A Sticky-Price View of Hoarding *

Christopher Hansman
Imperial College London

Harrison Hong
Columbia University

Aureo de Paula
University College London

Vishal Singh
New York University

June 17, 2023

Abstract

Hoarding disrupts the functioning of markets. Yet little is known about its determinants. We analyze major recent hoarding episodes through the lens of an optimal inventory model in which risk-averse agents hoard both as a precautionary hedge against price uncertainty and to speculate when prices are predictable. Using supermarket scanner data, we provide reduced-form evidence of the importance of the speculative motive due to sticky retail prices. We use our model to quantify that speculation accounts for a meaningful fraction of overall hoarding, although smaller than precaution in our episodes.

*This paper subsumes our earlier work on hoarding and commodity bubbles. We thank seminar participants at Shanghai University of Finance and Economics, CICF 2019, INSEAD, Aalto, Peking University, Warwick University, Cambridge University, NBER Universities Conference on Commodities, HKUST, New York University, PUC-Rio and Yale University for helpful comments. We are also grateful to Neng Wang, Quisha Peng, Pengfei Wang, Yuriy Gorodnichenko, Emi Nakamura, Hassan Afrousi, Michael Woodford, William Goetzmann, Hank Bessembinder, Manuel Arellano, Orazio Attanasio, Richard Blundell, Marcelo Fernandes, Bo Honoré, Guy Laroque, Valerie Lechene and Elie Tamer for useful conversations, and to Zhenhao Tan for excellent research assistance. de Paula acknowledges financial support from the European Research Council through Starting Grant 338187 and the Economic and Social Research Council through the ESRC Centre for Microdata Methods and Practice grant RES-589-28-0001.

1 Introduction

Hoarding is thought to disrupt market functioning by exacerbating shortages.[1] During the COVID-19 pandemic, stockpiling by households and firms was reported in a range of markets—from household staples to semiconductors—distorting prices and creating bullwhip effects along the supply chain (Shin, 2021).[2] Similar issues are well documented during the energy crisis of the 1970s, the commodity price boom in the 2000s, and numerous other episodes.[3] Despite the economic importance of hoarding, there has been relatively little work examining its determinants. We address this issue by analyzing household hoarding using high-frequency supermarket scanner data that covers prominent recent stockpiling incidents in the US.

Our approach is based on a simple optimal inventory model of a risk-averse (mean-variance) agent with storage costs. This model nests two incentives for stockpiling. The first is a hedging motive driven by price uncertainty. Hoarding today provides insurance against coming price fluctuations for a staple good. This precautionary motive — or its close cousin panic buying — forms the prevailing popular narrative of retail hoarding.[4]

The second is a speculative motive driven by the expected level of prices. If price increases are partially predictable, even risk averse agents face an incentive to speculate and cheaply secure their own supply (Scheinkman & Schechtman, 1983; Deaton & Laroque, 1992). While this mechanism has received less attention as a determinant of retail hoarding—perhaps because households are not thought to have an informational advantage over market prices in standard models—our framework allows us to derive a tractable test to quantify the relative importance of precaution and speculation.

The distinction is relevant for firms and policymakers. As the widespread adoption of anti-price gouging measures demonstrates, precautionary stockpiling is a salient concern for policymakers (see, e.g., Zwolinski, 2008; Giberson, 2011).[5] Whether anti-price gouging regulations are

[1]See, e.g., Sen (1983) on the role of hoarding in the Bengal famine of 1943 or Priest (2012) on the 1970s energy crisis.

[2]On consumer and firm behavior during the pandemic, see, e.g. “Coronavirus: The psychology of panic buying,” *BBC*, March 4, 2020, "Are bottlenecks bullwhip?", *Financial Times*, December 9, 2021, and “Carmakers hoarding semiconductors like 'toilet paper' risk prolonging the chip shortage”, *Euronews.next*, July 13, 2021.

[3]On the energy crisis, Priest (2012) writes: "Motorists, whose consumption of gasoline rose from 243 gallons per capita in 1950 to 463 gallons per capita in 1979, compounded supply problems by hoarding fuel, idling their engines in gas lines, and frantically topping off their tanks with frequent trips to the local filling station."

[4]For example, see Dawe (2012) and "How Fear Turned A Surplus into Scarcity," *NPR*, November 4, 2011 on the commodity crisis.

[5]The first US state law directed at price gouging was enacted in New York in 1979, in a period of commodity market

welfare enhancing is far from clear, particularly if consumers are driven by speculation. Unless cost increases are truly transitory, price caps can exacerbate speculative stockpiling and erode firm profits.

Our paper has three parts. In the first, we argue that a standard retail phenomenon—sticky shelf prices—makes the speculative motive present for a wide range of hoarding episodes. A large literature documents sluggishness in price adjustment (Nakamura & Steinsson, 2008; Kehoe & Midrigan, 2015), even after the cost shocks or disasters that typically precede consumer stockpiling (Cavallo *et al.*, 2014; Gagnon & Lopez-Salido, 2015). If households recognize that prices are sticky and likely to rise in the future, they will be motivated to shift demand intertemporally and stockpile storable goods.

To do so, we provide reduced form evidence that a sticky-price based speculative motive drove household purchases during a major recent episode: the 2008 global rice crisis. The crisis was sparked by a supply shock (from a US perspective)—a ban on Indian rice exports in the fall of 2007. This shock led rice prices in commodities markets to rise by roughly 300 percent, peaking in April and May of 2008. Retail prices did not catch up until later in 2008. The time series of consumer purchases shows that households hoarded at relatively low cost before shelf prices rose.

The cross-section of hoarding behavior further confirms the presence of a speculative motive. While prices were slow to adjust for nearly all rice products during the crisis (and in nearly all stores) there is considerable dispersion in the size of later price increases. These differences were generated by heterogeneity in the exposure of products to the wholesale cost shock, ex-ante pricing policies, and more. We show that consumers stockpiled rice most in the products and stores that later saw the greatest increases. These patterns do not appear to be the spurious result of reverse causality (retailers pricing in response to excess purchases), differences in price levels during the crisis itself, or hoarding acting as a signal of future demand.

Since households are likely to hoard due to *both* speculative and precautionary motives— in the second part of our paper—we use our model to derive a forecast test that quantifies their relative importance. The intuition underlying our two-step test is that a household's purchases encode their ex-ante beliefs about prices. As a result, the covariance between stockpiling and ex-post realized price changes reflects the degree to which consumers are driven by a speculative motive.

instability. These laws have proven popular—most states now have some type of regulation (Davis *et al.*, 2008).

The challenge is translating quantities purchased into a meaningful statement about future prices.

The basic insight of our first step is that the consumer response to standard, non-crisis, retail sales provides a benchmark for beliefs. Excess quantities purchased during a typical retail promotion capture households' sensitivity to known and certain future price increases.[6] By measuring consumer stockpiling during these promotions, we can back out the elasticity of purchases to future price changes under perfect foresight (which in our model is governed strictly by storage costs). Assuming this elasticity is stable, we are able to use observed stockpiling during the crisis—across products, brands and stores—to recover what we call risk-neutral price forecasts. These represent consumers' implied price expectations if all excess purchases were driven by a speculative motive.

In the second step of our test we estimate a linear regression of ex-post *realized* price changes on these forecasts. This is reminiscent of efficient forecast tests in the spirit of Mincer & Zarnowitz (1969); Nordhaus (1987); Keane & Runkle (1990) utilizing cross-sectional rather than time-series variation. The slope of the regression of realized price changes on expected price changes captures the efficiency of the price forecasts we recover. The model makes direct predictions on the coefficient from this regression. If all hoarding is generated by a speculative motive (e.g., because of risk neutrality) and consumers have rational expectations, we expect a slope of 1. The presence of a precautionary motive drives the coefficient below 1, tending to 0 as this incentive dominates. Across products and stores, we estimate a statistically significant slope of just over 0.1, indicating that speculation played a meaningful but smaller role than precaution in the hoarding episode we study.[7]

In the third and final portion of our paper, we re-apply our test to consider rice hoarding during the recent COVID-19 episode. The setting is a bit more complex than the 2008 rice crisis—consumers faced both changing prices and new, COVID specific demand shocks—and comprehensive data is not yet available. However, we are able to access a more limited scanner dataset that covers the first months of the pandemic. We find very similar qualitative patterns hold in this recent episode, albeit with a smaller role for speculation. These results complement early work

[6] For example, consumers are aware that a good marked down temporarily from \$100 to \$80 will experience a 25% increase when the promotion ends.

