

Small Price Responses to Large Demand Shocks*

Etienne Gagnon
Federal Reserve Board

David López-Salido
Federal Reserve Board

January 26, 2015

Abstract

We study the pricing response of U.S. supermarkets to large demand shocks triggered by labor conflicts, mass population relocation, and shopping sprees around major snowstorms and hurricanes. Our focus on demand shocks is novel in the empirical literature that uses large datasets of individual data to bridge micro price behavior and aggregate price dynamics. We find that large swings in demand have, at best, modest effects on the level of retail prices, consistent with flat short- to medium-term supply curves. This finding holds even when shocks are highly persistent and even though stores adjust prices frequently. We also uncover evidence of tit-for-tat behavior by which retailers with radically different demand shocks nonetheless seek to match their local competitors' pricing movements and recourse to sales and promotions.

JEL classification: E30, L11.

Keywords: Demand shocks, inflation, sales, labor conflicts, mass population displacement, severe weather events.

*The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or any other person associated with the Federal Reserve System. We are grateful to Information Resources Inc. (IRI) for providing the scanner data. All estimates and analyses in this paper based on IRI data are by the authors and not by IRI. We thank Jason Sockin for providing superb research assistance, and Laura Ardila, Andrew B. Giffin, and Teyanna Munyan for further assistance with the data. We also thank Charlie Gilbert for assisting us with the weekly seasonal filtering procedure. Comments and suggestions can be directed to etienne.gagnon@frb.gov and david.lopez-salido@frb.gov.

1 Introduction

The effect of movements in aggregate demand on prices and quantities remains a lingering question in macroeconomics. A key reason is that prices and quantities are simultaneously determined; analyses based on correlations between these two variables may thus suffer from biases in coefficients and not be used to support causal claims. Ultimately, of course, it is the causal relationship that most interests scholars.¹

In this paper, we document a number of instances in which large demand shocks can be identified and show that, in those instances, the level of prices and retailing strategies more broadly react proportionally little. In particular, we study the pricing response of U.S. supermarkets to large demand shocks—that is, shocks that either reallocate customers across retailers or that alter shoppers’ willingness to consume—brought about by labor conflicts, mass population relocation, and shopping sprees caused by severe weather events. As we shall argue, we see the occurrence of our shocks as exogenous to retailers’ pricing strategies and supply factors, allowing us to use them as instruments to assess the causal effect of broad-based variation in demand on prices.²

We perform our analysis using a large dataset of weekly scanner data made available to researchers by Information Resources Inc. (IRI). This dataset contains price and quantity information from 2001 through 2011 on 29 personal care, housekeeping, and food product categories, and is well suited for our endeavor for several reasons. First, it covers 50 U.S. metropolitan markets, making possible the study of shocks that had a large impact on demand at a regional level but a limited influence on national aggregates. Second, its exceptionally large size—about 4 million individual observations per week after our trimming—permits the computation of reasonably accurate statistics at the product category–market level; this feat would not be achievable using the considerably smaller micro datasets collected for the computation of the official CPI and national account statistics. Third, it contains the universe of items sold by stores within each product category, thus allowing us to track the evolution of spending on these categories. In conjunction with the price data, the spending data permit us to derive real (same-store) quantity measures using a methodology similar to that employed for the national accounts. Fourth, the weekly data frequency allows us to zoom in on shocks that have a short life-span. Fifth, the availability of data from multiple retail chains and stores permits us to contrast the pricing behavior of local competitors facing differing shocks to demand.

Our most crucial task is to identify events that can broadly move demand but that are otherwise exogenous to supply factors. We first consider two labor conflicts that both began in October 2003.

¹The literature on the identification of aggregate demand shocks and their effects is too vast for us to offer a comprehensive review in these pages. See Shea (1993) for a discussion of identification issues along with evidence of manufacturing demand shocks derived from shifts in input-output linkages. For monetary policy shocks, see the complementary approaches in Christiano, Eichenbaum, and Evans (1999), Romer and Romer (2004), and Bernanke, Boivin, and Elias (2005). For fiscal policy shocks, see Hall’s (2009) survey. For evidence on the aggregate demand effects of uncertainty shocks, see Leduc and Liu (2013).

²One could argue that both labor conflicts and natural disasters constitute variation in the treatment—that is, the variable (changes in demand) whose causal effect we would like to understand—as if they were generated by a random experiment (see, for instance, Angrist and Imbens (1995)).

The conflict in St. Louis, MO, was relatively short (less than a month) while the conflict in Southern California dragged on nearly five months to become the longest supermarket strike and lockout in U.S. history. Affected stores remained open during both conflicts. Because many shoppers would rather take their business elsewhere than cross picket lines, strikes and lockouts can profoundly reallocate store traffic within a market. Many stores in our sample saw their revenues collapse 50 percent or more for the duration of these conflicts while others experienced correspondingly large increases. Stores whose employees were not on strike or locked-out faced no supply constraints, making the rise in their demand akin to an exogenous demand shock. At the end of the conflicts, staff numbers normalized but store traffic did not always revert to pre-conflict levels, providing additional opportunities to study the effects of variation in demand.

We next consider the mass population displacement brought about by Hurricane Katrina, the most expensive natural disaster in U.S. history. The hurricane displaced about 1 million persons, many of whom took years to resettle. Most displaced households moved in with relatives or friends, boosting population density and store traffic in neighborhoods that were less affected by the storm. Stores in our sample from the New Orleans, Louisiana, market and the Mississippi market experienced a persistent rise in sales volumes of about 20 percent, on average, in the wake of the hurricane.

Finally, we consider shopping sprees around major snowstorms and hurricanes not associated with mass population displacement. Like Hurricane Katrina, their occurrence is unquestionably exogenous to retailing activities. But contrary to Hurricane Katrina, the typical effects of these storms on demand are short-lived and operate primarily through a rise in the demand of existing shoppers rather than through a reshuffling of the consumer base. Storms that result in the closing of schools and workplaces force households to consume a larger fraction of their meals at home, thus boosting demand for food items. Similarly, the demand for personal care and housekeeping products may rise as households engage in more home production or take advantage of their trip to the supermarket to purchase items other than food.

Our key finding is that large swings in demand appear to have, at best, a modest effect on the level of prices, consistent with a flat short- to medium-term supply curve in the retail industry (or, equivalently, consistent with the price elasticity of supply being high). This finding holds even in the case of our most persistent shocks for which stores adjusted the price of most items at least a couple of times. Put differently, the lack of a significant price response to large swings in demand seems inconsistent with the marginal cost of retailers being sensitive to the level of demand because of fixed factors of production (as in, for example, Burstein and Hellwig (2007)), or to retailers making large alterations to their markups in response to changes in demand. Our evidence is further consistent with tit-for-tat behavior in which retailers with radically different demand shocks nonetheless seek to match their local competitors' pricing movements and recourse to sales and promotions.

Our two labor conflicts feature the most dramatic shocks to demand in our study. For the Southern California strike and lockout, we split establishments into two broad groups: those that

saw a drop in revenues in excess of 10 percent and those that saw an increase in excess of 10 percent (henceforth generically referred to as “on strike” and “not on strike,” respectively). The group on strike saw quantities sold drop by half relative to the pre-strike period while the group not on strike saw quantities increase by a third. Despite this remarkable difference, the level of prices was similarly behaved between the two groups throughout the conflict, rising about two percent in the first couple months and erasing half of that increase in the remaining months. These price movements are not especially large relative to the historical volatility of the sample. Importantly, they are modest relative to quantity developments. Perhaps remarkably given limited staffing numbers, stores on strike put almost the same number of items on sale during the conflict, supporting the view that engaging in price discrimination is a central activity of retailers.

At the end of the Southern California conflict, stores that had been on strike recouped only 4 out of every 5 dollars in lost business. To regain market share, these stores increased the number of sales and promotions, a move that was largely mimicked by their local competitors. This strategy gradually brought customers back even though relatively aggressive discounts did not translate into lower transaction prices on average. This finding points to marketing having an influence on sales volumes that goes beyond what can be inferred from movements in price indexes alone.

The shorter St. Louis conflict had similarly large effects on quantities at stores on strike and stores not on strike than the longer Southern California conflict. However, the end of the strike brought an immediate return to pre-conflict activity levels even though the strike had arguably lasted long enough to induce consumers deterred by picket lines to shop at unaffected establishments. Perhaps as a result, stores that had been on strike did not launch a sales and promotion campaign in the wake of the conflict. Overall, these findings suggest that there is some stickiness in consumer preferences for particular points of purchase, that those preferences are not eroded by shopping elsewhere for a few weeks, but that longer displacements can lead to permanent changes.

We break down the effects of Hurricane Katrina on retail activities in two phases. In the week of the storm, store revenues in our New Orleans sample jumped by nearly a quarter from their average of the previous couple months. The jump was especially sizable for a number of personal care products such as toothbrushes, a sad reminder of the personal tragedies endured by many displaced families. The second phase extends from the weeks after the storm, when store revenues were about 20 percent above average, to well over a year later as displaced households slowly returned home or settled elsewhere. We find little if any evidence that retailers responded to the persistent increase in demand by raising prices. Although food prices rose a little more in New Orleans than the corresponding national average in the months after the disaster, the price of personal care and housekeeping products rose by somewhat less. Moreover, we cannot exclude the possibility that supply disruptions created some upward pressure on food prices in New Orleans, further reducing the contribution of higher demand to the observed food price increase.

Our last set of demand shocks—shopping sprees triggered by major snowstorms and hurricanes—has no apparent effect on prices and other elements of our retailers’ pricing strategies. This finding is perhaps unsurprising because weather forecasts suffer from considerable uncertainty even for

horizons as short as a few days. Moreover, the effect of these shocks on demand rapidly dissipates. The fleeting character of these shocks may simply leave retailers insufficient time to adjust prices accordingly. Nonetheless, they invite us to qualify earlier findings that retail prices tend to fall around periods of peak demand. Warner and Barsky (1995), MacDonald (2000), Chevalier, Kashyap, and Rossi (2003) report that prices of household appliances and some food items tend to decline around major holidays, which are also short-lived peaks in demand.³ This evidence, which is based on the behavior of narrow sets of products, is sometimes seen as suggestive of counter-cyclical markups at the macroeconomic level. Our conjecture is that the high predictability of demand peaks due to holidays or the passing of the seasons allows retailers to adjust their marketing strategies accordingly, making these events somewhat uninformative about how retailers might respond to other high-demand episodes.

To our knowledge, our focus on broad-based demand shocks is novel in the empirical literature that uses large datasets of individual data to bridge micro price behavior and aggregate price dynamics. Perhaps one exception is Coibion, Gorodnichenko, and Hong (2012), who look at the link between local labor market conditions and pricing. Otherwise, previous work has considered a number of supply-related factors such as changes in sales taxes (see, for example, Dhyne *et al.* (2005) and the references therein, Karadi and Reiff (2012), and Gagnon, López-Salido, and Vincent (2013)), imported good price shocks (for example, Gopinath and Itskhoki (2010)), and commodity price shocks (for example, Nakamura (2008) and Hong and Li (2013)). This literature has typically found moderate pass-through rates that contrast with the muted price responses to our shocks. Our use of natural disasters to identify exogenous changes in demand is also somewhat unusual as the literature has typically been interested in their disruptive effects on supply. In a related paper, Cavallo, Cavallo, and Rigobon (2013) use online price data to show that a 2010 earthquake in Chile and the 2011 Sendai earthquake and tsunami in Japan both led to widespread product unavailability but not to higher prices.

The paper is organized as follows. Section 2 details the construction of our dataset and presents key features of retail activities in normal times. Section 3 analyzes customer base displacement due to labor conflicts. Section 4 investigates shifts in demand due to major weather events, starting with the year-long population displacement caused by Hurricane Katrina and proceeding with transitory boosts to demand related to major snowstorms and hurricanes. Section 5 discusses what factors might contribute to retailers' muted price response to our large demand shocks. Section 6 concludes.

2 Dataset and some salient pricing facts in normal times

Our price and quantity analysis is performed using weekly scanner data made available to researchers for a nominal fee by IRI. The content of the dataset is detailed in Kruger and Pagni

³This finding also appears to apply at the product-category level in our IRI sample. For most product categories, we find a negative correlation between movements in our price and quantity indexes after removing variations at frequencies higher than 53 weeks.

(2012) and Bronnenberg, Kruger, and Mela (2008), so we shall give only a brief exposition. The data come from a large sample of over 1,500 U.S. supermarket stores belonging to a variety of retail chains and operating in 50 U.S. markets. An observation corresponds to the information about an item (that is, a barcode sold in a particular store) in a given week. The number of observations is exceptionally large at about 300 million per year. Available item information includes the number of units sold and total revenue during the week. As is customary with scanner data, we derive a unit price by dividing total revenue by the number of units sold. Although the data are limited to 29 food, housekeeping, and personal care product categories, they have the appealing feature of encompassing all item transactions within those product categories.⁴

To make the micro data suitable for our purposes, we drop items with suspiciously large price adjustments and apply various filters to extract regular price series and compute various moments of temporary sales. Our regular price filter is based on that proposed by Nakamura and Steinsson (2008).⁵ We then construct price and quantity indexes for each IRI market-product category combination using a methodology that closely matches that employed by the Bureau of Labor Statistics (BLS) and the Bureau of Economic Analysis (BEA) for the U.S. CPI and the U.S. National Income and Product Accounts, respectively. Our sample-wide statistics aggregate product-category statistics by weekly revenues. We relegate these technical details to appendix A.

Before proceeding with our key findings, we briefly describe the general behavior of retail prices in the dataset to provide us with some benchmark against which to compare pricing behavior in the presence of large demand shocks. Appendix B provides a more detailed presentation of salient pricing facts in normal times. The key message is threefold: retail prices are adjusted frequently, price adjustments are primarily associated with sales and promotions, and the bulk of barcode-level price variation is common across stores belonging to the same retail chain.

In particular, on average, 30.4 percent of items in our sample experience a price change every week.⁶ On a monthly basis, 42.3 percent of items experience a price change, a figure somewhat above corresponding estimates for the broader U.S. CPI, which range from 26 percent in Bils and Klenow (2004) and Nakamura and Steinsson (2008), to 36 percent in Klenow and Kryvtsov (2008). Regular price changes identified by our filter are less common, at 11.2 percent per week and 23.7 percent per month. Moreover, price behavior is consistent with pricing decisions originating principally at the retail chain level. To establish this fact, we define a weekly “chain price” as the modal price at which a barcode is sold across stores belonging to the same retail chain within a market. Overall, 71.2 percent of store prices equal their corresponding chain price. In addition, the

⁴The IRI dataset also has information on cigarettes and photographic supplies; we exclude these product categories because of regulatory restrictions on pricing and their gradual obsolescence, respectively.

