step ( $\tilde{b}_t$ ). The seventh column ( $\tilde{b}_t^f$ ) presents the “forward”  $\gamma_2 = 2$ -period cumulative minimum, while the eighth column ( $\tilde{b}_t^b$ ) presents the “backward”  $\gamma_2 = 2$ -period cumulative minimum. Finally, the last column ( $b_t$ ) presents the upper envelope of the two cumulative minimum series by taking the period-wise minimum.

In summary, the base price algorithm is a straightforward implementation of two filter-smoother steps that both have the interpretation of a minimax (maximin) approach. Both steps are parameterized by separate hyperparameters,  $\gamma_1$  and  $\gamma_2$ , respectively. The specific details of each step in the algorithm are provided below.

### Steps for Base Price Identification

For any given observed time series of prices  $(p_t)_{t=\tau}^{t=T}$  and a set of hyperparameters  $(\gamma_1, \gamma_2)$ ,

#### 1. **Minimax**( $\gamma_1$ ) filter-smoother:

- Construct  $\gamma_1$  period forward cumulative max of the observed price series,  $p_t^f$ .

$$p_t^f = \max\{p_t, p_{t-1}, \dots, p_{t-(\gamma_1-1)}\}, \quad \forall \tau + (\gamma_1 - 1) \leq t \leq T$$

- Construct  $\gamma_1$  period backward cumulative max of the observed price series,  $p_t^b$ .

$$p_t^b = \max\{p_t, p_{t+1}, \dots, p_{t+(\gamma_1-1)}\}, \quad \forall \tau \leq t \leq T + (\gamma_1 - 1)$$

- Take period-wise minimum of the two price series  $\tilde{b}_t = \min\{p_t^f, p_t^b\}$ .

$$\tilde{b}_t = \min\{p_t^f, p_t^b\}, \quad \forall \tau \leq t \leq T$$

#### 2. **Maximin**( $\gamma_2$ ) filter-smoother:

- Construct  $\gamma_2$  period forward cumulative min of the observed price series,  $\tilde{b}_t^f$ .

$$\tilde{b}_t^f = \min\{\tilde{b}_t, \tilde{b}_{t-1}, \dots, \tilde{b}_{t-(\gamma_2-1)}\}, \quad \forall \tau + (\gamma_2 - 1) \leq t \leq T$$

- Construct  $\gamma_2$  period backward cumulative min of the observed price series,  $\tilde{b}_t^b$ .

$$\tilde{b}_t^b = \min\{\tilde{b}_t, \tilde{b}_{t+1}, \dots, \tilde{b}_{t+(\gamma_2-1)}\}, \quad \forall \tau \leq t \leq T + (\gamma_2 - 1)$$

- Take period-wise maximum of the two price series  $b_t$ .

$$b_t = \max\{\tilde{b}_t^f, \tilde{b}_t^b\}, \quad \forall \tau \leq t \leq T$$

The resulting price series,  $b_t$ , is the base price for the algorithm with hyperparameters  $(\gamma_1, \gamma_2)$ .

## B.3 Base price algorithm validity check: The Dominick's dataset

In this section, we evaluate how closely the base price algorithm mimics the promotion practices of retailers. To provide such an assessment, we leverage the promotional activities reported in an alternative scanner dataset of the now defunct grocery retailer, Dominick's, that operated about 100 stores in the Chicagoland area until 2013.<sup>27</sup> Much like

<sup>27</sup>This dataset is also made available by the Kilts Center for Marketing at the University of Chicago, and was used in the analysis of Chevalier et al. (2003) and Nevo and Hatzitaskos (2005). For more information, see <https://www.chicagobooth.edu/research/kilts/datasets/dominicks>.