[7] Since classical measurement error may downward bias this slope, we view our estimate as a lower bound on the contribution of the speculative sticky-price motive to hoarding.

using Google search data (Keane & Neal, 2021).

Beyond providing a systematic study of the drivers of consumer hoarding, our paper contributes to a couple of other literatures. We provide a method to address a long-standing challenge in the commodities literature of distinguishing between speculative versus hedging demand (Fama & French (1987), Bessembinder (1992), Gorton *et al.* (2013), Kang *et al.* (2020)). Our test and empirical findings also contribute to a large body of research on the role of speculation in the 2000s commodity bubble (Kilian & Murphy, 2014; Hamilton, 2009; Tang & Xiong, 2010; Singleton, 2013; Acharya *et al.*, 2013). Finally, we provide evidence on the role of household speculation in a sticky-price or menu cost setting (Barro, 1972; Sheshinski & Weiss, 1977)), a long-standing issue in that literature that goes back to Benabou (1989). Such speculation informs specific firm-level cost to sticky-pricing policies (Gorodnichenko & Weber, 2016).

Our paper is organized as follows. In section 2 we describe our data. In section 3 we provide an overview of the 2008 rice crisis and show time series evidence on the presence of a speculative motive. In section 4 we provide evidence based on the cross-section of products and stores. Section 5 introduces our model and outlines our forecast test. Section 6 provides a discussion of hoarding in the COVID-19 crisis. We conclude in section 7.

2 Data

Our primary sources of household and store-level data are the Nielsen retail scanner and consumer panel datasets held at the Kilts center. In both, we consider weekly data from 2007-2009 and limit the sample to packaged and bulk rice products.[8]

Retailer Data

For our store-level data we consider food retail channels only. This leaves us with 10,561 unique stores. The data contain weekly store-UPC (Universal Product Code) level prices and quantities sold, as well as product and store characteristics. In various parts of our analysis, we consider aggregated store level data, store-UPC level data, and store-brand level data. To avoid rarely bought products when considering UPC or brands, much of our analysis restricts to store-UPC or

[8]UPCs with product module code 1319.

store-brand pairs with at least 5 units per week sold on average in 2007. This restriction leaves us with 71,952 store-UPC pairs representing 547 unique UPCs for our store-UPC level data and 43,953 store-brand pairs representing 154 unique brands for our store-brand level data.

When considering aggregate store level data, our primary quantity measure is the total volume of rice sold across all UPCs (measured in ounces). Our primary price measure is the sales weighted average price per 80 ounces across all UPCs. Our results are robust to alternative price definitions, for example considering equal weighted prices, fixing 2007 sales weights, or considering only the price of the most popular UPC within each store. Panel A of Table 1 presents summary statistics on store level aggregates. In our sample, the average store sold just over 8000 ounces of rice per week, with an average price of $5.39 per 80 ounces.

Panel B of Table 1 provides summary statistics for store-UPC level data. The average UPC represented in our data contains 52.6 ounces of rice (16, 32 and 80 ounces are all common sizes). 11.5 units per week were sold on average, representing just under 700 ounces. The average price per 80 ounces was $5.58. Panel C provides summary statistics at the store-brand level. On average 27 units were sold per week at the store-brand level, representing just under 1700 ounces.

Consumer Data

The consumer panel covers between 40,000-60,000 demographically balanced U.S. households each year who use hand-held scanners to record every bar-coded grocery item purchased. The broader dataset records every purchase made at the UPC level. There is also detailed demographic information. Appendix Figure A.I plots the distributions of various demographics of the Nielsen Panel. We restrict the panel to households we observe purchasing packed or bulk rice products at least once between 2007-2009. This leaves us with 42,441 unique households. Panel D of Table 1 presents summary statistics on this restricted household sample. The average quantity purchased by a household in a given week is just under 2 ounces, although households typically purchase about 72 ounces in weeks when they purchase.

Overall, the households in our data are similar to the general population in terms of income. The median household in our data earns $50,000-60,000 per year. The median for all US households in 2008 was $51,726 (Noss, 2010). Households in our data appear to be slightly better educated than the general population—roughly 53 percent of our sample has a college education or higher, compared to the 38 percent of adults over 25 reported to have an associates or bachelors degree in 2008.[9] Finally, our sample also has slightly lower fraction of Asian households: 3.4 percent of our sample is Asian, lower than the 5.6 percent reported in the 2010 census.

We also construct a balanced household-brand level panel dataset that considers only rice purchases at stores that also appear in the Nielsen retail scanner data. This panel contains just under 18,000 households purchasing 168 unique brands at 8,194 different stores. While this restriction meaningfully limits the set of households, it allows us to construct a consistent store-brand level price series to capture the prices faced by consumers in weeks they did not purchase rice. For this panel we define the store-brand level price as the equal weighted price per 80 ounces across the UPCs sold in that week for that brand and store. In the limited cases in which no UPC for a brand was sold in a given store and week, we impute using the price in the store in the previous week.

3 Aggregate Time Series Evidence on Sticky Store Prices and Hoarding

3.1 Overview of the Rice Crisis and Commodity Price Dynamics

Rice commodity prices skyrocketed in mid-2008. The price increase was accompanied by unrest in Haiti, Bangladesh, and elsewhere in the developing world, surges in purchasing globally, and new restrictions to ensure domestic supply for a number of exporters.[10] These events received widespread media attention in the US and across the world.

Retrospective overviews highlight political factors as the key trigger for the 2008 global rice crisis. According to Dawe & Slayton (2010) and Slayton (2009), the crisis began with India's electioneering driven 2007 ban of rice exports, was compounded by restrictions in Vietnam and elsewhere, and continued until Japan agreed to release rice reserves to global markets in mid-2008.[11] While the late 2000s saw instability in energy and other food commodity prices, the political nature of the rice crisis meant that spikes in rice prices had a "fundamentally different explanation" in comparison to fluctuations in the price of other major cereals (Dawe, 2012).

Figure 1 displays the dynamics of commodity prices during the crisis. The solid black line

[9] See the U.S. Census Bureau, Current Population Survey, 2008 Annual Social and Economic Supplement.

[10] See, e.g. `https://www.cnn.com/2008/WORLD/americas/04/14/world.food.crisis/` for contemporary coverage of unrest and Childs *et al.* (2009) on export restrictions.

[11] A World Trade Organization agreement had mandated that Japan import US rice while limiting re-export, generating significant stock in Japan. The US publicly provided permission to re-export in mid May of 2008.

shows a proxy for the global price of rice on commodities markets from the IMF, highlighting the crisis and associated events.[12] A sharp increase is evident following the first vertical line (a peak of around $1000 per metric ton), which represents the Indian ban on exports in October 2007, as is a correction following the second vertical line, which represents the late May 2008 news of Japan's agreement to release reserves. Even with this correction, the global price converged to a level well above the pre-ban average, rising from $332 per metric ton on average in 2007 to $589 per metric ton on average in 2009, a nearly 80 percent increase.

3.2 Sticky Store Prices

Despite the massive increase in commodity prices, prices on the shelf in the US were sticky—basically unchanged—for nearly all retailers through the peak of the crisis. Figure 2 displays the fraction of stores that updated retail prices in the wake of the shock to commodity prices. While there is no standard definition of price adjustment, we take what we believe to be a relatively conservative approach. We define a store to have updated if its price is greater than 125 percent of its 2007 average price.[13] As commodities prices rose through the beginning of 2008, a very small fraction of stores updated prices according to our metric. Even this limited fraction appears to be on trend with regular (and gradual) price increases relative to 2007.

Notably, the large majority of stores failed to update prices through the weeks of April 19th-May 10th (highlighted in gray), which we refer to as the *hoarding period*. This interval just before the Japan agreement represents the most intense period of the crisis, in which commodity and wholesale prices hit their peak, and—as we shall see in the next subsection—the most aggressive consumer hoarding took place. In the weeks following the hoarding period, stores updated rapidly to match the long run increase in commodities prices: half updated within a few weeks and more than 75 percent updated within a few months.

Price stickiness is perhaps more easily observed in Figure 3, which compares the dynamics of US wholesale prices and retail shelf prices. The black line captures a proxy for *US wholesale prices*, which largely track international rice prices—rising through the beginning of 2008 and peaking in

[12] This line presents the rice series from the IMFs Primary Commodity Price System. The series measures the Thailand nominal price quote for 5 percent broken milled white rice in USD per metric ton and is available at `https://www.imf.org/en/Research/commodity-prices`.