⁵The identification of regular prices based on store-reported weekly average transaction prices is subject to some degree of imprecision. As appendix A discusses, we experimented with a number of procedures to separate low- and high-frequency price movements. These procedures include a reference price filter, a V-shape sales filter, and the use of a 5-percent-or-more price reduction flag created by IRI. Often, filters that produce relatively infrequent regular price adjustments are associated with relatively frequent but shallow sales, possibly as they pick up transitory price movements related to the use of retailer coupons or price adjustments during the week. Although no method seems fully satisfactory, we believe that our broad conclusions about the effect of large demand shocks on prices are robust.

⁶Our sample-wide statistics aggregate product-category statistics by weekly revenues.

probability of an item experiencing a price change jumps to over 80 percent when we condition on a change in the item’s chain price. Together, these facts provide clear indications that item prices are coordinated to some degree across stores belonging to the same retail chain.

3 Labor conflicts as exogenous demand shocks

Labor conflicts can have substantial effects on store traffic because many shoppers would rather take their business elsewhere than cross employee picket lines. Thus, affected stores can lose most of their demand whereas unaffected stores may experience a jump in sales volume. This section looks at such rearrangements of the customer base during two major labor conflicts in the supermarket industry. The key finding is that pricing strategies at stores whose demand rose largely mimicked those at stores whose demand declined. This finding is consistent with the effects of radically different demand shocks on pricing being trumped by a desire to keep up with pricing movements by local competitors.

Because labor conflicts are man-made and affect staffing levels, it is appropriate to take a moment to discuss their exogeneity with respect to pricing strategies and their likely effects on stores’ ability to supply goods. For stores whose employees are not on strike or locked-out, there were no disruption to staffing numbers. The observed rise in their sales volumes was thus akin to an exogenous demand shock, supporting our use of the strike as an instrument. The normalization of staffing levels at the end of our conflicts provides other opportunities to study the effects of variation in demand in the absence of constraints on supply because store traffic did not always revert back to pre-crisis levels.

In the case of stores whose employees were on strike or locked-out, the ability to meet any given level of demand was certainly hindered by low staffing levels over the duration of the conflicts. As we will see shortly, the conflict was also accompanied by a large and arguably more important adverse shock to demand as many shoppers opted not to cross picket lines and took their business elsewhere. Strikes and lockouts are often due to disagreements over employee compensation, which enters a grocers’ marginal cost and should thus influence prices. Due to the censoring of store and retail chain information, we cannot assess the importance of compensation for our establishments.⁷ However, we know from financial statements of publicly-traded U.S. supermarket chains that compensation accounts for a relatively small share of overall supermarket revenues. For instance, the largest two publicly-traded U.S. retail chains with activities concentrated in supermarket products, Kroger and Safeway, reported goods acquisition costs equivalent to 77.5 percent of their combined revenues in 2012. By contrast, operating and administrative expenses, which include employee compensation as well as spending on store management, utilities, local advertising, etc., accounted for only 18.3 percent of revenues. If these figures are reflective of the cost structure of establishments in our sample, then the difference in bargaining positions between employers and employees is unlikely to have exceeded a small fraction of overall costs. For this reason, we see labor disputes as ultimately

⁷The terms of our contract with IRI prevent us from reporting any information that could lead to the identification of particular retail chains or establishments.

creating only limited uncertainty regarding prices paid by consumers. That said, even small changes to employee compensation can have a large impact on store profitability because food retailing is a low-margin, high-volume businesses (operating profits at Kroger and Safeway averaged only 2.6 percent of their combined revenues in 2012).

3.1 2003-2004 Southern California supermarket strike and lockout

On October 11, 2003, about 70,000 unionized supermarket workers in Southern California either went on strike or were put on lockout due to a disagreement with several retail chains over benefits and compensation. The dispute ended nearly five months later, on February 29, 2004, when workers voted in favor of a negotiated contract, resolving what had become the longest supermarket strike in U.S. history. Stores whose employees were on strike or on lockout remained in operation throughout the conflict, in part as retail chains reallocated managerial resources to the aisles and pursued alternative arrangements to ensure continued product availability. News organizations reporting on the conflict indicated that shelves were generally fully stocked but that few shoppers were strolling the aisles.⁸ For this reason, we believe that supply constraints were not the primary driver of the large drop in quantities observed at stores on strike, although supply constraints certainly would have made it challenging if not impossible to satisfy pre-conflict sales volumes.

Our dataset contains establishments that experienced marked drops in revenues during the strike as well as establishments that experienced marked increases. As noted earlier, the identity of stores and retail chains is censored in the sample and, in compliance with our terms of usage, we make no attempt to uncover them. To further preserve the anonymity of establishments and retail chains involved, we pool all data from the Los Angeles and San Diego markets and look at two broad groups of stores: those whose revenues dropped more than 10 percent (generically labeled “on strike”) and those whose revenues rose more than 10 percent (generically labeled “not on strike”) during the strike relative to the corresponding period of the previous year. Because our identification strategy uses only observed movements in revenues but no retail chain information, the two groups may not perfectly overlap with the sets of stores that were actually affected and unaffected by the strike. That said, persistent movements in weekly revenues of 10 percent or more are exceptional in the dataset, and the large movements that we attribute to stores being on strike or benefiting from the strike accord with the chronology of the conflict. Some stores also saw revenues change by less than 10 percent; we ignore them to focus on stores subject to large demand shocks. We finally compute separate statistics for each group to contrast the situation of stores with large positive and large negative shocks to demand.

⁸For example, the New York Times reported on October 14 that “store shelves appeared fully stocked but the aisles were largely empty of shoppers.” As an explanation, the article cited a striking inventory control clerk contending that “picketing had turned away about half of the market’s customers” while “the other half [was] getting a couple things.”

3.1.1 Overall impact on quantities and prices

Salient features of the strike’s impact on retailing activities are presented in figure 1. As the top panel shows, stores whose employees were on strike experienced, on average, a staggering 50 percent drop in sales volumes over the conflict’s duration. Some stores even saw sustained declines in quantities close to 80 percent. The drop was somewhat more pronounced in the first few weeks, suggesting that some shoppers who had initially declined to cross picket lines or to limit their purchases quickly returned. This modest rebound aside, the drop in revenues at establishments on strike was large and sustained for the duration of the labor conflict. In sharp contrast, establishments in our group of stores that benefited from the strike witnessed large increases in revenues—over 30 percent, on average, with some stores even seeing their revenues more than double. Contrary to stores on strike, those that benefitted continued to have their regular employees present to serve consumers, replenish the shelves, and move products from warehouses to stores. In the absence of disruptions to their ability to supply products, the sudden rise in demand for these stores is thus unambiguously interpretable as a demand shock.

The top panel of figure 1 further suggests that a majority of consumers displaced by the strike returned to their previous shopping location at the end of the conflict: Sales volumes immediately rebounded at stores that had been on strike and fell sharply at stores that had benefited from the strike. The normalization was incomplete, however. In the year that followed the end of the conflict, sales volumes at stores that had been on strike were still 10 percent below sales volumes before the strike, whereas stores that had benefited from the strike retained some of the customers displaced by the strike. This fact could be consistent with theories that emphasize consumer loyalty to brand, stores, or retail chains.⁹ The immediate return of sales volumes toward their pre-conflict levels suggests that consumer preferences for points of purchase may persist even after consumers have switched stores for nearly five months. It is also possible that other factors were at play. For instance, stores with significantly higher sales volumes during the strike may have been able to offer fresher produce as a result. Such benefits may, in turn, have helped them retain consumers at the end of the conflict.

As the middle panel shows, price movements during the conflict, at only a couple of percentage points, were more than an order of magnitude smaller than swings in sales volumes, suggestive of flat supply curves. For example, if we attribute the observed rise in prices at stores not on strike entirely to a positive demand shock, then the estimates in table 1 imply a supply elasticity equal to $\log(1.021) / \log(1.336) = 0.07$. This modest estimate falls slightly if we control for inflation over the strike period by measuring the demand shock’s effect as the rise in excess of that for the price index of the full IRI sample.¹⁰ Such small responses appear consistent with surveys of price-setting

⁹Most such models are specified in terms of loyalty to specific items or brands rather than to establishments or retail chains as seems appropriate in our context. These models include switching costs (for example, Klemperer (1995) and Kleshchelski and Vincent (2009)) and deep habits (Ravn, Schmitt-Grohé, and Uribe (2006)).

¹⁰We could further measure the supply elasticity by comparing relative price movements to relative quantity movements between stores on strike and stores not on strike, with the caveat that stores on strike may have suffered from supply disruptions in addition to a relative demand shock. The resulting elasticity is essentially zero at $\log(1.336/0.512) / \log(1.021/1.013) = 0.01$.

practices in which firms report being relatively unresponsive to positive shocks to demand conditions (see Fabiani *et al.* (2006)). These supply elasticity estimates for the retail sector are markedly lower than the 0.18 figure reported by Shea (1993) for the U.S. manufacturing sector. One explanation for the difference could be that Shea’s (1993) estimate applies to a one-year horizon whereas the Southern California strike ended after five months, leaving less time for prices to adjust to higher demand. However, our evidence for the longer post-strike period, which features persistently large differences in demand and no supply disruptions at both groups of stores, also points to little if any price response.¹¹

It is also apparent that price movements were similar between stores on strike and stores not on strike before, during, and after the conflict. After rising a couple percent in the spring and summer of 2003, prices declined a little in the months prior to the strike, then rose over 3 percent in the first half of the conflict before retracing half of that rise in the second half. Overall, divergences in price movements between the two groups of stores were short-lived and not exceptionally large in comparison to other relative price movements over our sample period. After the strike, constraints on stores’ ability to adjust prices and to supply goods vanished. Fair pricing motives that could have made retailers reluctant to boost prices amid exceptionally high demand would have greatly eased now that customers could shop freely at all stores. Despite some persistent differences in the quantity indexes, our price indexes show little economically meaningful divergence between the two groups of stores.

To provide a more direct comparison of the level of prices between stores that saw their demand soar and stores that saw their demand collapse, we next look at the cost of purchasing identical baskets of goods. We consider three such baskets. The first basket (the “fixed” basket) consists of all barcodes continuously available at both groups of stores over the two-year period displayed in figure 1, that is, over a period starting 26 weeks before and ending 78 weeks after the beginning of the conflict. The number of units purchased for each barcode is set to the average weekly number of units sold across all stores over the two-year period. This fixed-quantity basket is reasonably representative of overall purchases in Southern California, accounting for nearly 70 percent of total revenues over the period. At the product category level, the coverage of the basket ranges from 27 percent of total revenues for razors to 95 percent for peanut butter. The second basket (the “on-strike” basket) corresponds to the number of units purchased at stores on strike, again using only barcodes that are continuously available at both groups of stores. Similarly, the third basket (the “not-on-strike” basket) consists of the number of units sold at stores not on strike. Contrary to the fixed basket, the composition of the on-strike and not-on-strike baskets varies from week to week in line with shoppers’ actual consumption.

The lower panel of figure 1 reports the cost of purchasing each of the three baskets at the mean transaction price observed at stores on strike relative to that at stores not on strike. (See appendix A for the exact formulas.) All three ratios tell the same story: The cost of purchasing any of our

¹¹When we measure the medium-run elasticity of supply using relative movements in prices and quantities from the period before to the period after the strike, we obtain a mildly negative slope estimate, $\log(1.009/1.012) / \log(1.036/0.904) = -0.02$.

baskets at stores on strike relative to stores not on strike hovered near its pre-strike level for the duration of the conflict. A slight increase in the relative price of the baskets at stores that were on strike is apparent several months after the end of the conflict. We are reluctant to attribute this rise in the ratios to a price response to persistently lower demand given the historical variability of the sample.

We note that shoppers at stores on strike could have purchased identical baskets of goods for an equal or even lower price at stores that were not on strike at any point over our two-year period, a finding that suggests some insensitivity to the level of prices on the part of consumers. This conclusion comes with a number of caveats. Because the identity of stores and retail chains is censored, we cannot control for differences in factors such as income and sales taxes that may affect the level of prices across areas. Also, our baskets comprise solely barcodes that are simultaneously available in both groups of stores; by ignoring roughly 30 percent of store revenues in our product categories, we may be overlooking the effect of private labels on the effective costs of a typical basket. Finally, retailers sell items in product categories that are not covered by the IRI sample.

3.1.2 Price adjustments, discounts, and the labor conflict

Promotional discounting is a central feature of supermarket retailing (see, for example, Pesendorfer (2002)), including for the product categories in our sample. Stores can adjust prices paid by consumers in several ways, notably by changing the regular price or by altering the frequency and depth of promotional sales. Figure 2 presents some key statistics on supermarkets' before, during, and after the strike. The mean of each statistic over these periods is reported in table 1. On average, stores on strike saw their weekly frequency of price changes drop 4.3 percentage points during the conflict, with a reduction in regular price adjustments accounting for most of the drop. The fraction of items on sale (the middle-left panel of figure 2) declined 1.2 percentage point to 22.0 percent while the mean saving conditional on a sale (the middle-right panel) was unchanged. In short, stores on strike broadly managed to maintain the importance of sales during the conflict. Stores not on strike experienced a somewhat smaller drop in the frequency of price changes, 3.0 percentage points, that reflected small declines in both sales-related and regular price adjustments.

On the one hand, the continued use of sales by stores whose demand tumbled suggests that engaging in price discrimination is an important endeavor of retailers independently of their level of demand. Perhaps it also reflects a desire from stores on strike to project a business-as-usual image, or an implicit promise to offer bargain shoppers some opportunities to buy a portion of their basket at a discounted price every week. Indeed, it is conceivable that price-sensitive shoppers who continued to patronize stores on strike would have found it unfair to see their loyalty rewarded with higher prices. On the other hand, the reduction in regular price adjustments seems consistent with limited human resources hindering stores' ability to reprice. It is also possible that, given markedly lower sales volumes, stores on strike would have tolerated larger deviations of regular prices from their optimum because fixed repricing costs would have been spread over a smaller number of units

sold.¹² We also note that the level of individual prices continued to be coordinated within stores belonging to the same retail chain (the lower-right panel) despite variation across stores in the magnitude of the drop in demand and ability to use managerial staff to fill positions previously held by striking employees. In fact, the share of prices equal to the chain price edged up a couple of percentage points during the strike at both stores on strike and stores not on strike.