Table B.3: Summary statistics for Dominick’s price of Campbell’s Chicken Noodle Soup

Variable	Mean	Overall std.	Within std.
Price (\$/can)	0.57	0.096	0.094
Simple discount	0.08	0.281	0.280
Impute	0.03	0.180	0.168

Source: Kilts Marketing Center University of Chicago, Dominick’s Data. Notes: The table shows summary statistics for the 31,567 store-week observations of prices and promotional activities for the 10.75 oz can of Campbell’s Chicken Noodle Soup. The weekly price observations come from 85 stores over an average of 372 weeks (per store).

the Nielsen dataset, the Dominick’s data contains weekly information on the posted price and quantity sold for each product in each store in each week. In contrast to the Nielsen data, the Dominick’s scanner dataset also reports for each product-store-week whether or not the product was on promotion that week. In particular, the dataset provides promotion in three measures—“Bonus Buy,” “Coupon,” or “Simple Price Reduction.” Although we have discovered that sometimes promotional activities are under- or miss-reported in Dominick’s scanner data, this dataset offers a unique opportunity to validate our base price algorithm in a grocery retail environment.

In preliminary analysis, the canned soup module displayed the most variation at seasonal frequencies. Consequently, for this exercise, we focus on one of the most popular items in the canned soup category, Campbell’s Chicken Noodle Soup. The Dominick’s dataset never reports the use of “Coupon” promotional activity. Furthermore, preliminary regressions reveal that “Bonus Buy” promotional activity is not closely linked to changes in per unit prices. In fact, in a regression of the logarithm of the per unit price on indicator for “Bonus Buy” and “Simple Price Reduction,” together with a set of month-price tier fixed effects, the estimated size of the price effect of a product being on “Bonus Buy” in that store-week is less than the benchmark threshold (10 percent) used in our baseline analysis. We also strongly reject the null hypothesis that the effect on prices of the two promotional activities is the same. Given that the objective of our algorithm is to infer promotional activities through changes in prices, we focus our validation exercise in picking up “Simple Price Discounts.”

Table B.3 presents some stylized facts about the prices for Campbell’s Chicken Noodle Soup in the Dominick’s dataset. The sample consists of an unbalanced panel of 85 stores reporting prices and units sold from September 20, 1989 to May 7, 1997. During this time period the average price was 57 cents, and was put on a “simple price discount” in 8 percent of store-weeks, a modest decrease from the frequency of promotion that we found in our main analysis. In the aggregate time series, Campbell’s Chicken Noodle Soup was on discount in at least one Dominick’s store in 8 percent of weeks over the time period of our sample. Lastly, the extent of imputed prices due to a store-week not reporting any sales is minimal comprising 3 percent of store-weeks.

As was documented for retailers more generally by DellaVigna and Gentzkow (2019), Dominick’s also exhibited relatively homogenous pricing at the product level across its

Table B.4: Type I and type II error rates of base price algorithm

Hyperparamters	Discount		Discount Non-missing	
	Type I	Type II	Type I	Type II
$\gamma_1 = 4, \gamma_2 = 6$	0.10	0.04	0.09	0.04
$\gamma_1 = 5, \gamma_2 = 7$	0.10	0.04	0.10	0.04
$\gamma_1 = 6, \gamma_2 = 8$	0.10	0.04	0.09	0.04

Source: Kilts Marketing Center University of Chicago, Dominick’s Data. Notes: Type I and II error rates for identification of promotional prices of the 31,567 store-week observations of prices for the 10.75 oz can of Campbell’s Chicken Noodle Soup of the base price algorithm. The weekly price observations come from 85 stores over an average of 372 weeks (per store), together with the reported promotional activities reported by Dominick’s. Type I error rates are defined as the percent of weeks defined as a promo conditional on Dominick’s not reporting that store-week as having a promotion, while type II error rates are defined as the percent of weeks defined as non-promo week conditional on Dominick’s reporting that store-week as having a promotion.