[13] While this threshold is somewhat arbitrary, we get similar patterns when we consider different cut-offs, e.g. 110 percent.

the hoarding period (shown in gray).[14] Alternatively, retail shelf prices from our store level data, shown in blue, did not rise at all with wholesale prices, staying flat or even declining slightly until after the peak of hoarding.[15] After the hoarding period, shelf prices increased to and stabilized at a higher price, mirroring long run commodities price dynamics. The average price in the post hoarding period was 35 percent above the average price in the pre-hoarding period. Relative to the 80 percent increase in wholesale prices, this represents a pass-through of just over 40 percent, on par with findings for other storeable goods (Leibtag, 2007). These patterns are consistent with a large literature in macroeconomics: the retail or supermarket prices that consumers face are sticky and tend to lag changes in commodities prices.[16]

3.3 Consumer Hoarding Anticipated Retail Price Changes

We now show evidence of the key phenomenon of our paper: severe consumer hoarding during the rice crisis. The red line in Figure 3 shows the pattern of total quantity sold by stores in our sample, which spiked sharply during the crisis and reached its highest point in the hoarding period (April 19th-May 10th). The average store sold over 11,000 ounces of rice per week during this hoarding period compared to an average of just under 8,000 ounces in all other weeks of our sample. The most intense week featured average purchases that were more than 65 percent above average. Notably, this increase in store sales roughly coincided with or slightly followed the peak of global commodities prices. Similar or even more severe consumption patterns were noted internationally.[17]

The key pattern displayed by this figure is the timing of the spike in consumer purchases relative to the increase in retail prices. Effectively all excess purchases occurred in the weeks before store prices began to increase. As retail prices began to rise in mid-May, quantity sold returned to

[14] The proxy is based on the average f.o.b. price for long grain rice at selected milling centers in Southwest Louisiana. Data provided by USDA, based on data from Agricultural Marketing Service, National Weekly Rice Summary. We scale the series by its mean over the sample period

[15] We similarly scale this series by its mean over the sample period. One potential concern is that the observed delay in adjustment of our price index might be an artifact of consumer substitution across types or qualities of rice. For example, if retailers increased all rice prices but consumers responded by substituting to the cheapest products, the two effects might cancel out in our aggregated price index. To address this, Appendix Figure A.II replicates Figure 3 but includes a measure of prices that holds product types fixed. In particular, this figure shows the equal weighted average across all UPC-store pairs that appear consistently throughout our sample.

[16] Although McShane *et al.* (2016) show that a meaningful fraction of positive wholesale price changes are eventually passed on to consumers.

[17] See, e.g. https://www.reuters.com/article/uk-philippines-rice/philippines-arroyo-leads-crackdown-on-rice-hoarding-idUKMAN1898020080508.

levels similar to or slightly below those in the pre-crisis period. Appendix Figure A.III shows a similar pattern for purchases in our household sample: household inventories peaked in the hoarding period, prior to any meaningful increase in retail prices. Overall, consumer hoarding during the rice crisis coincided with or slightly lagged commodity and wholesale prices, but anticipated the rise in retail prices.

The basic patterns in commodity, wholesale and retail price dynamics, as well as in household and store sales, are summarized and quantified in Table 2. This table presents regressions of the time series of (i) IMF commodity prices, (ii) US wholesale prices, (iii) average household quantities purchased and prices paid, and (iv) average store level quantities sold and prices charged, on an indicator equal to one in the hoarding period. For (i) and (ii), which are monthly, the hoarding period is defined as April and May of 2008. The remaining series are weekly, and the hoarding period is defined as the weeks of April 19th-May 10th. Price time series are constructed as sales weighted across products within households or stores, and equal weighted across households or stores. As shown in the Figures discussed above, commodity prices, wholesale prices, and quantities sold were significantly above average during the hoarding period, while retail prices were not. In fact, because of the high retail prices in the post-hoarding period, the coefficient on the hoarding period indicator for both household and retail prices is negative.

3.4 Retailer and Consumer Awareness

Given the extent of press coverage, producers and stores were likely aware of the wholesale rice price increase—a cost shock from their perspective—emanating from the India ban. Even without the media, retailers could easily have aggregated this information from wholesale prices and rice futures. Both the global commodities price (shown in Figure 1) and rice futures for May, July, and September of 2008 (shown in Appendix Figure A.IV) rose steadily through the first months of 2008 before reaching a high in April. Prices for all three futures contracts peaked on April 23rd, in the midst of consumer hoarding. At this point, July futures prices exceeded May prices, suggesting that the market anticipated prices remaining high and even rising over the course of the next several months. Put simply, it is doubtful that sticky retail prices were the result of an information gap. Retailers could easily have recognized that prices were rising in the beginning of 2008, and—

at the peak—should reasonably have expected prices to remain high for several months.[18]

Similarly, there appears to have been heightened consumer awareness of the rice crisis. While it is unlikely that the average consumer closely tracks wholesale or rice futures prices, the blue line in Appendix Figure A.V shows a notable spike in Google searches for the term "Rice" during the hoarding period. This figure presents a search volume index representing the weekly intensity of Google searches in the US between 2007-2009 (normalized by the average over the sample period). The IMF commodities price and quantity sold at the store level are included for comparison in black and red, respectively. This elevated search volume was likely prompted by media reports and suggests higher than typical consumer attention on the rice market.

3.5 The Speculative Motive Generated by Sticky Prices

Given of the limited connection between rice price dynamics and more general commodity price fundamentals, prevailing views of the rice crisis suggest that panic or precaution driven hoarding generated artificial shortages and exacerbated the price shock.[19] Classic narratives along these lines have consumers panicking just as prices skyrocket, leading them to purchase large quantities at high prices.[20] A simple comparison of consumer purchases and commodities prices in our episode would support this view. Excess purchases were concentrated at the very peak of commodities prices.

However, excess purchases preempted any change in the *retail* prices consumers actually faced. Consumers hoarded while shelf prices were low, before they rose to a permanently higher level. This is consistent with an alternative mechanism driving hoarding: the speculative motive generated by sticky prices. The logic of this mechanism is simple: if there is a shock to wholesale prices, but retailers are slow to respond, consumers have an incentive to build up inventories of a storable good like rice before shelf-prices rise. The implicit discount generated by sticky prices—relative to a sustainable long run price—will cause consumers to shift demand dynamically and stock up, just as they would when facing a standard retail promotion or sale.

[18] Futures prices show daily close prices for rice futures with expiration in May 2008, July 2008 and September 2008 from the Chicago Mercantile Exchange. The futures contract is for 2,000 cwt (hundredweight), which corresponds to about 200,000 pounds or circa 91 metric tons, of rough rice, no. 2 or better, and the price quote is in cents per hundredweight.

[19] Much coverage of the episode emphasizes a precaution or fear narrative ("How Fear Turned A Surplus into Scarcity," *National Public Radio*, November 4, 2011 and "A Run on Rice in Asian Communities," *New York Times* May 1, 2008).

[20] See, e.g., https://spectrumlocalnews.com/nc/charlotte/news/2021/05/11/higher-prices--panic-buying--what-the-colonial-pipeline-hack-means-for-north-carolina

Retail prices in the hoarding period were approximately 22 percent below the post-hoarding average price. If consumers were forward looking and interpreted prices during the hoarding period as a 22 percent discount, standard promotional elasticities from the literature (and those that we estimate in Section 4) suggest that a speculative motive could explain a substantial fraction of the excess purchases that took place during the crisis.[21]

Of course, the fact that consumer hoarding came before the increase in retail prices is not conclusive evidence of a speculative motive at work. It is possible that consumers were driven by some form of panic or precaution—perhaps after hearing about potential shortages on the news—and that the subsequent growth in aggregate prices was coincidental. To rule out this possibility, the next section of this paper shows that speculation also appears to have driven hoarding in the cross-section of products and stores.

4 Cross-sectional Evidence on Sticky Store Prices and Hoarding

In this section we develop a series of tests using the cross-section of excess purchases and subsequent price changes for different rice products, brands, and stores to distinguish a speculative motive from panic or precaution. As long as consumers have some information about the coming cross-section of price changes—due to knowledge of the exposure of a particular product to the shock, familiarity with the pricing patterns of a particular store, or attention to pre-crisis price increases—excess purchases should predict later price increases. That is, under a speculative motive, we should expect more hoarding in the products or stores with the largest subsequent price increases. Alternatively, if hoarding is strictly driven by a precautionary hedging motive there should be no relationship between hoarding and future price changes in the cross-section.