As noted above, stores on strike recouped only 4 out of every 5 dollars in lost business once the conflict ended. To win back customers, they increased the frequency and depth of sales in the ensuing year.¹³ The middle panels of figure 2 show that the share of discounted barcodes and the mean saving on the dollar, which captures both the frequency and depth of sales, both rose a couple of percentage points before slowly edging back.¹⁴

To further explore the strategic use of sales during and after the conflict, we break down discounts into 10-percentage-point bins, starting with discounts that are below or equal to 10 percent, then below or equal to 20 percent but greater than 10 percent, and so on. Figure 3 displays the contribution of discounts in each bin to total revenues (the top row of panels) and to total savings on the dollar (the lower row of panels).¹⁵ The upper-left panel shows that the contribution of items to total revenues generally was declining in the size of the discounts extended. Before the strike, discounts up to 10 percent accounted for almost 8 cents out of every dollar in revenues at stores that were subsequently affected by the conflict, while items offered at discounts in excess of 50 percent accounted for only about 2 cents out of every dollar in revenues. The lower-left panel shows that discounts in the 40 percent to 50 percent range, which were both substantial and frequent, allowed shoppers to lock in the largest savings. By contrast, very small and very large discounts contributed little to actual consumer savings because they were either too small or too infrequent to have a large impact. The distribution of contributions to total revenues and to total savings of discounts had similar shapes in the pre-conflict period at stores unaffected by the strike than at stores affected by the strike, although discounts of medium sizes played a somewhat lesser role in their overall marketing strategy.

During the strike, the number of items on sales declined only a little at affected stores during the

¹²See Golosov and Lucas (2007) for an illustration. The widening of the price inaction region at stores whose demand fell should, in addition to lowering the frequency of price adjustments, boost the absolute size of their price changes. We do find an large increase (4.4 percentage points) in the average size of regular price adjustments. However, a notable increase (1.5 percentage points) is also present at stores that saw their demand jump.

¹³In addition to the evidence presented in figure 2, these claims are supported by communications between the affected retailers and their investors. For example, one retailer mentioned in its 2004:Q2 earnings report that “post-strike sales [had] improved, but [had] not reached pre-strike levels,” and that it “[continued] to focus on rebuilding sales through promotional pricing [...]”

¹⁴We measure the mean saving on the dollar as $(\sum_j (p_{j,t}^{reg} - p_{j,t}) q_{j,t}) / (\sum_j p_{j,t}^{reg} q_{j,t})$, where $p_{j,t}^{reg}$ is the regular price of item j in week t identified by our sales filter.

¹⁵To construct figure 3, we first pool all weekly observations in each of our three time periods (before, during, and after the strike). For the top panels, we then weigh weekly observations by their share of total revenues (that is, $\omega_{i,t} = p_{i,t} q_{i,t} / (\sum_{t \in \mathcal{T}} \sum_j p_{j,t} q_{j,t})$, where $\mathcal{T}$ is the set of weeks in the period of interest). The area under the curve integrates to the share of total revenues accounted by items on sale. Similarly, for the lower panels, we weigh weekly observations by their contribution to total savings over the period (that is, $\omega_{i,t} = (p_{i,t}^{reg} - p_{i,t}) q_{i,t} / (\sum_{t \in \mathcal{T}} \sum_j p_{j,t}^{reg} q_{j,t})$). The area under the curve then integrates to the mean saving on the dollar.

strike. The upper-left panel of figure 2 shows that the importance of small and medium discounts for total revenues fell most while that of discounts in excess of 50 percent actually rose some. This shift toward deep discounts may have been part of a marketing strategy to lure shoppers to stores despite the strike, as heavy discounts make for good advertising. That said, the contribution to total savings of heavily discounted items was offset by the lesser importance of small and medium discounts, leaving the mean saving on the entire basket about unchanged. For stores that were not on strike, we also witness a decline in the importance of small and medium discounts during the conflict period but do not find a corresponding rise in the importance of very large discounts.

Once the strike was over, the contribution of small and medium discounts to total revenues and to total savings rose above its pre-strike level at both groups of stores. The importance of discounts in excess of 50 percent also rose at stores that had not been on strike, catching up with the similar increase documented at stores on strike during the conflict. In sum, the evidence is consistent with stores that had been on strike offering somewhat more frequent discounts—small, medium, and large—than usual to win back customers, and stores that had avoided the strike doing the same to retain their new business. Set against the background of differing swings in demand during the strike and beyond, these similarities in the recourse to sales suggest that retailers attach much value to matching changes in their local competitors’ pricing strategies.

3.2 A shorter labor conflict: the 2003 St. Louis grocery strike and lockout

The 2003 St. Louis grocery strike and lockout paints a broadly similar portrait as the longer conflict in Southern California: Pricing strategies of stores on strike and stores not on strike broadly resembled each other despite highly diverging demand. It started on October 7 when employees at several supermarket chains went on strike or were locked-out after voting down a tentative labor agreement. The dispute ended 24 days later when negotiating teams reached an agreement that proved acceptable to all parties.

The main effects of the conflict on prices and quantities are shown in figure 4. Due to the smaller size of the St. Louis sample (22 stores compared to 107 stores in Southern California), we compare establishments whose revenues dropped more than 10 percent (generically referred to as “on strike”) to all other establishments in the sample (generically referred to as “not on strike”). Picketing was again effective at deterring shopping activities, with stores on strike seeing sales volumes plunge about 45 percent, on average, while stores not on strike experienced a modest increase. Reading through the weekly volatility, prices at stores not on strike rose a touch more around the time of the conflict than prices at stores on strike. However, this slight divergence was part of a broader trend over the two-year period displayed and thus difficultly attributable to the strike. Moreover, price movements on a like-for-like basis were much more similar between the two groups of stores over the period displayed and around the strike in particular; our relative price measures controlling for the composition of the basket, which are shown at the bottom of figure 4 and employ the same methodology as those computed earlier for Southern California, were very stable.

Of note, sales volumes at stores on strike fully recovered as soon as the labor conflict was over, in contrast with the customer base erosion apparent at stores on strike in Southern California. Given limited household inventories, a 24-day conflict is arguably too long for customers unwilling to cross picket lines at their usual supermarket not to go shopping elsewhere. The recovery in sales volumes thus suggests some stickiness in consumer preferences for particular establishments that are not captured in standard models of store switching costs such as Kleshchelski and Vincent (2009), where store changes are permanent. In any case, the quick recovery in sales volumes probably explains why we did not observe a rise in the frequency and depth of discounts (shown in the middle panels of figure 5 and documented in table 2) in the wake of the St. Louis conflict.

4 Major weather events as exogenous demand shocks

Our next set of demand shocks were created by Mother Nature; their occurrence was unambiguously exogenous to retail activities in general and to supermarkets' pricing strategies in particular. We first look at Hurricane Katrina that, in the span of a few tragic weeks in the summer of 2005, led to massive population displacement. Stores located in areas that received an inflow of refugees experienced a sharp rise in store traffic that persisted for well over a year. We then look at shopping sprees triggered by snowstorms and hurricanes. Contrary to strikes and Hurricane Katrina, these storms do not feature a reconfiguration of retailers' customer base but rather a temporary increase in the demand of all customers.

4.1 Hurricane Katrina

In August 2005, Hurricane Katrina created an estimated \$108 billion in property damages (in 2005 dollars), making it the most expensive natural disaster in U.S. history. It was also directly responsible for the tragic loss of about 1,200 lives and the displacement of roughly 1 million persons.¹⁶ The city of New Orleans, Louisiana, sustained the most extensive damage due to the failure of its levee system, which led to the flooding of approximately 80 percent of the city. The flooding of residential areas made it impossible for many displaced households to return home for several months or years, and some households even chose to permanently relocate elsewhere. According to research conducted at the U.S. Census Bureau and reported in Geaghan (2011), as of late 2009, 31,500 households in the New Orleans metropolitan area (7 percent of the area's total) did not consider themselves permanently resettled.¹⁷

The hurricane had disruptive consequences on retail activities in New Orleans and other affected states. Of the 23 New Orleans stores that participated in the IRI sample on the eve of the tragedy,

¹⁶These estimates are taken from Blake and Gibney (2011).

¹⁷Three weeks after Hurricane Katrina, authorities in Texas ordered 1.8 million persons to evacuate coastal areas along the Gulf of Mexico ahead of Hurricane Rita's landfall. Transitory jumps in sales volumes are apparent for the main markets that hosted refugees. We exclude this storm from our analysis because the effects of mass population relocation cannot be disentangled from those of shopping sprees ahead of the storm by households sheltering in place. We will explore this latter phenomenon in section 4.2.

five exited the sample as soon as the storm hit while one store ceased to report data for a period of eight months. Similarly, five out of the nine stores in the Mississippi sample did not report data for a week or two around the hurricane. Although the IRI dataset contains no information that would permit us to ascertain that these sample exits and missing reports were caused by the hurricane, we interpret their coincidental timing as strongly suggestive that they were.

As figure 6 shows, the hurricane also affected retail activities at stores that remained in business. In the weeks and months that followed the disaster, stores that continued to report data to the IRI experienced sales volumes that were, on average, about 20 percent higher than before the disaster. We found an equally large rise in the smaller Mississippi sample (not shown); a smaller effect is also apparent for the Houston, Texas, sample (also not shown). We interpret these persistent rises as evidence that the hurricane displaced shopper traffic toward establishments located in areas that were relatively unaffected by the storm.¹⁸ Supporting our interpretation is the fact that, according to Geaghan (2011), 74 percent of displaced New Orleans householders reported living with an acquaintance.¹⁹ In addition, the persistent increase in sales volumes of food products as well as housekeeping and personal care products, which are shown separately in figure 6, were of a similar proportion, consistent with mass relocation boosting the demand of supermarkets' entire product offering. Indeed, we observe increases in all 29 product categories in the sample. The upper-left panel of figure 6 also features a short-lived but outsized 60-percent surge in sales volumes of housekeeping supplies and personal care products the precise week that Hurricane Katrina hit. Revenues from product categories such as toothbrushes and razors witnessed transitory increases in excess of 100 percent, again consistent with population displacement being a key driver of retail activities over the period.

Our empirical evidence provides little if any support to the view that retailers took advantage of higher demand brought about by the hurricane to raise prices, either initially when spending on personal care products skyrocketed or over the medium run when store traffic was boosted by mass relocation. The upper-right panel of figure 6 shows that the price of food products and of personal care and housekeeping products both rose in the weeks that followed the storm before erasing some of these gains. When we average over the period covered by the federal emergency declaration, we find that prices were 1.4 percent higher overall than they were over the 26-week period before the hurricane. As table 3 reports, this increase is modestly larger than the average rise across IRI markets not directly affected by Hurricane Katrina (that is, excluding New Orleans, Mississippi, and Houston). Muted price movements translate into small estimates of the short-term elasticity of supply. If we measure the price impact as the rise in the New Orleans price index in excess of

¹⁸Whether market-wide demand also rose is less clear, at least initially. On the one hand, the arrival of a hurricane typically boosts spending on a wide range of products, a phenomenon that we document in section 4.2. On the other hand, mass population exodus in the months and years after Hurricane Katrina arguably depressed market-wide demand, even if demand at continuing stores may have been persistently higher than in the past.

¹⁹The fact that most refugees sought shelter with family members and friends suggests a similarity in consumer characteristics between displaced and hosting households. See Lach's (2007) analysis of a mass migration from the former Soviet Union to Israel for evidence that large influx of consumers with price elasticities and search costs differing from the native population can affect retail prices.

the average rise for IRI markets not directly affected by Hurricane Katrina, then our estimate of the (short-run) elasticity of supply is 0.03.

When we compute the elasticity of supply using price and quantity movements over the longer period after the federal emergency relative to the pre-storm period, we get a supply elasticity estimate of 0.13. However, we are reluctant to interpret this medium-run estimate as suggestive that retailers took advantage of higher demand to boost prices for two main reasons. First, this estimate may be biased upward by hurricane-related disruptions to the region's food supply. Second, price pressure was not broad based, as the price of housekeeping and personal care products rose in line with the national average.

The remaining panels of figure 6 provide further results regarding the impact of the hurricane on retailers' broader pricing strategy. The frequency of price changes edged down during the federal emergency period, and reverted to its pre-hurricane average for about a year before sliding in the fall of 2006. The hurricane had a somewhat more apparent effect on the mean saving on the dollar, which temporarily slid from between 5 and 6 percent of the regular price in the weeks before the hurricane to a low near 2 percent, before rebounding at the end of the federal emergency declaration. Much of this decline reflects a lower proportion of items on sales, especially in late September and early October, rather than a shift in customer spending toward items with lesser or no discounts, as hinted by the relative stability of both the mean saving on the dollar and the share of items on sale in the first couple weeks after the storm.

One may suspect that the absence of significant upward movements in prices during or after Hurricane Katrina could reflect legal impediments to price adjustments. Indeed, on August 26, 2005, Louisiana Governor Kathleen Babineaux Blanco proclaimed a 30-day state of emergency that automatically triggered a law prohibiting price gouging. The law stated that "prices charged or value received for goods and services sold [...] may not exceed the prices ordinarily charged for comparable goods and services in the same market area at or immediately before the time of the state of emergency." These price restrictions were allowed to lapse on September 25, 2005, even though the state of emergency was systematically renewed through 2007. Figure 6 does not suggest that they were a material constraint on prices, as their expiration was not associated with unusual upward movements in prices. For the period covered by the restrictions, it is conceivable that retailers had some room to adjust prices upward by limiting the number and the depth of sales and promotions. Indeed, as noted above, we do observe some decline in the mean saving on the dollar and share of items on sales early after the tragedy. However, we are reluctant to directly attribute these movements to a response to legal price restrictions because of a number of other factors at play. For instance, hurricane-related supply chain disruptions could have increased the cost of bringing goods to consumers; the law against price gouging explicitly permitted the pass-through of such costs to consumers. Similarly, stockouts of items on sale could have contributed to the fall in the mean saving on the dollar as households turned to consuming non-sale items in greater proportion.

4.2 Major snowstorms and hurricanes

Major weather events such as large snowstorms and hurricanes can affect the demand for supermarket products through several channels. Storms that result in the closing of schools and workplaces force households to consume a greater proportion of their meals at home, thus boosting demand for food items. Similarly, the demand for personal care and housekeeping products may rise as households engage in greater home production or take advantage of their trip to the supermarket to purchase items other than food. Storms may also displace consumption across time periods by making it difficult or impossible to shop on certain days. In particular, some households may seek to build their domestic inventories in anticipation of a storm to ensure continued supplies, while others may need to replenish their inventories once the storm is over.

4.2.1 Identification

We have identified 59 combinations of an IRI market and a major snow episode whose disruptive consequences were favorable to the triggering of a shopping spree over the IRI sample period. Many of these combinations feature a peak in average store revenues of 10 percent or more relative to the previous few weeks. While there were hundreds of smaller snowstorms in our sample, we expressly choose to leave them aside because storms whose disruptive effects last only a couple of days are less likely to leave a clear imprint in supermarket data collected weekly.