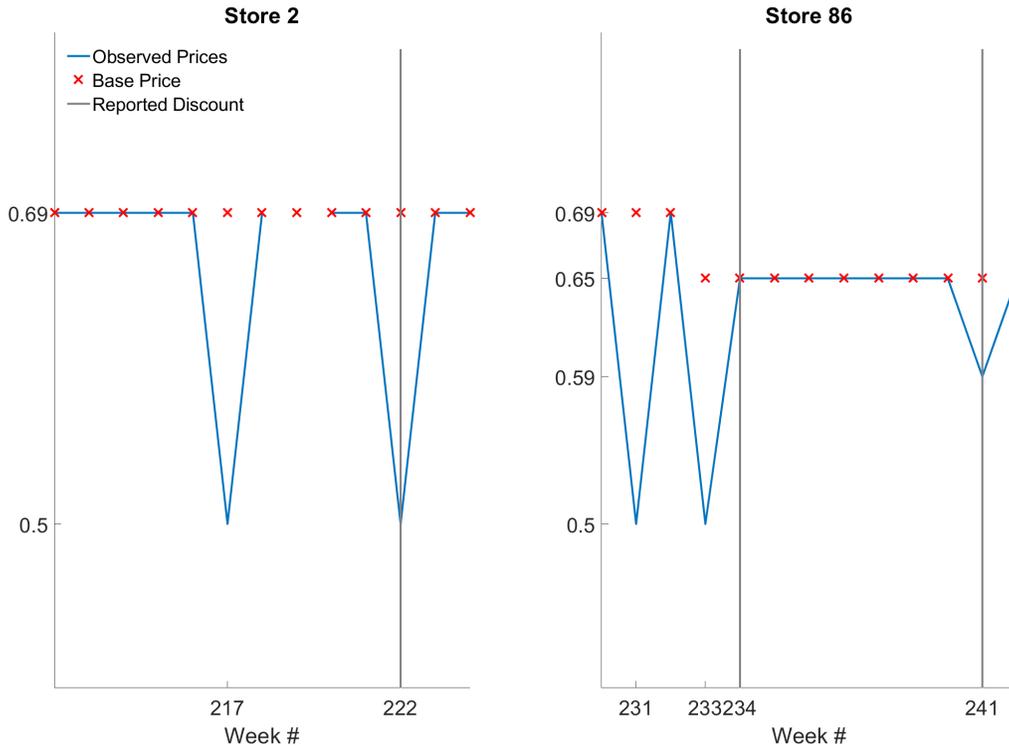
100 stores in any given week. To the extent that there are across-store differences in posted prices of a product for a given week, they are explained by the 12 price zones Dominick’s used to set their prices at four different price tiers. To this point, while store (price tier) fixed effects explain as little as 4 (2.6) percent of the variation in prices over the 85 stores and 400 weeks, a set of week-of-sample fixed effects explains as much 90 percent of the variation. We exploit the almost uniform pricing within a price tier to infer missing prices before we implement our algorithm.

The assessment of the algorithm is done by three steps. First, for store-weeks in which price is missing because no quantity is sold, we impute prices with the median price experienced in all the stores of that price tier. Second, we use our base price algorithm to impute base prices from the time series of weekly prices for a 10.75 ounce can of Campbell’s Chicken Noodle Soup (our modal product for the canned soup category) in each of the stores in the Dominick’s dataset. Third, we compare the implied promotional activities by the algorithm with those reported by Dominick’s for that product-store-week.

Table B.4 reports the type I and type II error rates of our base price algorithm for three sets of hyperparamters ( $\gamma_1, \gamma_2$ ). We define type I error rates as the percent of weeks defined as a promo by the algorithm conditional on Dominick’s not reporting that store-week as having a promotion. Type II error rates are defined as the percent of weeks defined as non-promo week by the algorithm conditional on Dominick’s reporting that store-week as having a promotion. The “Discount” columns report the error rates for all store-weeks, while in the last two columns we restrict the store-week observations to only those that did not require any price imputation. Due to the under- or false reporting from Dominick’s, the error rates are likely to be overestimated.

Several clear patterns emerge from table B.4. First, the performance of the algorithm appears relatively stable across the alternative choices of hyperparameters. The robustness of the algorithm to these hyperparameters is comforting, and was discussed in re-

Figure B.2: Dominick's case study examples



Source: Kilts Marketing Center University of Chicago, Dominick's Data. Notes: The figure reports the time series for two stores over several weeks of the observed price of a 10.75 oz can of Campbell's Chicken Noodle Soup (in blue), together with the base price (in red), as well as gray lines denoting the weeks in which Dominick's reported that a price discount was used in that store-week.

lation to our baseline sample in more detail earlier. Second, the type I and II error rates are modest, and given the likely measurement error in underreporting some promotional activity likely overestimates of the true false-positives and false-negatives. For instance, as soon as you restrict the evaluation to simple discounts and periods in which no prices had to be imputed, the maximum error rate (across both type I and II and sets of hyper-parameters) is 10 percent. As we will show next, it's highly likely some of these errors of the algorithm are due to measurement error in Dominick's data.