4.1 Dispersion in Price Changes

We begin by showing that there was substantial cross-sectional dispersion in price changes following the hoarding period. While prices increased for nearly all products after the crisis, the size of this increase varied significantly. Cross-sectional dispersion is the result of a number of factors,

[21] For example, Bergtold *et al.* (2004) estimate unconditional price elasticities of roughly -1 that do not factor in the additional dynamic incentive provided by temporary sales.

including differential exposure to the underlying cost shock, different degrees of ex-ante price adjustment (certain products had already experienced some amount of regular updating prior to the crisis), different ex-post pricing policies, and more.

Figure 4 displays this cross-sectional dispersion, showing histograms at various levels of aggregation. We present price changes as percentage increases relative to the hoarding period. Specifically, if we let $\bar{p}_i^h$ represent the average price of unit i during the hoarding period, and $\bar{p}_i^a$ represents the average price in the post-hoarding portion of our sample (from May 10th, 2008 to the end of 2009). We define

$$\text{Post-Hoarding Price Growth}_i = 100 \times \left(\frac{\bar{p}_i^a - \bar{p}_i^h}{\bar{p}_i^h} \right).$$

For example, if a price was \$8 during the hoarding period and \$10 in the post-hoarding period, this would translate to a realized increase of 25 percent.

In Panel A, we report a histogram of realized increases for store-UPC pairs in our sample. The median change is over 20% and nearly all store-UPCs saw an increase. There is considerable dispersion: the 25th percentile is below 15 percent while the 75th percentile is nearly 40 percent. The same patterns holds for the histogram of price discounts for store-brand pairs shown in Panel B, and for the set of stores and brands shown in Panels C and D. This substantial cross-sectional dispersion forms the basis for our tests.

4.2 Cross-Sectional Predictability

Consistent with a sticky-price motive, we next show that consumer purchases predict price increases in the cross-section. Consumers hoarded more aggressively in products that later saw larger price increases. This suggests that consumers were aware that prices were sticky and were able to predict the products that would experience the most significant price growth.

Hoarding Concentrated in Products With Large Post-Crisis Price Increases

We begin with graphical evidence that hoarding was concentrated in the products that later experienced large price increases. We categorize a store-UPC to have had a high post-hoarding price increase if it is in the top quartile of increases among all store-UPC pairs, and to have a low post-hoarding price increase if it is in the bottom quartile. Figure 5 plots quantities sold and price-per-ounce over our sample period for each of these two groups. All series are normalized by the group average over the full sample. The top panel shows that purchases during the hoarding period were more than double the average for store-UPCs in the top quartile of post-hoarding price increases. On the other hand, there was very little excess purchasing for store-UPCs in the bottom quartile of post-hoarding price increases. This is consistent with consumers being driven by a speculative motive.

Cross-sectional regression evidence

We now provide a more direct evaluation of this pattern—consumers hoarding more in the products that later increased prices sharply—at the product, brand, and store levels. We conduct our approach by regressing realized post-hoarding price growth on quantity increases *during* the hoarding period. Specifically, if Quantity Growth_i is the percentage increase in average weekly sales during the hoarding period relative to average weekly sales in the pre-hoarding period, we consider: [22]

$$\text{Post-Hoarding Price Growth}_i = \alpha + \lambda \text{Quantity Growth}_i + \varepsilon_i. \tag{1}$$

Our basic test asks whether the degree of hoarding—captured by Quantity Growth_i—contains information about coming price growth, that is, whether $\lambda > 0$. If consumers are forward looking, have information regarding coming price changes, and are motived by the speculative incentive generated by by sticky prices, then excess purchases should positively predict future prices, on average.

We present results of our tests in Table 3 and Figure 6. These results show estimates of Equation 1 at the store-UPC, store-brand, store, and brand levels. The binned scatterplots shown in Figure 6 present perhaps the most striking evidence of cross-sectional predictability. The relationship between forecasted discounts and realized discounts is roughly monotonic and close to linear for

[22] If q_i^b represents average weekly quantity sold before the hoarding period and q_i^h represents the average weekly quantity sold during the hoarding period, we define

$$\text{Quantity Growth}_i = 100 \times \left(\frac{\bar{q}_i^h - \bar{q}_i^b}{\bar{q}_i^b} \right).$$

We find similar results when considering excess purchases relative to weekly sales in the post-hoarding period or relative to the period as a whole. We focus on the pre-hoarding period to avoid purchase decisions influenced by inventories built up while hoarding.

store-UPCs, store-brands and stores. At each point in the distribution, consumers purchased more of the products that later experienced greater price increases. This suggests that consumers were aware of coming price changes, and were motivated to stock-up more for prices that later increased the most.

Table 3 presents the results of our regressions, which confirm a strong relationship between hoarding intensity and price growth. Columns 1-3 display our store-UPC level analysis. The first column presents the specification shown in Equation 1 exactly, and shows a positive and highly significant coefficient. The magnitude suggests that products that experienced 10 percentage points more hoarding during the crisis later saw increases in price that were one percentage point higher than other products. Columns 2 and 3 confirm that this results holds when including store fixed effects, and the combination of store and UPC fixed effects. This suggests that the cross-sectional relationship between hoarding and price increases is not simply an artifact of hoarding in certain specific stores or UPCs.

Columns 4-6 display our store-brand level analysis, and show similar results. Columns 7 and 8 show results aggregated to the store level and aggregated nationally to the brand level, respectively. In each of these cases we continue to see a strong and positive relationship between quantity growth during the hoarding period, and post-hoarding price growth. The takeaway is straightforward: consumers hoarded most in the individual products, brands, and stores that later saw the greatest price increases.

We interpret this pattern as evidence of a speculative motive at work. Because of prices stickiness, consumers were aware of coming prices changes during the hoarding period. They acted on this information by stocking up on products or in stores where they expected shelf prices to increase substantially. However, there are a series of alternative explanations that might generate similar patterns. We next discuss these and provide evidence in favor of the speculative motive.

Addressing Alternative Explanations

Lower prices during the crisis One alternative explanation for our cross-sectional evidence is that products (or stores) with particularly low prices during the hoarding period were differentially likely to have large price increases in the post-hoarding period. This could be, for example, because the product or store had not recently updated prices before the crisis. If consumers excessively purchased low price products during the crisis, this might generate the observed correlation. Panel A of Table 4 shows that this is not the case. This table repeats the analysis in Table 3, but explicitly includes a control for the level of average unit prices during the hoarding period. We see a significant negative coefficient on the price level, but the relationship between hoarding intensity and shelf price growth is effectively unchanged across specifications.

Reverse causality

An additional concern is that the correlations shown in Table 3 might reflect a causal channel, rather than consumer expectations about coming price increases. Specifically, that stores changed prices as a result of the degree of excess purchases during the crisis. Perhaps the most compelling evidence against this concern is the permanence of shelf price increases following the hoarding period. As Figure 3 shows, the aggregate increase in shelf prices persisted through the end of our sample period. A similar pattern holds at the disaggregated level. Products and stores that saw large price increases in the period immediately-post hoarding continued to have relatively high prices towards the end of our sample. This is likely because price changes reflected real changes in costs rather than a response to transitory hoarding purchases.

To highlight this point, Panel B of Table 4 presents an alternative version of Table 3 in which we consider the very long run change in prices. Specifically, we redefine our measure of post-hoarding price increase to be the percentage increase in the price when comparing the last week of our sample (the last week of 2009) to the average in the hoarding period. The logic behind this exercise is that price responses to transitory increases in demand should themselves be transitory. We find results that are largely similar to our baseline specifications. This suggests that, for the causal channel to be a concern, it must then be the case that highly transitory increases in purchases during the hoarding period determine differences in prices more than 18 months later. While perhaps remotely plausible given sticky prices, the length of time elapsed begins to strain credulity.

Hoarding as a signal of future demand

An alternative but related possibility is that the degree of transitory hoarding during the crisis was a signal of a permanent increase in demand, and that heterogeneity in price increases reflected these permanent changes. This does not appear to be the case. For example, define store level post-hoarding quantity growth as $100 \times \left(\frac{\bar{q}_i^a - \bar{q}_i^b}{\bar{q}_i^b}\right)$, where q_i^a represents average quantity sold in the post-hoarding period. If we regress this measure on Quantity Growth$_i$ during the crisis, we get a slight negative coefficient (≈ -0.05). If anything, this indicates that stores facing greater hoarding during the crisis actually saw slightly less demand in the future (perhaps because customers had already build up inventories).