In identifying disruptive snowstorm episodes, we account for the fact that some localities have a greater ability to cope with snowfall than others. Snow accumulations that have crippling effects in Southern states, where snowstorms are scarce and snow plowing equipment is in short supply, may have only limited disruptive effects in Northern states, where local authorities are accustomed to clearing snow off the streets rapidly. To do so, we match our IRI scanner data with the U.S. National Oceanic and Atmospheric Administration’s Regional Snowfall Index (RSI) and the Federal Emergency Management Agency’s (FEMA) list of federal disaster declarations. The RSI controls for differences in historical snow precipitation and the authorities’ ability to cope with snow through the setting of precipitation thresholds that are specific to nine U.S. regions (the West Coast of the United States is not covered due to insufficiently frequent snowstorms). Storms whose social impact is roughly in the top half of historical storms in their region are given a rank between 1 and 5. Many of the snow episodes in our sample were ranked “3–major,” “4–crippling,” or “5–exceptional,” placing them in the top 5 percent of storms in terms of regional-level disruptiveness. Many storms were also granted an “emergency” or a “disaster relief” status by FEMA because they were of “such severity and magnitude that effective response [was] beyond the capabilities of the State and the local governments and that Federal assistance [was] necessary.”²⁰

We validate our list of snow episodes against daily snowfall measurements reported by local weather stations to avoid situations in which a disruptive storm at the regional level results in

²⁰Disaster declarations put into motion short- to long-term federal relief, some of which may be directed to individuals. Emergency declarations are more limited in scope and seek to meet specific emergency needs or to help prevent major disasters from occurring. Both types of declaration require presidential approval.

little snow accumulation or passes as rain in a particular market. We also use local daily snowfall measurements to include a number of storms that likely had large localized effects but whose regional impact, as measured by the RSI, was small. In a few cases, our snow episodes cover two snowstorms rather than one because the separate meteorological systems are indistinguishable in weekly data. Finally, we incorporate a few major snowstorms from the U.S. West Coast for which RSI scores are not available. The list of snowstorms, along with their cumulative snowfall, RIS classification, and FEMA declaration, is provided in an online appendix.

We follow a similar strategy for identifying hurricanes that are likely to induce shopping sprees. We look at all “emergency” and “major disaster” declarations by FEMA that are attributed to a hurricane. We then validate our list against daily rainfall and maximum wind speed measurements from local weather stations.²¹ In total, we have identified 21 combinations of an IRI market and a hurricane, which we also list in our online appendix along with their characteristics.

4.2.2 Illustration: 2009–2010 winter in Washington, D.C.

Figure 7 illustrates some effects of major storms on retail activities using two snow episodes that hit Washington, D.C., during the 2009–2010 winter. The first episode began on Friday, December 18, 2009, and left 41.7 centimeters (16.4 inches) of snow at D.C.’s Reagan National Airport. Federal offices were closed the following Monday and operated on an unscheduled leave basis for two more days due to impracticable roads in parts of the metropolitan area. The second snow episode was more disruptive, consisting of two back-to-back blizzards that together blanketed the U.S. capital with 72.6 centimeters (28.6 inches) of snow. Federal offices closed early as the first blizzard moved in on Friday, February 5, 2010, remained closed through February 12, and then operated on an unscheduled leave basis through February 16. As is typical of the episodes in our sample, the National Weather Service and local media began reporting on the approaching snowstorms several days ahead of their occurrence.

The upper-left panel of figure 7 shows that sales volumes peaked 20 to 45 percent above their trend in the December 2009 and February 2010 snow episodes. The timing of the surge in quantities differs between the two episodes, although the use of weekly retail data limits our ability to identify their timing precisely (note that IRI weeks run from Monday to Sunday). As is apparent for both episodes, the quantity of food products and of personal care and housekeeping products spiked around the storms, supporting our treatment of major snowstorms as shocks to the overall demand of supermarkets rather than as shocks specific to some product categories. In the first episode, quantities of both groups of products rose 10 percent in the week of the storm relative to their recent trend, and rose even more in the ensuing week. In the second episode, sales volumes surged 27 percent for personal care and housekeeping products in the week encompassing the beginning of the second episode and 45 percent for food products. The demand for food products remained

²¹We exclude Hurricane Katrina and Hurricane Rita due to their exceptional mass population displacement, a phenomenon that we analyzed separately in section 4.1. We also drop most observations related to Hurricane Ike and Hurricane Lili because they came on the heel of other disruptive hurricanes that required federal assistance. Our regression results are robust to keeping these latter hurricanes in the sample.

above its recent trend during the ensuing week. The large impact of the second storm may be due to the anticipation of greater snow totals (and, if our experience is representative, it could also reflect some learning from the first episode that shopping in advance of the storm avoids some headaches).

The two snowstorm episodes left no apparent imprint on price indexes, as shown in the upper-right panel. Similarly, the frequency of price changes, the mean saving on the dollar, the share of items on sales, and the share of store revenues derived from items on sales were all within the range experienced over the 2009–2010 winter. Our econometric analysis below, on the broader sample of snowstorms and hurricanes, confirms these impressions.

4.2.3 Econometric analysis

There is some uncertainty regarding the precise timing of when weekly retail activities should feel the effects of storms most strongly. Storms that hit early in the week or whose disruptive effects are anticipated may affect retail activities prior to the storm. Similarly, storms that hit late in the week or that require domestic inventory rebuilding by households afterwards could have some effect in the week after the storm. As a first step into our investigation, we ignore this timing issue and compare various retailing statistics in the week corresponding to the observed peak in quantities to their respective average in the weeks prior to the storm. More precisely, for a storm beginning in week t , we identify the peak in quantity over the weeks $t - 1$, t , and $t + 1$ and then compare the statistic of interest for that peak week by its average over the weeks $t - 4$ to $t - 2$ (the “pre-storm period”).²²

Our sample of snow episodes supports the patterns apparent in figure 7, namely that major snowstorms boost consumer spending on supermarket items while having little if any influence on pricing. The mean peak in quantities around snowstorms is 12.9 percent higher than the average over the pre-storm period, whereas the peak in prices is only 0.1 percent higher. A simple statistical test that the population mean of the ratio across snow episodes equals 1 is rejected for quantities but not for prices. The peak quantity and price responses, as well as their statistical significance, are nearly identical for hurricanes. For both types of storms, we also find similar statistics on broader features of our stores’ pricing strategies between the week of the peak in quantities and the pre-storm period. This remarkable finding suggests that retailers broadly retain their usual pricing strategies.

We next follow a regression-based approach similar to that used by Chevalier, Kashyap, and Rossi (2003) in their study of demand peaks around holidays. We regress our statistics of interest (illustrated here with the log of our market-specific price index) on a quadratic time trend and a set of dummies marking the week immediately before ($W_{c,t}^{pre}$), the week during ($W_{c,t}^{during}$), and the

²²We normalize the quantity and price indexes in each market by dividing their value during the quantity peak week by their average during the reference period.

week after ($W_{c,t}^{after}$) a storm,

$$\log(P_{c,t}^F) = \alpha_{c,0} + \alpha_{c,1}t + \alpha_{c,2}t^2 + \beta_{-1}W_{c,t}^{pre} + \beta_0W_{c,t}^{during} + \beta_1W_{c,t}^{after} + \varepsilon_{c,t}.$$

The estimated coefficients on the week dummies can be interpreted as the average movement in the left-hand side statistic (relative to its trend) across the various storm episodes in the sample. To ensure that our quadratic trend fits the level of the individual series properly, we retain only data from the first week in October through the last week in March in the case of snowstorms and from the first week of July to the last week of December in the case of hurricanes (there are no qualifying storms outside of these broadly-defined snowstorm and hurricane seasons). We next fit a quadratic time trend that is specific to each IRI market and 6-month season combination, represented by the subscript c in the above equation. We then run separate regressions for snowstorms and hurricanes.

Table 4 reports the results for our key statistics of interest. For major snowstorms, we find that the boost to quantities occurs almost entirely in the week of the storm. We also observe a statistically-significant small decline in quantities of about 1.9 percent in the week after the storm. This decline is suggestive that greater shopping activities during snowstorms result in the bringing forward of some household expenditures, leading to a small pull back in the week immediately after. The corresponding estimates for prices point to a small increase during the week of the storm that is not statistically significant. In the case of hurricanes, we find a statistically significant boost to spending in both the week immediately before and the week of the storm. This finding suggests that hurricanes primarily pull forward consumption expenditures, perhaps due to a greater predictability of these events and a greater risk that households could experience reduced supplies for a protracted period. A pull back is also observed in the week immediately after a hurricane, although it is not statistically significant at standard confidence levels.

The absence of a price response does not appear to be due to legal impediments against price increases. Indeed, Zwolinski (2008) lists 33 U.S. states restricting price rises during declared emergencies. The laws typically seek to prevent opportunistic rises in the price of gasoline, hotel rooms, construction materials, and food items that are most susceptible to cause public uproar. Although restrictive, nearly all of the laws explicitly permit price increases that are attributable to additional costs or expenses related to an emergency; 12 states even explicitly permit price increases that do not exceed some value between 10 to 25 percent. Price setters may also have some room to raise prices by reducing the depth or frequency of sales. Price increases are typically measured against the prices “ordinarily charged” or observed over some period prior to the disaster. Therefore, a retailer may not break the law if it were to raise effective prices by diminishing the frequency or depth of sales and promotions. In any case, as a check of the importance of price gouging laws, we added to our regression a set of law dummies interacted $W_{c,t}^{pre}$, $W_{c,t}^{during}$, and $W_{c,t}^{after}$. We could not reject the hypothesis that price responses in markets not subject to laws against price gouging were the same as price responses in markets subject to such law.

The observation that retailers deviate little if at all from their usual pricing strategies during major storms contrasts with the finding that prices tend to fall during periods of peak demand, in

particular short-lived booms around holidays (see Warner and Barsky (1995), MacDonald (2000), and Chevalier, Kashyap, and Rossi (2003)). The difference could reflect a number of factors. Notably, whereas the timing of holidays is perfectly predictable, the occurrence of a major snowstorm or a hurricane can be anticipated at most a week or so in advance, and only with great uncertainty. This limited predictability may not leave enough time for manufacturers to adjust production or for retailers to implement “loss-leader” pricing strategies to lure customers into their stores. Circulars may already have been printed and, in any case, there can be lags of several months in the planning of sales and promotions (see Anderson *et al.* (2013) for a discussion).

5 What factors account for a flat retail supply curve?

We have shown that supermarkets do not boost prices in response to peaks in demand brought about by labor conflicts, mass population relocation, and shopping sprees around storms and hurricanes. This evidence points to flat short- to medium-run supply curves in the retail sector. In this section, we discuss some factors that might contribute to this flatness.

We first note that information friction theories do not seem to offer a plausible explanation. Our shocks were easily observable, benefiting from extensive media coverage. For this reason, it seems inconceivable that supermarkets would have failed to adjust prices in response to a labor conflict or a major storm because they were either unaware its occurrence (as would be the case with Mankiw and Reis’ (2002) sticky information model) or because they were overly focused on idiosyncratic shocks (as would be the case in Maćkowiak and Wiederholt’s (2009) rational inattention model). Our evidence also does not support the idea that price gouging laws, where present, were important impediments to price increases. In the case of Hurricane Katrina, the expiration of anti-gouging provisions four weeks after the disaster struck was not accompanied by significant price rises. We also find no evidence in support of larger price responses to snowstorms in markets where no such laws exist relative to markets where they do.

The applicability of countercyclical markup theories based on countercyclical collusion (Rotemberg and Saloner (1986)) or economies of scale in consumer search (for example, Bills (1999) and Warner and Barsky (1995)) seems limited in our context. In the case of countercyclical collusion, the retail markup falls when demand exogenously rises because retailers have greater incentives to deviate from cartel prices. In contrast with this prediction, the number of establishments where households were either willing or able to shop arguably declined during our labor conflicts and Hurricane Katrina, which could in principle have made collusion easier to sustain. In the case of economies of scale in consumer search, consumers increased their search intensity in periods of high demand, such as during holiday seasons, making their demand more elastic than in normal times and thus limiting retailers’ ability to enjoy high margins. Although shopping intensity jumped around our major snowstorms and hurricanes, households arguably had limited time to search at multiple establishments. As a result, their demand may have been less elastic than otherwise, in contrast with the assumption at the heart of the theory.

A number of other theories seem consistent with some but not all of our findings. Notably, the lack of a price response to unanticipated, short-lived demand shocks could be consistent with standard sticky-price and sticky-plan theories, and is supported by evidence of lags in retailers' planning of sales and promotions and infrequent regular prices adjustments presented in Anderson *et al.* (2013). That said, rigidities in prices or in pricing strategies do not appear to be a fully satisfactory explanation for the lack of a stronger price response to our most persistent shocks—namely the Southern California labor conflict and Hurricane Katrina—for which supermarkets often adjusted item prices several times but did not, on net, raise prices significantly.

Theories emphasizing the customer base as an asset for retailers also get some partial backing from our findings to the extent that one interprete the assumption of consumer loyalty to specific items or brands in terms of loyalty to particular establishments or retail chains. These theories include switching costs (for example, Klemperer (1995) and Kleshchelski and Vincent (2009)) and models of deep habits (Ravn, Schmitt-Grohé, and Uribe (2006)). Consistent with these theories, the effect of the Southern California strike and lockout on store traffic was not fully reversed at the end of the conflict. However, the effect was much less pronounced than posited by these theories. Moreover, customer base losses can take time to materialize, as evidenced by the full bounce back in sales volumes at stores impacted by the relatively short St. Louis conflict. In addition, these theories could have difficulties explaining the lack of a strong response to our shortest-lived shocks, as they imply that demand is much less elastic in the short run than over the long run.

Overall, we retain two explanations that appear broadly consistent with our findings. The first explanation is that retailers' marginal costs and markups are insensitive to the level of demand. Although we do not observe these variables in our dataset, a number of elements support the plausibility of this explanation. As noted in section 3, goods acquisition costs account for the bulk of supermarket expenses. Given the regional character of our shocks and the ability of most of our retailers' to move stocks across regions, it is conceivable that acquisition costs could remain fairly stable even as demand fluctuates, especially if shocks are not too persistent. With respect to markups, Eichenbaum, Jaimovich, and Rebelo (2011) show that item-level markups vary little around their mean using data from a large U.S. supermarket chain. Relatedly, Fabiani *et al.* (2006) provide evidence of the pervasiveness of constant markup rules based on a large survey of European firms. The micro evidence on the sources of variations in exchange rates is further supportive of limited variation in item markups; using data from a leading North American retailer, Gopinath, Gourinchas, Hsieh, and Li (2011) show that variation in markups account for none of the cross-border variation in retail prices between Canada and the United States.