To provide some casual evidence to support our view that the error rates reported in table B.4 are conservative, we also report some particular instances of the observed time series of prices as well as the imputed base price for a couple of stores. In particular, figure B.2 reports the time series for two stores (stores 2 and 86) over several weeks of the sample, as well as the promotional activity reported by Dominick's over those weeks.

Several instances of false reporting on the part of Dominick's are evident from this figure. In particular, a price discount in week 217 for store 2, and week 231 and 233 for store 86, seems is not reported by Dominick's. Contrast those periods, with the reported price discount that Dominick's reported as having occurred in week 234 in store 86. Of course,

if the \$0.65 price in week 234 was in fact a price discount, it seems difficult to reconcile why the remaining weeks of that month would not also have been considered a price discount. These patterns indicate that measurement error may exist in the Dominick’s reporting of promotional activities, resulting in inflated type I and II error rates of our base price algorithm. In this light, we view the stable and strong performance of the base price algorithm in this sample of retailer behavior as providing substantial support for its use more systematically in the Nielsen Retail Scanner dataset used in the baseline analysis.

## C Empirical Bayes Shrinkage Procedure Algorithm

Our measures of monthly shifts in quantities ( $\alpha_m$ ), elasticities ( $\eta_m$ ), and prices ( $\beta_m$ ) are estimated with sampling error. If uncorrected, this sampling error means that we overstate the min-max difference and therefore end up with biased estimates of seasonality. We use an Empirical Bayes Shrinkage to adjust for this sampling error. The procedure we use goes back to the economics of education literature (where it is used to adjust differences in teacher- or school-specific value-added measures, see e.g. Kane and Staiger (2008); Jacob and Lefgren (2008); Angrist et al. (2017)). This approach is also used by DellaVigna and Gentzkow (2019) for their store-level elasticity estimates; we follow their implementation and description closely.

Let  $\theta_{im}$  denote a generic parameter estimate ( $\alpha_{im}$ ,  $\eta_{im}$ , or  $\beta_{im}$ ) for product  $i$  and month  $m$ . We define the shrunk estimate  $\tilde{\theta}_{mi}$  as

$$\tilde{\theta}_{mi} = \left( \frac{\sigma_{mi}^2}{\sigma_{mi}^2 + Var(e_{mi})} \right) \hat{\theta}_m^i + \left( \frac{Var(e_{mi})}{\sigma_i^2 + Var(e_{mi})} \right) \bar{\theta}^i,$$

where  $\bar{\theta}^i$  and  $\sigma_i^2$  are the prior mean and variance (at the category level), and  $e_{jce}$  is the estimation error in  $\hat{\rho}_{jce}$ . We define  $Var(e_{jce})$  as the estimate of the asymptotic variance of the pass-through, from Equation 1 or 5. We set  $\bar{\theta}_i = 0$  (which is exactly the sample average of  $\hat{\theta}_{im}$ ). We measure  $\sigma_i^2$  as  $Var(\hat{\theta}_i)$  minus the average estimation variance for that product (or zero, if that is larger). This procedure shrinks each product’s month-specific deviation from their annual average towards zero, i.e. the null of zero seasonality.

## D Lost Profit Calculations

This section briefly outlines our approach to calculating the lost profits reported in section Section 4. Our approach, which follows DellaVigna and Gentzkow (2019) closely, proceeds in two steps. First, we back out an estimate of marginal cost at the product-store-year level assuming that the average price is set optimally. Second, we calculate profits with optimally time-varying prices, with constant prices, and with prices at the observed path.