5 Quantifying Speculative versus Precautionary Motives

The combination of our time-series and cross-sectional evidence indicates that consumer hoarding was, at least to some extent, motivated by speculation on sticky prices. However, this does not rule out the possibility that precautionary concerns also play a role. In this section we build a forecast test based in a simple theoretical model that allows us to quantify the relative importance of speculative and precautionary motives. We then implement this test in the context of the rice hoarding episode.

5.1 Model

We base our test on a model of optimal household hoarding that combines both speculation (because stickiness makes price changes partially forecastable) and hedging (because households require a baseline level of the staple, they are effectively *short* rice and exposed to price risk). Suppose a household can choose a level of inventory (or hoarding) I by purchasing the good today at known p_0. Tomorrow, the household faces an uncertain price p_1 and must consume $b > 0$ units of rice.[23] We assume that the risk-free rate is zero, the households initial wealth is w_0, and wealth next period is given by:

$$w_1 = (w_0 - Ip_0) + Ip_1 - bp_1.$$

Here, b is the households exposure to price risk.

We assume a simple mean-variance utility specification with a coefficient of risk aversion of γ,

[23] This assumption captures, in a stylized way, the fact that rice is a staple good.

and that households face quadratic storage costs $C(I) \equiv cI^2/2$. The household problem is then

$$\text{Max}_I \left\{ \text{E}[I(p_1 - p_0) - bp_1] - \frac{\gamma}{2}\text{Var}[I(p_1 - p_0) - bp_1)] - \frac{c}{2}I^2 \right\}.$$

The first-order condition with respect to I gives:

$$I^* = \frac{\text{E}[(p_1 - p_0)] + \overbrace{\gamma b \text{Cov}[(p_1 - p_0), p_1]}^{\equiv h_0 \text{ (hedging motive)}}}{c + \gamma \text{Var}[(p_1 - p_0)]} = \frac{\text{E}[(p_1 - p_0)] + \overbrace{\gamma b \sigma_p^2}^{h_0}}{c + \gamma \sigma_p^2}, \tag{2}$$

where I^* is the optimal hoarding level and $\sigma_p^2 = \text{Var}[p_1]$.

From equation (2) it is immediately clear that the optimal level of hoarding is increasing in the expected price change, which will tend to be positive when prices are sticky. This is the speculative motive. In particular:

$$\frac{\partial I^*}{\partial E[(p_1 - p_0)]} = \frac{1}{c + \gamma \sigma_p^2} > 0.$$

If there is no uncertainty, the optimal inventory level is given by

$$I^* = \frac{\text{E}[p_1 - p_0]}{c} = \frac{p_1 - p_0}{c}. \tag{3}$$

So when uncertainty is low, for example, during a regular retail promotion, the sensitivity of inventories to expected price changes is modulated solely by the storage cost c.

Hoarding also arises from precautionary or hedging demand as captured by the h_0 in equation (2). As $b > 0$ grows, households will hoard more because I^* is a hedge against uncertainty in p_1. Furthermore, as long as inventories lie below b, an increase in uncertainty about future prices will lead to more hoarding for risk averse households:

$$\frac{\partial I^*}{\partial \sigma_p^2} = \frac{\gamma}{c + \gamma \sigma_p^2}(b - I^*)$$

which is positive if $b > I^*$. This is the precautionary motive for household hoarding.

5.2 A Two-Step Test

Based on this simple model, we derive an empirical test in two steps to assess the relative importance of speculative versus precautionary motives for hoarding. We introduce a subscript i for cross-household variation that captures households facing different prices p_{1i} and having different exposure to rice price risk b_i. For simplicity, we assume that storage costs and risk aversion are constant across households.

Step 1: Recovering Risk Neutral Forecasts. In the first step, we recover the key parameter governing storage costs (c), which determines consumer sensitivity to future price changes when there is no uncertainty. To do so, we examine situations in which coming price changes are clearly advertised and well understood by households: retail sales during normal, non-crisis, periods. The basic idea is to regress observed stockpiling when a product goes on sale against the size of the sale (measured as the post-sale price increase). Under the assumption that there is no uncertainty about the size of the sale, and that consumers therefore know $p_{1i} - p_{0i}$, this allows us to recover an estimate of the convex storage cost $\hat{c}$.[24]

We can then use our estimate of $\hat{c}$ to generate what we call a *risk neutral forecast* of future prices in any period. If consumers are risk neutral ($\gamma = 0$), and hence not driven to stockpile by any precautionary or hedging motive, an estimate of their expectations can be recovered as

$$\mathrm{E}[\widehat{p_{1i} - p_{0i}}] = \hat{c} I_i^*.$$

The key assumption here is that the storage cost c is stable over time, including across crisis and non-crisis periods.

Step 2: A Forecast Test. Notice from equation (2), which governs optimal hoarding, that

$$\mathrm{E}[p_{1i} - p_{0i}] = -h_{0i} + \left(c + \gamma \sigma_{p_i}^2\right) I_i^*.$$

[24]If $\sigma_{p_i}^2 = 0$, the first order condition gives $I_i* = \frac{1}{c}(p_{1i} - p_{0i})$.

If consumers have rational expectations, this implies that

$$p_{1i} - p_{0i} = -h_{0i} + \left(c + \gamma\sigma_{p_i}^2\right) I_i^* + \epsilon_i,$$

where $\text{Cov}(\epsilon_i, I_i^*) = 0$. Given this, a linear regression of ex-post price changes during the crisis period, $p_{1i} - p_{0i}$ on our risk neutral price forecasts (i.e., $\text{E}[\widehat{p_{1i} - p_{0i}}]$ measured using (i) using non-crisis estimates of storage costs $\hat{c}$ and (ii) observable stockpiling during the crisis I_i^*), provides coefficient:

$$\begin{aligned}\beta &= \frac{\text{Cov}[-h_{0i} + \left(c + \gamma\sigma_{p_i}^2\right) I_i^* + \epsilon_i, cI_i^*]}{\text{Var}(\hat{c}I_i^*)} \\ &= 1 + \sigma_{p_i}^2 \frac{\gamma}{c} - \frac{\text{Cov}[h_{0,i}, cI_i^*]}{\text{Var}(cI_i^*)}\end{aligned}$$

The slope β is our object of interest. It is composed of three terms. If households are risk neutral ($\gamma = 0$) and motivated only by speculation we have only the first term: $\beta = 1$. This is a standard rational forecast result, and provides us a benchmark.

If consumers are risk averse ($\gamma > 0$) but have no precautionary hedging incentive ($b_i = 0$), we have the first two terms: $\beta = (1 + \sigma_{p_i}^2 \frac{\gamma}{c}) > 1$. A risk averse household will speculate less than a risk neutral households, so will increase quantities by by a smaller amount for the same forecasted increase in prices. If $b_i > 0$ the third term enters. Because $\text{Cov}[h_{0,i}, cI_i^*] > 0$, this term pushes β downward. A larger precautionary or hedging demand leads to a flatter slope that is closer to 0. As a consequence, a slope of β between 0 and 1 provides a quantification of the relative importance of precautionary versus speculative motives.

This second step is reminiscent of rational forecast tests following Mincer & Zarnowitz (1969); Nordhaus (1987); Keane & Runkle (1990). The slope of the regression of realized price adjustments on expected price adjustments captures the efficiency of household forecasts. A slope close to zero means household hoarding is driven by precautionary purchases and hence forecasts are not as efficient.

5.3 Implementing Step 1

We now turn to implementing our test using data from the 2008 rice crisis. We conduct our analysis at 4 levels of aggregation: the store-UPC level, the store-brand level, the store level, and the brand level. To build a risk neutral forecast, we require two objects. The first is c, the storage cost for households. We recover this by estimating consumer responsiveness to typical retail sales following (Hendel & Nevo, 2006). We provide details of this estimation and various robustness checks in Appendix B. We generally find an elasticity to the value of a retail sale of around 1.2. If consumers know that the price of a given store-UPC will increase by 10% when a promotion ends, they increase purchases by 12%. This elasticity allows us to back out $\hat{c}$, our estimate of storage costs.

The second is an estimate of the degree of stockpiling that took place in the crisis for each unit i. We measure this as Quantity Growth$_i$, the percentage increase in quantity for unit i relative to a benchmark. In our primary specifications, we estimate Quantity Growth$_i$ by comparing average weekly sales in the hoarding period to a baseline of average weekly sales in the pre-hoarding period.[25] We construct this (and estimate $\hat{c}$) separately at the store-UPC, store-brand, store, and brand levels. We can then construct our risk neutral price forecast as:

$$\text{Forecast}_i = \hat{c} \times \text{Quantity Growth}_i.$$

We next turn to comparing these to price realizations.