The second explanation is that fairness motives may be at play. Boosting prices to take advantage of high demand could risk antagonizing customers, leading to longer-term losses that more than offset any short-run benefits. For instance, Kahneman, Knetsch, and Thaler (1986) analyze responses to public opinion surveys and conclude that households perceive firms that raise prices in response to shifts in demand as acting unfairly.²³ For our major snowstorms and hurricanes,

²³This conclusion is supported by the subsequent literature; see the review by Xia, Monroe, and Cox (2004). See

retailers may have wanted to avoid being perceived as unjustly profiting off their customers' unusually high marginal utility for fear of losing their business in the future. That said, we note that our sample excludes products whose consumption is essential during storms, such as de-icing salt, batteries, or snow shovels. Our sample instead contains products consumed year-round and for which several brands are typically available. A retailer that would want to take advantage of temporarily high demand may thus have to raise prices over a broad product offering. This feat could be achieved by reducing the number or depth of sales rather than by raising regular prices but our results offer only limited support for these channels.

6 Conclusion

We have shown that the level of supermarket prices responds little to large swings in demand brought about by labor conflicts, mass population relocation, and shopping sprees around storms and hurricanes. This evidence is consistent with flat short- to medium-term supply curves in the retail sector. In particular, the evidence seems inconsistent with the marginal cost of retailers being sensitive to the level of demand because of fixed factors of production. It seems sensible to conjecture that our low supply elasticity estimates could reflect relatively steady marginal costs and markups, or possibly a fear of antagonizing customers through price changes perceived as unfair.

If our retailers responded little to large demand shocks overall, they did seem concerned with keeping up with the pricing strategies of their local competitors. This fact was most clearly seen during and after the Southern California supermarket strike and lockout, when retailers with radically different demand shocks nonetheless matched their local competitors' pricing movements and recourse to sales of various depths.

The above observations invite a reconsideration of the place occupied by the retail sector in macro models. Many modelers conflate the notion of producers and retailers, and then calibrate their model to match the pricing features of retailers only. Our analysis suggests that the retail sector's short- to medium-term supply curve is perhaps much flatter than that of other sectors such as manufacturing; this finding is in line with the low-margin, high-volume nature of the retail industry. If so, then a better understanding of pricing behavior and deviations from constant returns to scale at lower levels of the production chain seem much needed.

also Rotemberg (2005, 2011) for an exploration of the pricing and promotional sales implications of fairness within the context of a modern microfounded pricing model.

References

- [1] Anderson, Eric, Emi Nakamura, Duncan Simester, and Jón Steinsson (2013). “Informational Rigidities and the Stickiness of Temporary Sales,” NBER Working Papers 19350, National Bureau of Economic Research, Inc.
- [2] Angrist, Joshua D. and Guido W. Imbens (1995). “Two-Stage Least Squares Estimation of the Average Causal Effects in Models with Variable Treatment Intensity,” *Journal of the American Statistical Association*, vol. 90(430), pages 431–442.
- [3] Barro, Robert (1972). “A Theory of Monopolistic Price Adjustment,” *Review of Economic Studies*, vol. 39(1), pages 17–26.
- [4] Bernanke, Ben, Jean Boivin, and Piotr S. Elias (2005). “Measuring the Effects of Monetary Policy: A Factor-augmented Vector Autoregressive (FAVAR) Approach,” *Quarterly Journal of Economics*, vol. 120(1), pages 387–422.
- [5] Bills, Mark (1989). “Pricing in a Customer Market,” *Quarterly Journal of Economics*, vol. 104(4), pages 699–717.
- [6] Bills, Mark and Peter J. Klenow (2004). “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, vol. 112(5), pages 947–985.
- [7] Blake, Eric S. and Ethan J. Gibney (2011). “The Deadliest, Costliest, and Most Intense United States Tropical Cyclones from 1851 to 2010 (and Other Frequently Requested Hurricane Facts),” NOAA Technical Memorandum NWS NHC-6, U.S. National Oceanic and Atmospheric Administration.
- [8] Bradley, Ralph (2012). “Price Index Estimation Using Price Imputation for Unsold Items,” chapter in: *Scanner Data and Price Indexes*, NBER Studies in Income and Wealth, vol. 64, pages 349–382, University of Chicago Press.
- [9] Bronnenberg, Bart J., Michael W. Kruger, and Carl F. Mela (2008). “Database Paper: The IRI Marketing Data Set,” *Marketing Science*, vol. 27(4), pages 745–748.
- [10] Hellwig, Christian and Ariel Burstein (2007). “Prices and Market Shares in a Menu Cost Model,” NBER Working Papers 13455, National Bureau of Economic Research, Inc.
- [11] Cavallo, Alberto, Eduardo Cavallo, and Roberto Rigobon (2013). “Prices and Supply Disruptions during Natural Disasters,” NBER Working Papers 19474, National Bureau of Economic Research, Inc.
- [12] Chevalier, Judith A. and Anil K. Kashyap (2013). “Best Prices: Price Discrimination and Consumer Substitution,” NBER Working Papers 16680, National Bureau of Economic Research, Inc.
- [13] Chevalier, Judith A., Anil K. Kashyap, and Peter E. Rossi (2003). “Why Don’t Prices Rise During Periods of Peak Demand? Evidence from Scanner Data,” *American Economic Review*, vol. 93(1), pages 15–37.
- [14] Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (1999). “Monetary policy shocks: What have we learned and to what end?,” *Handbook of Macroeconomics*, in: John B. Taylor and Michael Woodford (ed.), vol. 1, pages 65–148, Elsevier.

- [15] Cleveland, William P. and Stuart Scott (2007), “Seasonal Adjustment of Weekly Time Series with Application to Unemployment Insurance Claims and Steel Production,” *Journal of Official Statistics*, vol. 23, pages 209–221.
- [16] Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong (2012). “The Cyclicalities of Sales, Regular and Effective Prices: Business Cycle and Policy Implications,” NBER Working Papers 18273, National Bureau of Economic Research, Inc.
- [17] Dhyne, Emmanuel, Luis J. Álvarez, Hervé Le Bihan, Giovanni Veronese, Daniel Dias, Johannes Hoffmann, Nicole Jonker, Patrick Lünemann, Fabio Rünstler, and Jouko Vilmunen (2005). “Price Setting in the Euro Area: Some Stylized Facts from Individual Consumer Price Data,” Working Paper 524, European Central Bank.
- [18] Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo (2011). “Reference Prices, Costs, and Nominal Rigidities,” *American Economic Review*, vol. 101(1), pages 234–262.
- [19] Fabiani, Silvia, Martine Druant, Ignacio Hernando, Claudia Kwapil, Bettina Landau, Claire Loupias, Fernando Martins, Thomas Mathä, Roberto Sabbatini, Harald Stahl, and Ad Stokman (2006). “What Firms’ Surveys Tell Us about Price-Setting Behavior in the Euro Area,” *International Journal of Central Banking*, vol. 2(3), pages 3–47.
- [20] Federal Emergency Management Agency, “Declared Disasters by Year or State,” last accessed December 2, 2011, http://www.fema.gov/news/disaster_totals_annual.fema.
- [21] Gagnon, Etienne, David López-Salido, and Nicolas Vincent (2013). “Individual Price Adjustment along the Extensive Margin,” chapter in: *NBER Macroeconomics Annual 2012*, vol. 27, pages 2450–2486.
- [22] Geaghan, Kimberly A. (2011). “Forced to Move: An Analysis of Hurricane Katrina Movers 2009 American Housing Survey: New Orleans,” Social, Economic, and Housing Statistics Division Working Paper Number 2011-17, U.S. Census Bureau.
- [23] Golosov, Mikhail and Robert E. Lucas Jr. (2007). “Menu Costs and Phillips Curves,” *Journal of Political Economy*, vol. 115, pages 171–199.
- [24] Gopinath, Gita, Pierre-Olivier Gourinchas, Chang-Tai Hsieh, and Nicholas Li (2011). “International Prices, Costs, and Markup Differences,” *American Economic Review*, vol. 101(6), pages 234–262.
- [25] Gopinath, Gita and Oleg Itskhoki (2010). “Frequency of Price Adjustment and Pass-Through,” *Quarterly Journal of Economics*, vol. 125 (2), pages 675–727.
- [26] Hall, Robert E. (2009). “By How Much Does GDP Rise If the Government Buys More Output?,” *Brookings Papers on Economic Activity*, vol. 40(2), pages 183–249.
- [27] Hong, Gee Hee and Nicholas Li (2014). “Market Structure and Cost Pass-Through in Retail,” Working Papers tecipa-470, University of Toronto, Department of Economics.
- [28] Kahneman, Daniel, Jack L. Knetsch, and Richard Thaler (1986). “Fairness as a Constraint on Profit Seeking: Entitlements in the Market,” *American Economic Review*, vol. 76(4), pages 728–41.

- [29] Karadi, Peter and Adam Reiff (2012). “Large shocks in menu cost models,” Working Paper Series 1453, European Central Bank.
- [30] Klemperer, Paul, 1995. “Competition When Consumers Have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade,” *Review of Economic Studies*, vol. 62(4), pages 515–539.
- [31] Klenow, Peter J. and Oleksiy Kryvtsov (2008). “State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?,” *Quarterly Journal of Economics*, vol. 123(3), pages 863–904.
- [32] Kleshchelski, Isaac and Nicolas Vincent (2009). “Market Share and Price Rigidity,” *Journal of Monetary Economics*, vol. 56(3), pages 344–352.
- [33] Kruger, Michael W. and Daniel Pagni (2012). *IRI Academic Data Set Description*, version 2.0, Chicago: Information Resources Incorporated.
- [34] Lach, Saul (2007). “Immigration and Prices,” *Journal of Political Economy*, vol. 115(4), pages 548–587.
- [35] Liu, Zheng and Sylvain Leduc (2013). “Uncertainty Shocks Are Aggregate Demand Shocks,” 2013 Meeting Papers 270, Society for Economic Dynamics.
- [36] Lee, Lung-Fei and Mark M. Pitt (1986). “Microeconomic Demand Systems with Binding Nonnegativity Constraints: The Dual Approach,” *Econometrica*, Econometric Society, vol. 54(5), pages 1237–1242.
- [37] MacDonald, James M. (2000). “Demand, Information, and Competition: Why Do Food Prices Fall at Seasonal Demand Peaks?,” *Journal of Industrial Economics*, vol. 48(1), pages 27–45.
- [38] Mackowiak, Bartosz and Mirko Wiederholt (2009). “Optimal Sticky Prices under Rational Inattention,” *American Economic Review*, vol. 99(3), pages 769–803.
- [39] Mankiw, N. Gregory and Ricardo Reis (2002). “Sticky Information Versus Sticky Prices: A Proposal To Replace The New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, vol. 117(4), pages 1295–1328.
- [40] Nakamura, Emi (2008). “Pass-Through in Retail and Wholesale,” *American Economic Review: Papers & Proceedings*, vol. 98(2), pages 430–437.
- [41] Nakamura, Emi and Jón Steinsson (2008). “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics*, vol. 123(4), pages 1415–1464.
- [42] National Climate Data Center, “NCDC Climate Data Online: Daily Surface Data,” last accessed December 5, 2011, <http://cdo.ncdc.noaa.gov/pls/plclimprod/somdmain.somdwrapper>.
- [43] Martin Pesendorfer (2002). “Retail Sales: A Study of Pricing Behavior in Supermarkets,” *Journal of Business*, vol. 75(1), pages 33–66.
- [44] Ravn, Morten, Stephanie Schmitt-Grohé, and Martín Uribe (2006). “Deep Habits,” *Review of Economic Studies*, Oxford University Press, vol. 73(1), pages 195–218.
- [45] Romer, Christina D. and David H. Romer (2004). “A New Measure of Monetary Shocks: Derivation and Implications,” *American Economic Review*, vol. 94(4), pages 1055–1084.

- [46] Rotemberg, Julio J. (2005). “Customer anger at price increases, changes in the frequency of price adjustment and monetary policy,” *Journal of Monetary Economics*, vol. 52(4), pages 829–852.
- [47] Rotemberg, Julio J. (2011). “Fair Pricing,” *Journal of the European Economic Association*, vol. 9(5), pages 952–981.
- [48] Rotemberg, Julio J. and Garth Saloner (1986). “A Supergame-Theoretic Model of Price Wars during Booms,” *American Economic Review*, vol. 76(3), pages 390–407.
- [49] Shea, John (1993). “Do Supply Curves Slope Up?,” *Quarterly Journal of Economics*, vol. 108(1), pages 1–32.
- [50] Warner, Elizabeth J. and Robert B. Barsky (1995). “The Timing and Magnitude of Retail Store Markdowns: Evidence from Weekends and Holidays,” *Quarterly Journal of Economics*, vol. 110(2), pages 321–352.
- [51] Xia, Lan, Kent B. Monroe, and Jennifer L. Cox (2004). “The Price Is Unfair! A Conceptual Framework of Price Fairness Perceptions,” *Journal of Marketing*, vol. 68(4), pages 1–15.
- [52] Zwolinski, Matt (2008). “The Ethics of Price Gouging,” *Business Ethics Quarterly*, vol. 18(3), pages 347–378.

Appendix A: Cleaning and organizing the IRI data

A.1 Data trimming and filtering

We perform a number of data cleaning steps to create a sample suitable for our purposes. We only use observations pertaining to grocery stores to ensure maximum comparability across stores and retail chains. The IRI sample also includes observations from drugstores but the number of such stores is more limited than for grocery stores. We exclude two product categories from the sample, cigarettes and photo supplies, because their prices and advertising are heavily regulated and they gradually became obsolete over our sample period, respectively. As is common with scanner data, we obtain a weekly transaction price by dividing total revenue by the number of units sold. In a tiny proportion of cases, the division yields a price with fractional cents, which Kruger and Pagni (2012) attribute to floating point conversions. For New York City, fractional prices represent less than 0.2 percent of non-missing prices. We correct all fractional prices for which a one-cent difference in item revenues results in a non-fractional price, and round the remaining fractional prices to the nearest cent as recommended in the IRI documentation.