**Recovering marginal costs** We assume that the demand for product  $i$  in store  $j$ , week  $t$ , year  $y$ , and month  $m$  is given by  $Q_{ijt} = D_{ijt} P_{ijt}^{\eta_i + \eta_{im}}$ , where  $\eta_i + \eta_{im}$  corresponds to the price elasticity of demand for product  $i$  in month  $m$ . This specification parallels our

empirical approach: the price elasticity  $\eta_i + \eta_{im}$  varies at the product-month level but not more finely; and,  $D_{ijt}$  is an arbitrary demand shifter. We further assume that marginal cost for a given product-store-year is constant and given by  $c_{ijy}$ .

To recover marginal costs, we assume that stores price correctly on average, choosing the average price to maximize variable profits. That is, let  $\bar{P}_{ijy}$  be the average price charged for product  $i$ , by store  $j$ , in year  $y$ . We assume that:

$$\bar{P}_{ijy} = \operatorname{argmax}_P \sum_{t \in y} (P - c_{ijy}) D_{ijt} P^{\eta_i + \eta_{im}}. \quad (\text{D.10})$$

Rearranging the first-order condition for  $P$ , we obtain an expression for marginal costs:

$$c_{ijy} = \sum_{t \in y} \left[ (1 + \eta_i + \eta_{im}) D_{ijt} \bar{P}_{ijy}^{\eta_i + \eta_{im}} \right] / \sum_{t \in y} \left[ (\eta_i + \eta_{im}) D_{ijt} \bar{P}_{ijy}^{\eta_i + \eta_{im} - 1} \right].$$

To back out  $c_{ijy}$  and calculate profits, we require estimates of  $\eta_i + \eta_{im}$  and  $D_{ijt}$ , together with each product-store-year's average price  $\bar{P}_{ijy}$ . For  $\eta_i + \eta_{im}$ , we take the monthly elasticity implied by the (shrunk) estimates of Equation 1, (i.e.,  $\hat{\eta}_i + \tilde{\eta}_{im}$ ). Then, our estimate of the demand shifter is given by  $\hat{D}_{ijt} = Q_{ijt} / P_{ijt}^{\hat{\eta}_i + \tilde{\eta}_{im}}$ .<sup>28</sup>

**Profits under alternative pricing rules** We consider profits under optimal prices, as well as under the observed price path and under constant prices. Optimal prices are given by maximizing profits week-by-week, and adjusting prices to accommodate the fluctuations in the price elasticity of demand over the year. We use the average price at the product-store-year level as our measure of constant prices to be consistent with our estimate of marginal costs. To obtain the observed price path, we let prices fluctuate throughout the year around the constant level, with fluctuations given by our estimates of the seasonal fluctuations of prices. Specifically, along the "observed" path, we set the price in week  $t$  for product  $i$  store  $j$  and year  $y$  as

$$\tilde{P}_{ijt}^{obs} = \bar{P}_{ijy} \exp \{ K_{ijy} \} \exp \{ \tilde{\beta}_{im} \},$$

where  $\tilde{\beta}_{im}$  are the (shrunk) estimated coefficients on the month-of-year effects from Equation 5 (when log price is the dependent variable), and  $K_{ijy}$  is a product-store-year specific constant that ensures the average of the seasonally fluctuating price equals  $\bar{P}_{ijy}$ . This expression takes the average price and moves it up or down during the year according to observed seasonal movements outlined in our previous results.

Each chain offers many more products than those found in our sample. To provide a sense of the overall magnitude our lost profit estimates might have on chain level decisions, we scale up our product-level estimates to the aggregate chain level. To scale up our estimates, we assume that the profit lost from seasonal miss-pricing is a constant share of revenue across all products within a chain, and use the share of revenue coming from our sample of products as the basis for the scaling up our lost profit estimates. Panel

<sup>28</sup>This means that in store-weeks with zero sales,  $\hat{D}_{ijt} = 0$ , and price changes do not affect quantities. We do not believe this is important for our results because these store-weeks would be predicted to have very low quantities, and hence very low profits, under any alternative approach.

B of Table 3 reports the results of this calculation, for the median chain-year. Although this procedure is crude, we believe that it provides a conservative upper bound on profits lost from failing to more accurately time price deviations with seasonal fluctuations in the price elasticity of demand, because our sample of products represent a diverse and wide set products that span both ex-ante highly seasonal demand products as well as products that turned out not to be, as well as the most widely available products.