5.4 Implementing Step 2

The second step of the test involves regressing realized shelf price changes on our risk neutral forecasts. For our measure of realized price increases, we once again use Post-Hoarding Price Growth$_i$, measured at the relevant level of aggregation (e.g. store-UPC). Our regressions are then:

$$\text{Post-Hoarding Price Growth}_i = \alpha + \beta \text{Forecast}_i + u_i. \tag{4}$$

[25] This follows the definition in footnote 22. An alternative would be to use the post-hoarding period, which might better capture typical static weekly purchases at the post-hoarding price. However, we chose to use pre-hoarding period sales due to concerns that purchases in the immediate post-hoarding period would be low due to dynamic reallocations. Ultimately, the results are qualitatively similar whether we use the pre-hoarding period, the post-hoarding period, or the full sample excluding the hoarding period.

Because an estimate of $\hat{c}$ is necessary to compute Forecast$_i$, we bootstrap over both steps to construct standard errors, with clusters drawn at the unit level over 1000 repetitions.

5.5 Test Results

We present results of our efficiency tests in Table 5. Our key object of interest is the coefficient β. Across store-UPCs, store-brands, and stores, we find consistent estimates of β that are between 0.1 and 0.15 and statistically significant. Our brand level coefficient is also positive, although it is smaller and not statistically significant (which may be unsurprising, given the very small number of brands and substantially higher standard errors).

The fact that β is positive provides evidence that the speculative motive is present, directly mirroring the cross-sectional regressions shown in Section 4. Excess purchases, and hence our risk-neutral forecasts, contain information about the cross section of coming price increases. The fact that β is below 1, and in fact is closer to 0.1, indicates that the precautionary motive is also at work, and perhaps dominant. These results are consistent at the product, brand, and store level.

Appendix Figure A.VII presents binned scatter plots of the same relationships. These plots (and our tests) provide a scaled version of the cross-sectional relationships shown in Figure 6, with the slope β scaled so as to be interpreted in accordance with our forecast test. The imposed black line 45-degree line provides an efficient risk-neutral benchmark, with a slope of 1. Again, we see positive correlations across the board, albeit with slopes that are meaningfully below 1. All in, the results of our test suggest that the speculative motive is responsible for a non-trivial portion of observed hoarding during the crisis, but that the precautionary motive is a more dominant driver in this episode.

6 Rice Hoarding in the COVID Crisis

Our analysis suggests that the speculative motive generated by sticky prices was a significant, if perhaps not dominant, driver of household hoarding in the 2008 rice crisis. A natural question is whether the same motive is relevant in other episodes. As a final step, we consider the role of sticky prices in perhaps the most salient hoarding experience in recent memory: the run on food, staples, and other goods at the beginning of the COVID-19 pandemic.

Unfortunately, the pandemic does not provide quite as clean a laboratory as the 2008 rice crisis for studying the role of sticky prices. While COVID-19 and the associated regulatory restrictions certainly entailed a supply shock (which ultimately impacted prices) they also surely had non-trivial impacts on household demand. Consumers faced a need for new, pandemic specific goods such as hand sanitizer and masks, changes in income due to job loss, furlough, or government support measures, and different consumption needs to suit daily life while locked-down or working from home. Consequently, abnormal consumer purchase patterns during the crisis likely involve a more complex interaction between supply and demand side shocks.

Despite these caveats, this section presents evidence that an implicit promotion motive was a driver of rice hoarding in the early stage of the COVID-19 pandemic. We do so using a more limited regional scanner database that covers approximately 1000 stores belonging to three major supermarket chains in the south.[26] Our data is derived from a frequent-shopper card database that records daily transactions for all customers visiting the store (transactions include non-card users and employee cards). We observe detailed information on quantities, prices, and promotions at the UPC level. In addition, we observe detailed product information (package size, brand name) as well as exact locations of stores. For our analysis, we aggregate the individual card transactions to the UPC-Store-Week level and use data from Jan 1, 2020 to July 1, 2020. We focus only on UPCs listed in the rice category.

We begin by presenting simple, time-series evidence that mirrors the analysis conducted in Section 3. Specifically, Panel A of Figure 7 plots prices and quantities for rice in the first 6 months of 2020. The red line shows total weekly quantity sold in our sample (normalized by the mean over the sample period. The blue line shows the sales weighted unit-price of rice, again normalized by its mean over the period. There are two key takeaways. First, there was substantial hoarding of rice at the outset of the COVID-19 crisis. There is an extreme spike in quantities sold the week of March 11th-17th, 2020, which coincides with the proclamation of a national state of emergency on March 13th, 2020. Second, there is a modest, but noticeable increase in the price of rice in the weeks following this aggressive hoarding, but no meaningful increase the week of the 11th-17th itself. Indeed, during the 11th-17th, the (sales weighted) price was marginally below the average

[26] Appendix Figure A.VI shows the geographical locations of the stores in the data along with demographic information for the Zip codes (from 2019 Census ACS-5yr files).

for the first 10 weeks. The following week (the 18th-24th) prices rose by roughly 15 percent and stayed consistently high through the remainder of our sample period.

In other words, extensive household hoarding appears to have preempted the shelf-price increase generated by the crisis, just as in the 2008 episode. This is consistent with sticky prices driving consumers to stockpile due to an implicit promotion motive. To confirm this intuition, we next turn to the cross-section of products and stores. Following Section 4, we once again consider the relationship between excess purchases at the UPC$\times$Store level during the hoarding period (which we define as the week of March 11th-17th), and subsequent price increases *after* the hoarding period. If the intensity of hoarding predicts later price hikes in the cross-section, we take this as evidence that consumers are driven by the implicit promotion created by sticky prices.

As the upward slope in Panel B of Figure 7 shows, we find that consumers hoarded more aggressively in the products that later experienced larger price increases, just as in the 2008 rice crisis. Products with greater excess purchases the week of the 11th-17th (in percentage terms), saw larger price hikes in the subsequent weeks. This pattern suggests that consumers had some information about coming price changes at the product level and that this information drove their hoarding patterns, at least to some extent, during the early portion of the COVID crisis.

Taken together, this evidence suggests that sticky prices are a relevant driver of hoarding across different episodes, and played a role in the stockpiling that occurred in March of 2020. However, a few points of qualification are valuable to note. First, the magnitude of excess purchases during this episode (seen in Panel A of Figure 7) exceeds that in the 2008 crisis, particularly relative to the scale of the later price hike. This suggest that there may have been substantial additional drivers of hoarding during this period, in addition to the speculative motive. This is unsurprising given the demand shifts discussed earlier and the deep uncertainty that existed at the time. Second, efficiency tests, in the vein of those conducted in Section 5 further confirm this notion. The β coefficient on these tests is an order of magnitude below the coefficient estimated for the rice crisis (approximately 1%, versus over 10%, using the same elasticity). Again, this suggests that other factors: precaution, panic and demand, were more dominant drivers of observed hoarding during the COVID episode.

7 Conclusion

Little is known about the determinants of hoarding despite its importance in disrupting the well-functioning of markets. To make progress on this issue, we analyze major recent hoarding episodes using an optimal inventory model in which risk-averse agents hoard both to hedge against price uncertainty and to speculate when prices are predictable. Using supermarket scanner data from two recent hoarding episodes, the 2008 rice crisis and COVID-19, we provide reduced-form evidence of the importance of the speculative motive due to sticky retail prices. Our quantification suggests that speculation accounts for a meaningful fraction of overall hoarding, although smaller than precaution in our episodes. There are a number of paths for future research, including working out the implications of these two types of motives for anti-price gouging regulation, firm price setting, and supply chain risk management.

References

ACHARYA, VIRAL V, LOCHSTOER, LARS A, & RAMADORAI, TARUN. 2013. Limits to arbitrage and hedging: Evidence from commodity markets. *Journal of Financial Economics*, **109**(2), 441–465.

BARRO, ROBERT J. 1972. A theory of monopolistic price adjustment. *The Review of Economic Studies*, **39**(1), 17–26.

BENABOU, ROLAND. 1989. Optimal price dynamics and speculation with a storable good. *Econometrica: Journal of the Econometric Society*, 41–80.

BERGTOLD, JASON, AKOBUNDU, EBERE, & PETERSON, EVERETT B. 2004. The FAST method: Estimating unconditional demand elasticities for processed foods in the presence of fixed effects. *Journal of Agricultural and Resource Economics*, 276–295.

BESSEMBINDER, HENDRIK. 1992. Systematic risk, hedging pressure, and risk premiums in futures markets. *The Review of Financial Studies*, **5**(4), 637–667.