In addition to analyzing prices paid by consumers, we consider measures of regular prices to capture lower-frequency movements in prices, in particular temporary sales and promotions. Our regular price filter is based on that proposed by Nakamura and Steinsson (2008) in section A of their technical appendix. It removes temporary price drops that are followed by an increase to a price at or above the previous price or to a new regular price. Our implementation uses the parameters $J = 6$, $K = 3$, and $L = 3$.²⁴ The adequacy of this parameterization for identifying regular and sales prices necessarily involves an judgmental element on our part. Contrary to the CPI research database maintained by the BLS, the IRI dataset contains no flag that an item is subject to an on-going sales or promotion based on in-store observations (although some information about displays and circulars may be provided). Filtered data thus cannot be validated against actual regular prices. Moreover, reported revenues typically comprise retailer discounts, may or may not include discounts extended through frequent purchase programs, and may be subject to reporting errors. For these reasons, some high-frequency price movements that are unrelated to sales could be present. As an investigative check, we consider a number of alternative filters, including a reference price filter based on a variant of that of Eichenbaum, Jaimovich, and Rebelo (2011). This filter defines the reference price of an item in week t as the modal price observed over a 13-week centered window. We also consider a filter removing solely V-shaped sales as well as a filter based on a price reduction flag computed by IRI a proprietary formula that, like our own filters, uses reported average price series rather than in-store observations. Typically, filters that produce relatively infrequent regular price adjustments are associated with relatively frequent but shallower sales, suggesting the possible presence of tradeoffs between accurately estimating regular price adjustments and accurately measuring the frequency and depth of temporary sales. Importantly

²⁴Nakamura and Steinsson’s (2008) technical appendix does not discuss the issue of missing prices, which are quite common in our weekly scanner data. We carried forward the last observed price to fill in missing prices before running the filter. Regular prices associated with no units sold were censored after running the filter.

for our purposes, the choice of a filter does not appear to alter our general conclusions about the effects of large demand shocks on prices.

To compute monthly price statistics, we retain only weekly observations that encompass the 15th day of each month to match the BLS’ practice of sampling in the same narrow time window every month. Finally, all weekly and monthly statistics employ only observations from frequently traded items that are free of anomalous price movements, as we now detail in the context of deriving price and quantity indexes.

A.2 Constructing indexes of total revenues, prices, and quantities

We first compute indexes of total store revenues at the IRI market–product category level (the “stratum” level). This task is made possible by the IRI dataset’s comprehensive coverage of items available for sale at participating stores within each of the 29 product categories. This feature makes it possible to track the evolution of a store’s total revenues by product category over time and to monitor the entry and exit of items in its offering. Occasionally, all observations from a store may be coded as missing. In such problematic weeks, we impute missing store’s total revenues using average revenue growth for the subset of stores in the same IRI market whose revenues for the product category are observed. Formally, let s index stores in market m and product category c that are present in the sample at week t and for which total sales in week $t - 1$ are known (because they were either directly observed or imputed from the preceding period). Let also $I_{s,t}^{m,c}$ be an indicator that store s ’ data are not systematically missing in the period. Whenever all observations are missing at some stores, we first compute the revenue growth for stores that are reporting correctly in period t ,

$$g_t^{m,c} = \left(\sum_{s|I_{s,t}^{m,c}=1} revenues_{s,t}^{m,c} \right) / \left(\sum_{s|I_{s,t}^{m,c}=1} revenues_{s,t-1}^{m,c} \right),$$

and set total revenues at problematic stores equal to $g_t^{m,c} \cdot revenues_{s,t-1}^{m,c}$. The estimated growth rate of revenues in the stratum is then

$$\frac{revenues_t^{m,c}}{revenues_{t-1}^{m,c}} = \frac{\sum_s revenues_{s,t}^{m,c}}{\sum_s revenues_{s,t-1}^{m,c}}.$$

This estimate is transformed into an index of the level of total revenues by computing its cumulative product over time. The index is scaled so that its geometric mean in 2005 matches average weekly nominal sales observed in the sample that year.

We follow a methodology similar to that employed by the BLS for the U.S. CPI to compute disaggregated price indexes for the IRI sample. We first obtain an estimate of inflation at the IRI market–product category level using a geometric mean formula,

$$\pi_t^{m,c} = \prod_{i \in I_t^{m,c}} \left(\frac{p_{i,t}}{p_{i,t-1}} \right)^{w_{i,t}},$$

where $\pi_t^{m,c}$ is (gross) inflation in period t for market m and product category c , $w_{i,t}$ is the relative weight of item i , and $I_t^{m,c}$ is the set of admissible items. In practice, the BLS adjusts item weights to ensure that the CPI sample is representative of U.S. households' shopping habits as established through its Point-of-Purchase Survey. For the IRI sample, we use uniform weights across admissible items in part due to the greater similarity in our sample of stores, which all pertain to a retail chain. Like the BLS, we restrict the set of items admissible to compute inflation to those that are actively traded, a restriction which, in the case of the IRI sample, has the added benefit of limiting inference issues created by the censoring of price observations that have no transactions. In particular, we restrict admissible price observations to those pertaining to blocks of 52 consecutive weeks with at most 15 percent of missing observations (that is, at most 8 missing weekly prices). Like the BLS, we impute missing item prices by incrementing their previous weekly price by the average log price increase of items pertaining to the same stratum whose prices are not missing in the current period.²⁵ To limit measurement error, we exclude all items featuring a movement in their weekly log price in excess of 1.9 (such a movement corresponds to a drop in excess of 85 percent in the case of negative adjustments). We also discard every item's first and last 12 price observations, thus shaving off roughly one quarter worth of price changes at each end of item price histories. This latter trimming seeks to prevent selection effects associated with the IRI's practice of dropping from its sample observations with missing prices at either the beginning or the end of price trajectories. Unless indicated otherwise, all statistics reported in the paper are computed using the subsample of frequently traded items that are free of suspiciously large price movements and whose price histories are trimmed at both ends.

We transform our inflation estimates at the stratum into measures of the price level by calculating these estimates' cumulative product over time. We then scale the price indexes so that their geometric mean equals unity in 2005. We finally obtain a real same-store quantity index at the stratum level by deflating our total revenue index by our geometric price indexes.²⁶ This approach of computing a quantity index by effectively pooling price and spending data from separate surveys mimics that employed by the BEA for disaggregated real personal consumption expenditures in the U.S. National Income and Product Accounts.

Also echoing the BEA methodology, we aggregate stratum statistics across product categories using a chain-weighted approach. We compute a Fisher quantity index, $Q_t^F = \sqrt{Q_t^P Q_t^L}$, as the geometric average of a Paasche quantity index,

$$Q_t^P = \frac{\sum_j p_{j,t} q_{j,t}}{\sum_j p_{j,t} q_{j,t-1}},$$

²⁵We also experimented with carrying forward an item's last observed price whenever its current price is missing. The resulting price indexes were nearly identical.

²⁶The interpretation of our indexes is complicated by a change in the treatment of frequent shopper card discounts. Starting January 2002, IRI is reporting dollars spending net of loyalty program discounts for the subset of stores that makes use of these programs and report the value of these rebates to IRI. Before that date, participating stores only reported total sales before rebates are applied. We ignore for now possible breaks in our indexes created by this methodological change but our analysis of natural disasters will flag the handful cases for which this change could be problematic.

and a Laspeyres quantity index,

$$Q_t^L = \frac{\sum_j p_{j,t-1} q_{j,t}}{\sum_j p_{j,t-1} q_{j,t-1}}.$$

The index j applies to all stratum statistics being aggregated. Similarly, we calculate a Fisher price index, $P_t^F = \sqrt{P_t^P P_t^L}$, as the geometric average of a Paasche price index,

$$P_t^P = \frac{\sum_j p_{j,t} q_{j,t}}{\sum_j p_{j,t-1} q_{j,t}},$$

and a Laspeyres price index,

$$P_t^L = \frac{\sum_j p_{j,t} q_{j,t-1}}{\sum_j p_{j,t-1} q_{j,t-1}}.^{27}$$

A.3 Comparing the level of prices

One attractive feature of the IRI sample is that we observe item barcodes, which allows us to directly compare the level of prices across stores for specific items or for baskets of items. Let $\tilde{I}_t$ be a set of items sold simultaneously in stores s and s' in period t . Let $\tilde{q}_{i,t}$ be the number of units of item i in the comparison basket.²⁸ The price of purchasing the basket in store s relative to the price of purchasing it in store s' is given by

$$R_t^{s,s'} = \frac{\sum_{i \in \tilde{I}_t} p_{i,s,t} \tilde{q}_{i,t}}{\sum_{i \in \tilde{I}_t} p_{i,s',t} \tilde{q}_{i,t}},$$

where $p_{i,s,t}$ denotes the price of item i in store s at time t . This formula can be used to compare the level of price across groups of stores by redefining $p_{i,s,t}$ and $p_{i,s',t}$ as the average unit price of item i in groups of stores s and s' , respectively.

A.4 Time series adjustments for breaks and seasonal patterns

Several IRI product categories display a jump in the number of barcodes per store at the beginning of 2007. The origin of this jump is not explicated in IRI documentation but it likely reflects a broadening of product category definitions rather than an actual increase in the number of barcodes offered by stores. To prevent spurious breaks in our disaggregated total sales and quantity indexes, we impute total sales growth in the problematic week at the stratum level using average total

²⁷Contrary to geomeans price indexes, Fisher price indexes have the attractive property of taking into account variation in spending shares across periods. We contemplated the computation of Fisher formulas directly on micro price and quantity data rather than on our disaggregated indexes. Unfortunately, the presence of weeks without transactions and, more generally, missing observations makes this approach largely impractical on a large scale. See Lee and Pitt (1986) for a discussion of zero transaction data in an industrial organization context and Bradley (2003) for an exploration of alternative imputation approaches.

²⁸The computation of our geomeans price indexes assumed that the comparison basket was fixed over time. Here we allow the composition to vary to be able to compare, say, how much the basket purchased by the consumers of store s would have cost had they instead shopped in store s' .

sales growth of items that are present in both the last week of 2006 and the first week of 2007.²⁹ In addition, a number of stores exiting the IRI sample early report exceptionally large sales—sometimes over ten times normal volumes—during their last few weeks in the sample. The reason for such skyrocketing reported sales is unclear to us but their occurrence seems largely orthogonal to the object of our study. For this reason, we systematically drop the last four weeks in the sample of stores that exit early to remove any associated breaks in the level of our series.

In addition to controlling for breaks, we often want to control for seasonal variation and holidays when comparing retailing activities during, say, natural disasters, to those during periods in which they are absent. To this end, we filter our stratum-level indexes with a procedure developed by Cleveland and Scott (2007) for weekly time series. The filter allows us to account for the precise timing of holidays and to exclude particular periods from the computation of the seasonal factors. Our analysis of labor conflicts excludes the strike periods while that of Hurricane Katrina drops the week immediately before the storm and the entire federal disaster declaration period. For shopping spree around major snowstorms and hurricanes, we exclude the weeks immediately before, during, and after the storms. National aggregates use all weekly observations. Price and quantity indexes use multiplicative seasonal factors whereas all other indicators, which are in percent, use additive factors.

²⁹The imputation proceeds by regressing total sales growth for all items on a constant and total sales growth for continuing items in the stratum over a centered 52-week window around problematic periods. We exclude observations in the window for which the difference in total sale growth between the full sample and the subsample of continuing observations exceed 5 percent. We then compute the best linear predictor for problematic periods. The first and last 12 weekly observations of each item history are excluded from the calculation of sales growth for continuing items to avoid selection effects related to the trimming of items with no transactions.

Appendix B: Central features of retail pricing in normal times

We begin by describing the general behavior of retail prices in the dataset to provide us with a benchmark against which to compare pricing behavior in the presence of large demand shocks. The key message is threefold: retail prices are adjusted frequently, price adjustments are primarily associated with sales and promotions, and the bulk of barcode-level price variation is common across stores belonging to the same retail chain.

Fact #1: Retail prices are adjusted frequently

Table B.1 presents some key statistics on posted price adjustments broken down by product category.³⁰ On average, we are able to compute a posted price change for nearly 4 million items each week, which are associated with weekly revenues of about \$110 million. Of these items, 30.4 percent experience a (nonzero) posted price adjustment.³¹ The weekly frequency of price changes is nontrivial for all product categories, ranging from 13.7 percent for sugar substitutes to 43.9 percent for carbonated beverages. When we use only weekly price observations corresponding to the 15th day of each month, we derive a mean monthly frequency of price changes of 42.3 percent. Again, all product categories display frequent price adjustments, with sugar substitutes having the lowest monthly frequency, at 27.5 percent, and frozen dinner prices having the highest monthly frequency, at 55.2 percent. The mean monthly frequency of price changes in our sample is noticeably higher than corresponding estimates for the U.S. CPI, which range from about 26 percent in Bils and Klenow (2004) and Nakamura and Steinsson (2008), to about 36 percent in Klenow and Kryvtsov (2008). This difference is not due to the IRI sample being tilted toward product categories that have relatively high frequencies of price changes; if we restrict the CPI sample to personal care and processed food categories, we find a mean frequency more similar to that of the full IRI basket. Instead, the IRI sample’s relatively high frequency of price changes likely reflects a selection of stores that are actively changing prices and possibly other influences on average weekly prices, such as retailer coupons, that may not be picked up in official statistics.

Fact #2: Retail price adjustments are primarily driven by sales

A majority of price changes in our sample are associated with sales. To establish this fact, we identify sales using a version of the filter proposed by Nakamura and Steinsson (2008) that eliminates temporary downward price movements. This filter was described in appendix A. Table B.2 shows that a relatively modest share (11.2 percent) of regular prices are adjusted in a typical week, consistent with two out of every three price adjustments in our sample being driven by sales and promotions. The pre-eminence of sales in price adjustments is found in all product categories. On a monthly basis, 23.7 percent of regular prices are adjusted, so that sales account for about half of monthly price adjustments. The fact that sales represent a noticeably smaller fraction of

³⁰Our product-category statistics differ some from those reported by Coibion, Gorodnichenko, and Hong (2012) for the same dataset for a number of reasons. Their sample stops in 2007 whereas ours runs through 2011, and our respective sales filters are not identical. Their analysis includes drugstores whereas we focus on grocery stores to ensure homogeneous establishments and because drugstores have sparse food coverage. Furthermore, as our appendix A detailed, we sought to adopt a methodology similar to that in use for official statistics.

³¹Our sample-wide statistics aggregate product-category statistics by weekly sales.

monthly price adjustments than of weekly price adjustments indicates that many sales-related price movements are missed when data are collected monthly.

We also experimented with identifying low-frequency price movements using the concept popularized by Eichenbaum, Jaimovich, and Rebelo (2011) of a “reference price,” which we define as an item’s modal price over a centered 13-week window. Deviations from reference prices are even more common than deviations from regular prices. At a weekly frequency, reference prices are adjusted only 4.5 percent of the time, implying that roughly six out of seven price adjustments in our sample are transitory deviations from the reference prices. At a monthly frequency, reference prices are adjusted 15.6 percent of the time, consistent with over two thirds of monthly price adjustments being transitory deviations from reference prices.

Fact #3: Most prices are set at the retail chain level

Price behavior in the IRI sample is consistent with pricing decisions originating principally at the retail chain level. To establish this fact, we define a weekly “chain price” as the modal price at which a barcode is sold across stores belonging to the same retail chain within a market. To ensure that we have a meaningfully large sample of stores to compute a chain price, we require that a barcode’s weekly posted price be observed at a minimum of five stores belonging to the same combination of a retail chain and market, in addition to imposing that items are traded frequently (see appendix A for the details of our data cleaning). If there is more than one modal price, then we use the largest one as the chain price. These methodological choices are immaterial for our finding that most pricing decisions do not originate at the store level.

Table B.3 presents our key statistics on retail chain pricing. Overall, 71.2 percent of posted prices equal their corresponding chain price, a clear indication that the level of posted price is coordinated to some degree across stores belonging to the same retail chain. This conclusion applies to all product categories in our sample, with the lowest proportion being a still-substantial 65.8 percent in the case of carbonated beverages. It also applies to regular prices, for which we find similar proportions of prices matching the modal regular price of their barcode across stores belonging to the same retail chain. We also find evidence of an appreciable degree of coordination in the timing of price adjustments. As noted earlier, about 30 percent of posted prices are adjusted in a typical week. However, the probability of an item experiencing a price change jumps to over 80 percent when we condition on a change in the item’s chain price. (Here we exclude the item’s own price from the computation of the chain price to avoid creating a spurious increase in the price change probability.) In sharp contrast, if there is no adjustment in the chain price, then the item has only a 12.3 percent probability of being adjusted. As was the case with the level of prices, the finding of strong coordination in price changes across stores belonging to the same retail chain applies to all product categories and is found for both posted and regular prices.

Summing up, the strong synchronization across stores in the level of prices and the timing of their adjustments indicates that stores play a relatively minor role in the determination of prices in our sample. This finding is consistent with Nakamura (2008), who explores similar issues using a multi-market sample of about 100 barcodes in 2004 from AC Nielsen. She reports that 65 percent

of the variation in the level of prices is common across stores belonging to the same retail chain, while only 17 percent is specific to the store and product. Additionally, because the time series variation in item prices is greater than the variation in barcode-level averages, she argues that the large shocks driving retail prices generally do not arise at the manufacturer level.

Other aspects of the data

Tables B.1 and B.2 show that downward nominal price adjustments are nearly as common as upward price adjustments, a finding that holds true whether one focuses on either posted prices or our measure of regular prices (with the caveat that the precise proportion can be sensitive to the filtering procedure used to remove transitory sales). In addition to being frequent, posted price adjustments are also large. On average, the mean weekly posted price increase is 17.6 percent whereas the mean posted price decrease is a little larger at 18.7 percent. Regular price increases and decreases are somewhat smaller, at 12.0 percent and 14.5 percent, respectively.

Table 1: Mean statistics before, during, and after the 2003–2004 Southern California supermarket strike

Statistic	On Strike			Not on Strike		
	Before	During	After	Before	During	After
Quantity index	1.000 (0.004)	0.513 (0.020)	0.904 (0.004)	1.000 (0.005)	1.336 (0.018)	1.036 (0.006)
Price index	1.000 (0.002)	1.013 (0.003)	1.012 (0.001)	1.000 (0.002)	1.021 (0.002)	1.009 (0.001)
Frequency of posted price changes	45.9 (0.4)	41.6 (0.6)	48.1 (0.3)	34.1 (0.3)	31.1 (0.4)	32.5 (0.3)
Frequency of regular price changes	20.0 (0.3)	16.1 (0.4)	19.4 (0.3)	15.2 (0.2)	13.4 (0.3)	12.7 (0.2)
Share of items on sale	23.2 (0.3)	22.0 (0.5)	25.0 (0.2)	18.1 (0.2)	16.8 (0.3)	18.9 (0.2)
Mean saving on the dollar	9.4 (0.1)	9.4 (0.2)	10.2 (0.2)	8.3 (0.1)	6.8 (0.2)	8.5 (0.1)
Share of revenues from items on sale	30.4 (0.4)	28.2 (0.7)	31.4 (0.3)	27.3 (0.4)	24.2 (0.4)	28.1 (0.3)
Share of prices equal to chain price	53.2 (0.4)	55.8 (0.5)	51.9 (0.3)	59.9 (0.3)	62.4 (0.3)	63.5 (0.2)
Addendum						
<i>Price index, IRI sample ex. striking markets</i>	<i>1.000</i>	<i>1.004</i>	<i>1.019</i>	<i>1.000</i>	<i>1.004</i>	<i>1.019</i>

Notes: The table reports the mean of the time-series statistics shown in figures 1 and 2, along with their standard deviation in parentheses, over three distinct periods. The label “before” corresponds to the 26-week period immediately before the strike, the label “during” to the 20-week period of the strike, and the label “after” to the 59-week period immediately after the strike. The price and quantity indexes are scaled such that their geometric mean equals 1 in the “before” period; all other statistics are in percent. The price index in addendum applies to the full IRI sample excluding the Los Angeles, San Diego, and St. Louis markets.

Table 2: Mean Statistics before, during, and after the 2003 St. Louis supermarket strike

Statistic	On Strike			Not on Strike		
	Before	During	After	Before	During	After
Quantity index	1.001 (0.007)	0.566 (0.060)	0.967 (0.004)	1.002 (0.011)	1.053 (0.022)	0.939 (0.008)
Price index	1.000 (0.001)	1.001 (0.002)	1.005 (0.001)	1.000 (0.001)	1.011 (0.001)	1.032 (0.001)
Frequency of posted price changes	24.7 (0.3)	22.5 (3.1)	24.9 (0.2)	22.6 (0.3)	21.1 (0.8)	22.6 (0.2)
Frequency of regular price changes	6.3 (0.3)	5.6 (1.4)	6.5 (0.2)	5.6 (0.3)	3.9 (0.3)	5.1 (0.1)
Share of items on sale	19.7 (0.2)	17.9 (0.5)	19.5 (0.1)	17.8 (0.3)	18.4 (0.1)	18.2 (0.2)
Mean saving on the dollar	9.7 (0.2)	9.1 (0.5)	10.3 (0.1)	9.6 (0.1)	10.8 (0.4)	10.8 (0.1)
Share of revenues from items on sale	28.5 (0.4)	23.5 (0.9)	28.8 (0.2)	26.5 (0.4)	28.3 (0.5)	28.1 (0.3)
Share of prices equal to chain price	94.3 (0.2)	92.7 (1.1)	94.0 (0.2)	87.5 (0.4)	87.3 (0.4)	87.7 (0.3)
Addendum						
<i>Price index, IRI sample ex. striking markets</i>	<i>1.000</i>	<i>1.001</i>	<i>1.017</i>	<i>1.000</i>	<i>1.001</i>	<i>1.017</i>

Notes: The table reports the mean of the time-series statistics shown in figures 4 and 5, along with their standard deviation in parentheses, over three distinct periods. The label “before” corresponds to the 26-week period immediately before the strike, the label “during” to the 4-week period of the strike, and the label “after” to the 75-week period immediately after the strike. The price and quantity indexes are scaled such that their geometric mean equals 1 in the “before” period; all other statistics are in percent. The price index in addendum applies to the full IRI sample excluding the Los Angeles, San Diego, and St. Louis markets.

Table 3: Mean statistics before, during, and after 2005 Hurricane Katrina in New Orleans, Louisiana

Statistic	All			Food			Housekeeping and Personal Care		
	Before	During	After	Before	During	After	Before	During	After
Quantity index	1.001 (0.008)	1.220 (0.019)	1.135 (0.009)	1.001 (0.008)	1.224 (0.021)	1.136 (0.010)	1.001 (0.010)	1.188 (0.049)	1.131 (0.007)
Price index	1.000 (0.001)	1.014 (0.002)	1.030 (0.002)	1.000 (0.001)	1.015 (0.002)	1.031 (0.002)	1.000 (0.002)	1.004 (0.003)	1.025 (0.003)
Frequency of posted price changes	30.6 (0.3)	28.9 (0.5)	30.0 (0.3)	30.9 (0.4)	29.4 (0.6)	30.3 (0.3)	28.7 (0.6)	25.6 (0.8)	28.4 (0.5)
Mean saving on the dollar	5.2 (0.1)	3.9 (0.3)	5.0 (0.1)	4.9 (0.1)	4.0 (0.4)	4.8 (0.1)	7.2 (0.4)	3.2 (0.5)	6.6 (0.2)
Share of items on sale	14.1 (0.3)	13.4 (0.6)	14.1 (0.1)	14.1 (0.3)	13.6 (0.6)	13.9 (0.2)	13.7 (0.4)	11.3 (0.8)	15.5 (0.3)
Share of revenues from items on sale	18.1 (0.4)	15.7 (1.0)	17.4 (0.2)	17.6 (0.4)	15.8 (1.1)	16.6 (0.2)	21.9 (0.8)	15.2 (1.3)	22.9 (0.5)
Addendum									
<i>Price index, IRI sample ex. Katrina</i>	1.000	1.007	1.017	1.000	1.007	1.016	1.000	1.007	1.026
<i>Price index, matched BLS sample</i>	1.000	1.008	1.013	1.000	1.007	1.008	1.000	1.010	1.033

Notes: The table reports the mean of each statistic shown in figure 6, along with their standard deviation in parentheses, over three distinct periods. The label “before” corresponds to the 26-week period immediately before Hurricane Katrina struck New Orleans, the label “during” to the 10-week federal emergency period declared by FEMA, and the label “after” to the 69-week period immediately after the lifting of the federal emergency. The price and quantity indexes are scaled such that their geometric mean equals 1 in the “before” period; all other statistics are in percent. The price index in addendum applies to the full IRI sample excluding the New Orleans and Mississippi markets.

Table 4: Effect of major snowstorms and hurricanes on retailing activities

	Quantity index (in logs*100)	Price index (in logs*100)	Frequency of price changes (percent)	Frequency of regular price changes (percent)	Mean saving on the dollar (percent)	Share of items on sale (percent)	Share of revenues from items on sale (percent)	Share of prices equal to chain price (percent)
Major snowstorms								
Quantity peak analysis								
Pre-storm average	0.00	0.00	31.96	10.77	8.91	16.97	26.86	70.02
Peak week value	12.88	0.12	31.53	10.30	8.46	16.85	25.93	70.15
t-test of zero difference	14.14	1.51	-1.36	-2.70	-3.46	-0.68	-3.11	0.37
Regression results								
Estimated coefficients								
Week before storm	0.74 (1.47)	0.09 (1.85)	-0.64 (-2.81)	-0.20 (-1.52)	-0.17 (-1.51)	-0.29 (-1.86)	-0.47 (-1.92)	0.29 (1.09)
Week during storm	11.70 (23.36)	0.00 (-0.02)	-0.33 (-1.47)	-0.49 (-3.69)	-0.32 (-2.88)	-0.03 (-0.17)	-0.73 (-3.04)	0.22 (0.84)
Week after storm	-1.85 (-3.69)	0.07 (1.28)	-0.61 (-2.69)	-0.44 (-3.27)	-0.26 (-2.34)	-0.36 (-2.29)	-0.91 (-3.78)	0.38 (1.42)
p-value, joint significance	0.00	0.20	0.00	0.00	0.00	0.05	0.00	0.33
Number of snowstorms	59	59	59	59	59	59	59	50
Observations	1,456	1,456	1,456	1,456	1,456	1,456	1,456	1,213
Hurricanes								
Quantity peak analysis								
Pre-storm average	0.00	0.00	30.19	11.05	7.22	15.73	24.08	72.91
Peak week value	12.84	0.06	30.64	11.60	7.22	16.10	24.54	71.84
t-test of zero difference	9.14	0.45	0.97	1.22	0.01	1.35	1.18	-2.33
Regression results								
Estimated coefficients								
Week before hurricane	6.44 (7.36)	-0.18 (-2.31)	-0.20 (-0.55)	0.41 (1.68)	0.09 (0.49)	-0.10 (-0.44)	-0.35 (-0.91)	-0.12 (-0.29)
Week during hurricane	6.03 (6.90)	0.13 (1.67)	0.71 (1.96)	0.82 (3.33)	-0.34 (-2.59)	-0.34 (-1.50)	-0.64 (-1.66)	-1.02 (-2.60)
Week after hurricane	-0.97 (-1.11)	0.09 (1.17)	0.99 (2.72)	1.05 (4.29)	-0.39 (-2.24)	-0.39 (-1.71)	-1.38 (-3.61)	-0.17 (-0.43)
p-value, joint significance	0.00	0.01	0.01	0.00	0.01	0.20	0.00	0.08
Number of hurricanes	21	21	21	21	21	21	21	14
Observations	476	476	476	476	476	476	476	300

Notes: For a storm beginning in week t , we define the “peak week” as the week corresponding to the observed peak in quantities over the period $t - 1$, t , and $t + 1$. Our “response at peak analysis” reports the mean across storms of each statistic during the peak week and during a “pre-storm period” covering weeks $t - 4$ to $t - 2$. The mean of the logged quantity and price indexes are normalized to zero during the pre-storm period. The t-statistics test the hypothesis that the difference between the value during the peak week and the average during the pre-storm period is zero across storms in the sample. The “regression results” are reported for a regression of each statistic on a quadratic time trend specific to each 6-month snowstorm or hurricane season and market and a set of dummies for the weeks around the storms. The p-values are for F-tests of the hypothesis that the coefficients on the week before, week during, and week after dummies are jointly different from zero. Not all markets have a sufficiently large number of stores per retail chain to allow for the computation of chain prices.