CAVALLO, ALBERTO, CAVALLO, EDUARDO, & RIGOBON, ROBERTO. 2014. Prices and supply disruptions during natural disasters. *Review of Income and Wealth*, **60**, S449–S471.

CHILDS, NATHAN W, KIAWU, JAMES, *et al.* 2009. *Factors behind the rise in global rice prices in 2008*. US Department of Agriculture, Economic Research Service.

DAVIS, CALE WREN, *et al.* 2008. *An analysis of the enactment of anti-price gouging laws*. Ph.D. thesis, Montana State University-Bozeman, College of Agriculture.

DAWE, DAVID. 2012. *The rice crisis: Markets, policies and food security*. Routledge.

DAWE, DAVID, & SLAYTON, TOM. 2010. The world rice market crisis of 2007-2008. *Pages 15–29 of:* DAWE, DAVID (ed), *The rice crisis: Markets, policies and food security.* London, Earthsan and FAO.

DEATON, ANGUS, & LAROQUE, GUY. 1992. On the behaviour of commodity prices. *The review of economic studies*, **59**(1), 1–23.

DELLAVIGNA, STEFANO, & GENTZKOW, MATTHEW. 2017. *Uniform pricing in us retail chains*. Tech. rept. National Bureau of Economic Research.

FAMA, EUGENE F, & FRENCH, KENNETH R. 1987. Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *Journal of Business*, **60**(1), 55–73.

GAGNON, ETIENNE, & LOPEZ-SALIDO, DAVID. 2015. Small price responses to large demand shocks. *Journal of the European Economic Association*.

GIBERSON, MICHAEL. 2011. The problem with price gouging laws. *Regulation*, **34**, 48.

GORODNICHENKO, YURIY, & WEBER, MICHAEL. 2016. Are sticky prices costly? Evidence from the stock market. *American Economic Review*, **106**(1), 165–99.

GORTON, GARY B, HAYASHI, FUMIO, & ROUWENHORST, K GEERT. 2013. The fundamentals of commodity futures returns. *Review of Finance*, **17**, 35–105.

HAMILTON, JAMES D. 2009. *Causes and consequences of the oil shock of 2007-08*. Tech. rept. National Bureau of Economic Research.

HAUSMAN, JERRY A. 1996. Valuation of new goods under perfect and imperfect competition. *Pages 207–248 of: The economics of new goods*. University of Chicago Press.

HENDEL, IGAL, & NEVO, AVIV. 2006. Sales and consumer inventory. *The RAND Journal of Economics*, **37**(3), 543–561.

KANG, WENJIN, ROUWENHORST, K GEERT, & TANG, KE. 2020. A tale of two premiums: The role of hedgers and speculators in commodity futures markets. *The Journal of Finance*, **75**(1), 377–417.

KEANE, MICHAEL, & NEAL, TIMOTHY. 2021. Consumer panic in the COVID-19 pandemic. *Journal of econometrics*, **220**(1), 86–105.

KEANE, MICHAEL P, & RUNKLE, DAVID E. 1990. Testing the rationality of price forecasts: New evidence from panel data. *The American Economic Review*, 714–735.

KEHOE, PATRICK, & MIDRIGAN, VIRGILIU. 2015. Prices are sticky after all. *Journal of Monetary Economics*, **75**, 35–53.

KILIAN, LUTZ, & MURPHY, DANIEL P. 2014. The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, **29**(3), 454–478.

LEIBTAG, EPHRAIM. 2007. *Cost pass-through in the US coffee industry*. USDA Economic Research Serv.

MCSHANE, BLAKELEY B, CHEN, CHAOQUN, ANDERSON, ERIC T, & SIMESTER, DUNCAN I. 2016. Decision stages and asymmetries in regular retail price pass-through. *Marketing Science*, **35**(4), 619–639.

MINCER, JACOB A, & ZARNOWITZ, VICTOR. 1969. The evaluation of economic forecasts. *Pages 3–46 of: Economic forecasts and expectations: Analysis of forecasting behavior and performance*. NBER.

NAKAMURA, EMI, & STEINSSON, JÓN. 2008. Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics*, **123**(4), 1415–1464.

NORDHAUS, WILLIAM D. 1987. Forecasting efficiency: concepts and applications. *The Review of Economics and Statistics*, 667–674.

NOSS, A. 2010. Household income for states: 2008 and 2009–American Community Survey briefs. *US Census Bureau*.

PRIEST, TYLER. 2012. The Dilemmas of Oil Empire. *Journal of American History*, **99**(1), 236–251.

SCHEINKMAN, JOSE A, & SCHECHTMAN, JACK. 1983. A simple competitive model with production and storage. *The Review of Economic Studies*, **50**(3), 427–441.

SEN, AMARTYA. 1983. *Poverty and famines: An essay on entitlement and deprivation*. Oxford: Oxford University Press.

SHESHINSKI, EYTAN, & WEISS, YORAM. 1977. Inflation and costs of price adjustment. *The Review of Economic Studies*, **44**(2), 287–303.

SHIN, HYUN SONG. 2021. Bottlenecks, labour markets and inflation in the wake of the pandemic. *Bank of International Settlements Working Paper*.

SINGLETON, KENNETH J. 2013. Investor flows and the 2008 boom/bust in oil prices. *Management Science*, **60**(2), 300–318.

SLAYTON, TOM. 2009. *Rice crisis forensics: How Asian governments carelessly set the world rice market on fire*. Tech. rept.

TANG, KE, & XIONG, WEI. 2010. *Index investment and financialization of commodities*. Tech. rept. National Bureau of Economic Research.

ZWOLINSKI, MATT. 2008. The ethics of price gouging. *Business Ethics Quarterly*, **18**(3), 347–378.

FIGURE 1: GLOBAL RICE COMMODITY PRICES RISE FOLLOWING INDIA EXPORT BAN

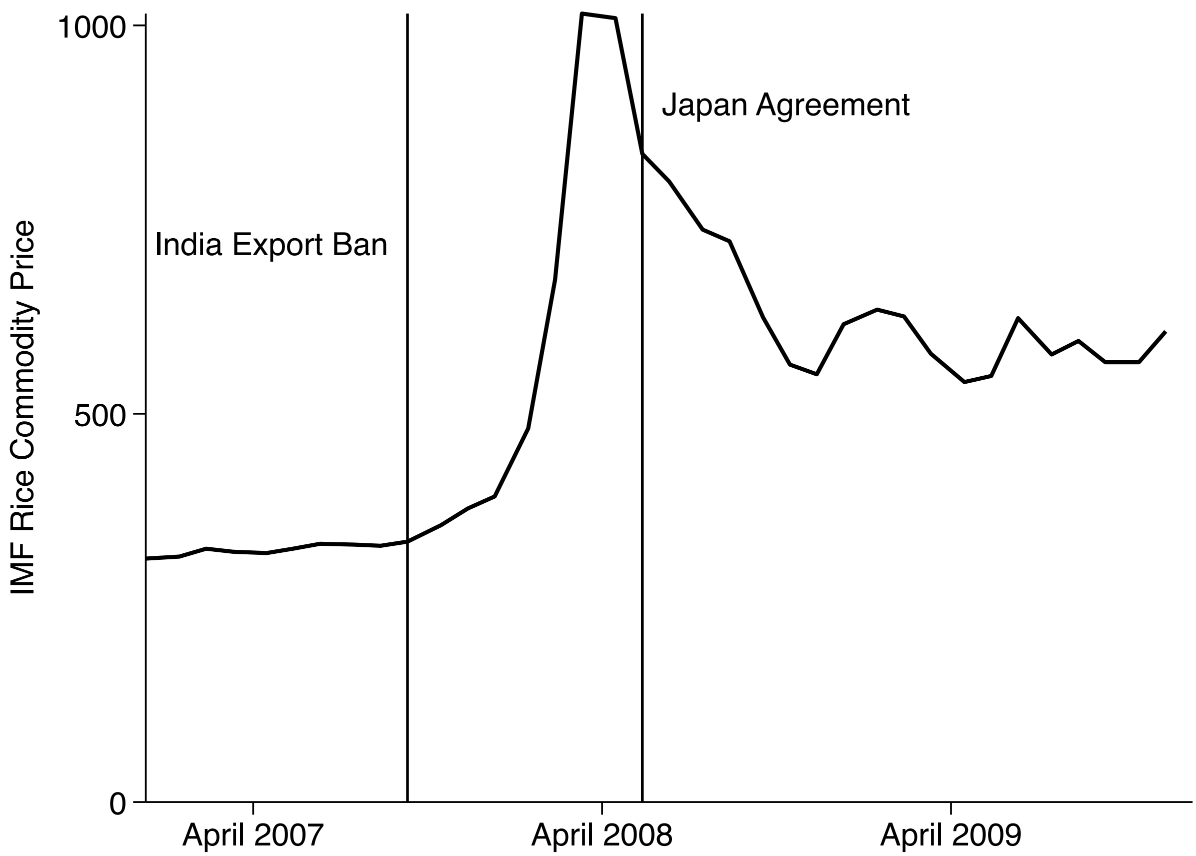

Notes: The black line displays the rice series from the IMFs Primary Commodity Price System. The series measures the Thailand nominal price quote for 5 percent broken milled white rice in USD per metric ton. The first vertical line denotes India's ban on rice exports in October 2007, while the second vertical line denotes June 2008, to approximate the timing of Japan's agreement to release rice reserves.