Table B.1: Price adjustments statistics

Product Category	Mean Weekly Observations	Mean Weekly Revenues	Weekly Statistics (in percent)				Monthly Statistics (in percent)			
			Frequency	Share decreases	Mean increase	Mean decrease	Frequency	Share decreases	Mean increase	Mean decrease
Beer	196,241	10,343,511	26.7	47.3	8.3	8.9	42.3	46.9	10.3	10.7
Blades	46,370	817,394	19.7	47.0	15.4	17.1	32.6	44.8	16.3	18.7
Carbonated Beverages	360,736	15,299,109	46.2	49.2	19.6	20.0	52.2	48.2	19.5	20.1
Coffee	155,018	3,607,684	27.1	47.7	16.6	17.5	42.7	45.9	17.8	18.8
Cold Cereal	274,122	8,378,019	29.8	47.6	22.4	24.5	44.0	47.3	25.4	27.7
Condiments	138,514	862,443	22.5	49.4	18.8	19.5	34.2	49.5	20.1	20.6
Deodorant	58,757	1,641,793	30.7	49.7	10.8	11.0	42.6	50.3	11.5	11.6
Diapers	35,398	1,132,845	33.1	49.1	16.9	17.4	46.1	48.5	19.4	20.2
Facial Tissue	386,859	6,777,587	39.1	49.1	22.6	23.3	54.4	49.1	24.0	24.8
Frozen Dinner	131,859	3,675,692	43.1	49.0	20.2	21.0	56.5	49.1	21.4	22.1
Frozen Pizza	82,452	968,742	20.7	48.4	15.1	15.9	36.2	47.8	17.9	19.1
Hot Dogs	52,262	2,104,729	39.0	48.1	27.0	28.7	47.9	46.9	26.3	28.8
Household Cleaners	111,315	3,665,885	31.0	48.3	18.2	19.5	44.5	48.9	19.8	20.7
Laudry Detergent	68,545	1,643,571	27.7	47.2	17.7	18.9	44.0	45.5	19.9	21.5
Margarine/Butter	47,692	1,410,520	22.9	45.6	16.8	18.8	34.8	42.2	17.6	20.6
Mayonnaise	113,907	13,604,378	26.6	46.8	10.3	11.2	40.7	45.1	11.8	13.2
Milk	64,673	980,263	19.8	47.5	14.0	14.9	31.9	45.7	16.8	18.6
Paper Towels	40,881	2,663,599	28.3	47.9	14.6	15.6	41.6	47.0	17.5	18.6
Peanut Butter	49,822	1,180,378	24.8	47.8	13.5	14.4	39.7	46.8	16.4	17.7
Razors	7,441	112,845	25.4	48.4	13.2	14.0	36.8	48.6	14.5	15.2
Salty Snack	364,113	9,503,297	31.2	48.7	19.1	19.9	42.1	48.3	20.7	21.4
Shampoo	132,914	975,987	24.9	49.3	17.6	18.4	38.6	49.6	18.7	19.3
Soup	316,278	4,445,857	23.7	47.5	20.3	22.0	37.3	46.3	22.7	24.9
Spaghetti Sauce	131,097	2,067,445	30.4	48.8	20.0	20.9	46.5	48.3	22.2	23.4
Sugar Substitutes	26,204	438,470	15.1	47.4	9.2	10.0	28.9	45.8	11.9	13.3
Toilet Tissue	53,098	4,066,702	33.6	47.8	14.8	15.8	46.4	47.1	17.6	18.7
Tooth Brushes	69,545	519,410	27.9	49.3	21.3	22.0	39.0	49.4	24.7	25.7
Tooth Paste	111,184	1,113,780	27.1	49.1	20.1	21.0	40.6	49.1	22.4	23.2
Yogurt	246,566	4,939,522	35.6	49.3	19.8	20.4	46.9	48.7	20.7	21.6
Total	3,873,862	108,941,456	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Mean	133,581	3,756,602	28.7	48.2	17.0	18.0	41.8	47.5	18.8	20.0
Weighted Mean	n.a.	n.a.	32.2	48.1	17.0	17.9	44.8	47.4	18.5	19.7

Notes: “Mean weekly observations” reports the number of nonmissing observations database. “Mean” statistics across product categories use uniform weights whereas “weighted mean” statistics use mean weekly sales as weights. Monthly statistics use observations for weeks that encompass the 15th day of each month.

Table B.2: Regular price adjustments statistics

Product Category	Mean Weekly Observations	Mean Weekly Revenues	Weekly Statistics (in percent)				Monthly Statistics (in percent)			
			Frequency	Share decreases	Mean increase	Mean decrease	Frequency	Share decreases	Mean increase	Mean decrease
Beer	196,241	10,343,511	11.4	45.1	6.6	7.3	24.0	44.1	8.5	9.1
Blades	46,370	817,394	6.4	41.2	9.1	11.4	16.6	38.6	10.8	13.8
Carbonated Beverages	360,736	15,299,109	19.8	48.2	15.5	16.2	32.5	46.9	16.4	17.4
Coffee	155,018	3,607,684	10.2	43.4	10.9	12.3	24.3	41.8	12.9	14.4
Cold Cereal	274,122	8,378,019	9.5	44.0	13.0	16.1	20.9	43.1	15.6	19.3
Condiments	138,514	862,443	4.5	48.5	11.9	13.3	12.2	48.0	13.7	15.2
Deodorant	58,757	1,641,793	9.3	49.8	7.2	7.6	20.5	51.0	8.6	8.8
Diapers	35,398	1,132,845	11.7	47.6	11.8	12.8	23.9	46.7	15.0	16.4
Facial Tissue	386,859	6,777,587	11.3	47.6	15.5	17.0	25.0	47.5	18.3	20.1
Frozen Dinner	131,859	3,675,692	13.9	47.6	14.7	16.1	28.2	47.3	17.4	18.9
Frozen Pizza	82,452	968,742	6.0	45.0	10.6	12.4	15.0	43.8	13.1	15.5
Hot Dogs	52,262	2,104,729	10.2	43.5	15.8	19.4	21.2	41.3	17.5	21.8
Household Cleaners	111,315	3,665,885	9.6	46.7	11.7	13.3	21.6	47.3	13.8	15.3
Laudry Detergent	68,545	1,643,571	9.5	41.4	10.8	12.8	21.9	39.5	13.1	15.5
Margarine/Butter	47,692	1,410,520	9.2	39.4	10.5	13.1	20.8	36.2	12.2	15.6
Mayonnaise	113,907	13,604,378	12.2	43.0	6.9	8.0	25.2	41.4	8.2	9.7
Milk	64,673	980,263	7.9	44.2	10.5	12.1	18.7	42.4	13.4	15.5
Paper Towels	40,881	2,663,599	10.3	45.0	9.8	11.2	22.1	43.8	12.7	14.5
Peanut Butter	49,822	1,180,378	9.6	44.8	9.2	10.5	21.2	43.4	11.3	13.1
Razors	7,441	112,845	6.8	46.5	9.0	10.0	16.8	46.4	10.7	11.9
Salty Snack	364,113	9,503,297	10.2	46.9	14.4	15.7	19.2	45.4	16.7	18.2
Shampoo	132,914	975,987	5.4	48.9	11.4	12.9	14.9	48.6	13.3	14.8
Soup	316,278	4,445,857	8.1	43.7	13.0	15.3	19.6	42.2	15.5	18.1
Spaghetti Sauce	131,097	2,067,445	9.7	46.5	13.0	14.7	22.8	45.8	15.5	17.5
Sugar Substitutes	26,204	438,470	5.1	42.2	6.9	8.6	13.5	40.3	9.3	11.6
Toilet Tissue	53,098	4,066,702	13.9	45.7	10.1	11.3	25.8	44.3	12.7	14.3
Tooth Brushes	69,545	519,410	10.7	48.9	16.8	17.9	18.1	48.5	19.6	21.4
Tooth Paste	111,184	1,113,780	6.6	47.0	13.1	15.0	16.4	46.8	16.0	17.9
Yogurt	246,566	4,939,522	9.3	47.2	12.3	13.7	20.5	46.3	13.7	15.5
Total	3,873,862	108,941,456	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Mean	133,581	3,756,602	9.6	45.5	11.4	13.0	20.8	44.5	13.6	15.6
Weighted Mean	n.a.	n.a.	11.8	45.6	11.7	13.2	23.8	44.6	13.7	15.4


Notes: “Mean weekly observations” reports the number of nonmissing observations database. “Mean” statistics across product categories use uniform weights whereas “weighted mean” statistics use mean weekly sales as weights. Monthly statistics use observations for weeks that encompass the 15th day of each month.

Table B.3: Basic chain prices statistics

Product Category	Posted Prices				Regular Prices			
	Equal to chain price (in percent)	Adjustement probability (in percent)			Equal to chain price (in percent)	Adjustement probability (in percent)		
		Unconditional	Conditional on chain price change	Conditional on no chain price change		Unconditional	Conditional on chain price change	Conditional on no chain price change
Beer	66.6	24.9	73.0	10.3	67.3	12.0	49.2	6.2
Blades	77.0	18.6	70.6	9.2	76.1	9.8	39.7	4.1
Carbonated Beverages	65.8	42.8	90.0	17.4	68.6	17.6	67.9	11.4
Coffee	71.5	25.1	79.4	11.9	72.3	11.9	54.1	6.2
Cold Cereal	68.1	27.2	83.7	11.8	72.0	9.7	57.9	5.4
Condiments	78.6	24.1	73.8	10.2	74.5	11.2	30.9	4.1
Deodorant	76.1	27.8	78.4	13.7	75.5	11.0	43.1	6.2
Diapers	70.0	27.4	84.4	13.2	73.1	10.4	60.5	6.3
Facial Tissue	72.8	34.6	86.5	13.6	74.8	10.7	55.1	6.4
Frozen Dinner	71.3	37.5	87.7	17.6	74.4	11.9	58.4	8.3
Frozen Pizza	74.4	19.2	72.8	9.3	75.5	8.8	44.4	3.7
Hot Dogs	72.6	35.5	90.5	13.1	79.5	9.3	69.5	5.8
Household Cleaners	72.0	29.1	81.9	13.0	73.5	10.5	49.0	6.3
Laudry Detergent	71.3	24.6	84.0	9.3	74.2	9.3	63.2	4.2
Margarine/Butter	70.2	21.0	79.9	8.7	72.5	9.6	58.2	4.6
Mayonnaise	67.6	25.1	81.6	11.6	70.4	13.3	64.8	6.1
Milk	72.3	17.9	74.7	8.9	74.0	9.1	53.6	4.1
Paper Towels	71.2	23.7	83.2	11.9	74.6	9.5	61.4	5.6
Peanut Butter	69.1	22.4	78.9	10.7	71.1	10.7	57.1	5.3
Razors	77.1	26.5	75.4	12.7	73.5	13.0	34.0	5.9
Salty Snack	76.4	29.5	89.6	12.4	81.4	10.0	67.5	6.3
Shampoo	79.9	25.4	77.0	10.8	77.1	10.3	35.5	4.5
Soup	72.2	22.9	79.8	9.7	74.4	9.6	54.7	4.4
Spaghetti Sauce	72.2	27.9	83.1	11.8	73.9	10.2	55.3	5.6
Sugar Substitutes	73.8	14.0	67.0	7.0	75.2	7.4	44.9	2.7
Toilet Tissue	65.8	29.2	85.2	13.7	69.5	12.9	65.3	7.3
Tooth Brushes	70.2	28.2	78.0	16.1	68.7	16.0	48.7	8.6
Tooth Paste	77.8	24.6	78.6	11.1	77.0	9.2	40.4	4.5
Yogurt	72.6	32.3	87.5	12.0	76.9	8.8	61.5	5.1
Mean	72.4	26.5	80.6	11.8	73.8	10.8	53.3	5.7
Weighted Mean	71.2	28.4	82.3	12.3	73.2	11.3	58.0	6.4
Observations	n.a.	1,193,827,794	838,593,893	355,233,901	n.a.	1,193,827,794	1,058,562,164	135,265,593


Notes: We define a weekly “chain price” as the modal posted price of a barcode across stores belonging to the same retail chain and IRI market. We require a minimum of five store prices to compute a chain price. We exclude an item’s own price from the computation of the chain price before computing its adjustment probability conditional on a chain price adjustment or conditional on no chain price adjustment.

Figure 1: Impact of 2003-2004 Southern California supermarket strike on quantities and prices




Notes: The shaded areas indicate the strike period and the tick marks correspond to the first day of each month. All series combine data from the San Diego and Los Angeles markets and, with the exception of the price ratios, are seasonally adjusted. The series labeled "on strike" and "not on strike" correspond to stores that experienced a drop in revenues of 10 percent or more and an increase in revenues of 10 percent or more, respectively. The price and quantity indexes are scaled so that their geometric mean equals 1 in the 26-week period immediately before the strike. See appendix A for methodological details.

Figure 2: Some key pricing features of the 2003-2004 Southern California supermarket strike




Notes: The shaded areas indicate the strike period and the tick marks correspond to the first day of each month. All series combine data from the San Diego and Los Angeles markets and seasonally adjusted. The series labeled "on strike" and "not on strike" correspond to stores that experienced a drop in revenues of 10 percent or more and an increase in revenues of 10 percent or more, respectively. The price and quantity indexes are scaled so that their geometric mean equals 1 in the 26-week period immediately before the strike. See appendix A for methodological details.

Figure 3: Importance of discounts before, during, and after the 2003-2004 Southern California supermarket strike




Notes: The panels show the distribution of discounts measured in percent of the regular price identified by our sales filter. The discounts are aggregated in 10-percentage-point bins, starting with discounts in the interval (0,10] percent, then (10,20] percent, and so on. Observations in the top panels are weighted by each item's contribution to total revenues (that is, $\omega_{i,t} = p_{i,t}q_{i,t} / \sum_j p_{j,t}q_{j,t}$), so that the area under the curve integrates to the share of total revenues derived from items on sale. Observations in the lower panels are weighted by each item's total dollar savings over the period (that is, $\omega_{i,t} = (p_{i,t}^{reg} - p_{i,t})q_{i,t} / \sum_j p_{j,t}^{reg}q_{j,t}$), so that the area under the curve integrates to the average discount offered by retailers. All statistics combine data from the San Diego and Los Angeles markets and do not control for seasonality. The series labeled "on strike" and "not on strike" correspond to stores that experienced a drop in revenues of 10 percent or more and an increase in revenues of 10 percent or more, respectively. The label "before" corresponds to the 26-week period immediately before the strike, the label "during" to the 21-week period of the strike, and the label "after" to the 58-week period immediately after the strike.

Figure 4: Impact of 2003 St. Louis supermarket strike on pricing decisions




Notes: All series are seasonally adjusted. The shaded areas indicate the strike period. The price and quantity indexes are scaled so that their geometric mean equals 1 in the 26-week immediately period before the strike. The series labeled "on strike" and "not on strike" correspond to stores that experienced a drop in revenues of 10 percent or more and an increase in revenues of 10 percent or more, respectively. See appendix A for methodological details.

Figure 5: Impact of 2003 St. Louis supermarket strike on pricing decisions




Notes: All series are seasonally adjusted. The shaded areas indicate the strike period. The price and quantity indexes are scaled so that their geometric mean equals 1 in the 26-week period immediately before the strike. See appendix A for methodological details.

Figure 6: Impact of 2005 Hurricane Katrina on pricing decisions in New Orleans, Louisiana



Notes: All series are seasonally adjusted. The price and quantity indexes are scaled so that their geometric mean equals 1 in the 26-week period immediately prior to the strike. The shaded areas indicate the period (August 29, 2005, to November 1, 2005) covered by FEMA's major disaster declaration. See appendix A for methodological details.

Figure 7: Illustration of snowstorms effects: winter 2009–2010 in Washington, D.C.



Notes: The leftmost shaded areas mark the December 18–19, 2009, nor'easter while the rightmost shaded areas mark a pair of blizzards that hit the U.S. capital on February 5–10, 2010. The price and quantity indexes are scaled so that their geometric mean equals 1 in November 2009. All series are seasonally adjusted. See appendix A for methodological details.