FIGURE 2: DELAYED PRICE UPDATING BY US RETAIL STORES

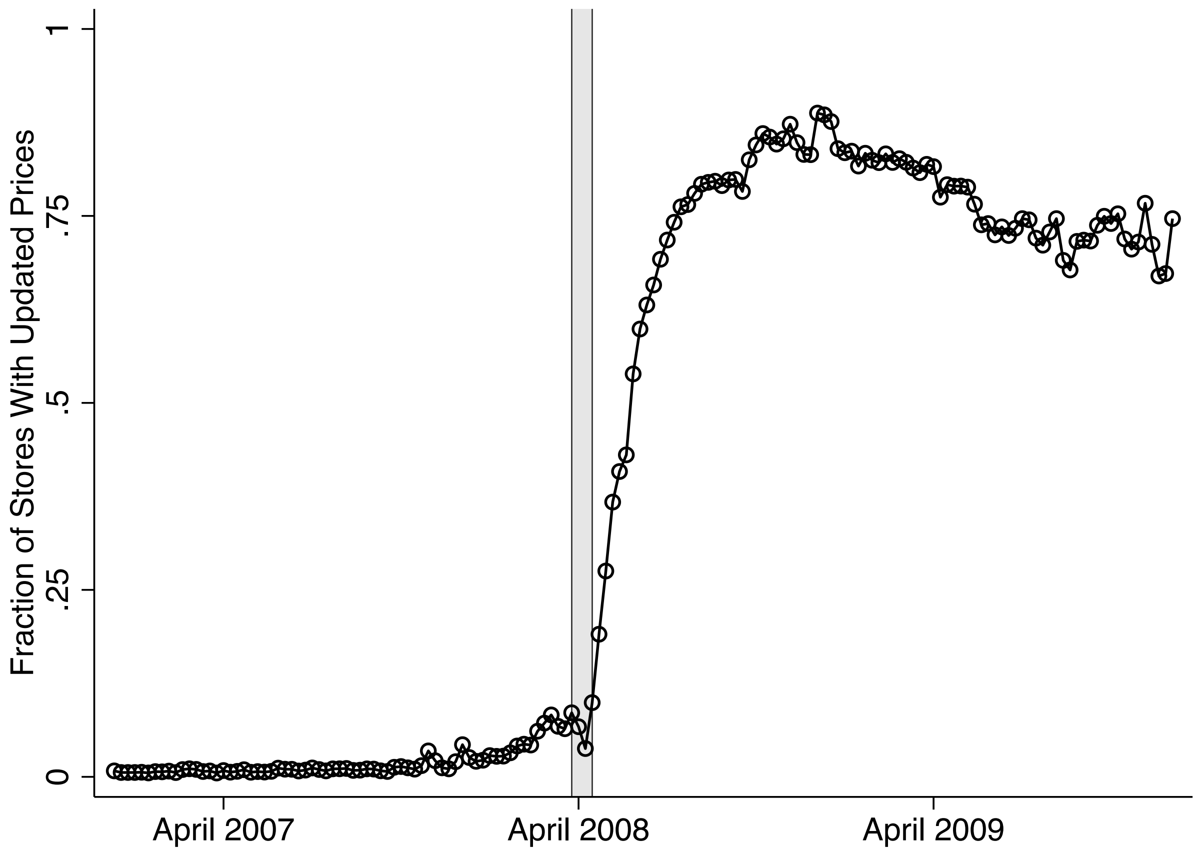

Notes: Plot displays the fraction of stores that have *updated prices* in the wake of the shock to international prices. A store is determined to have updated its price if the price is greater than 125 percent of the 2007 average. Grey region denotes our designated hoarding period, the weeks of April 19th-May 10th.

FIGURE 3: HOARDING ANTICIPATES CHANGE IN SHELF PRICES

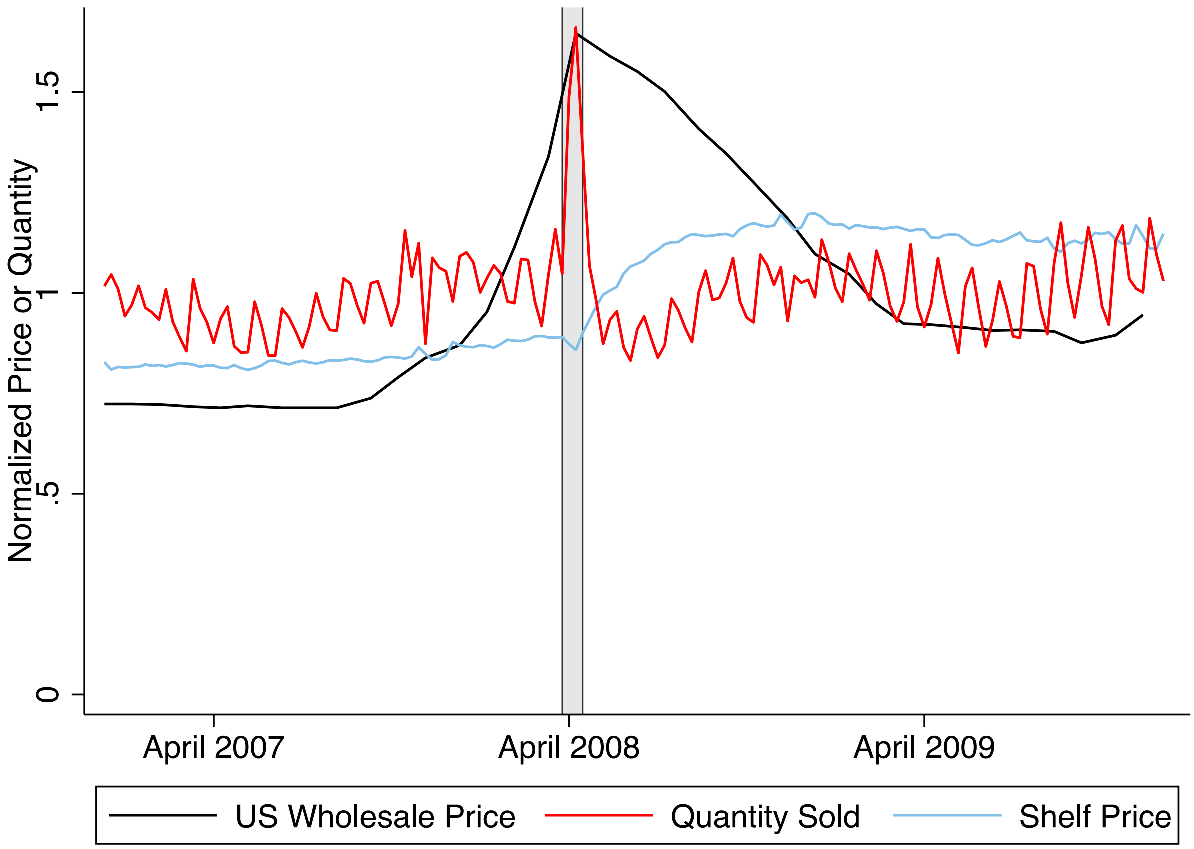

Notes: The black line displays a proxy for the US wholesale rice price at the monthly level. The proxy is based on average f.o.b. price for long grain rice at selected milling centers in Southwest Louisiana. Data provided by USDA, based on data from Agricultural Marketing Service, *National Weekly Rice Summary*. The red line displays average weekly sales at the store level, based on scanner data. The blue line displays the weekly average shelf price based on our store level rice prices. All variables are normalized by the average over the period shown: 2007-2009. Grey region denotes our designated hoarding period, the weeks of April 19th-May 10th.

FIGURE 4: SUBSTANTIAL DISPERSION IN POST-HOARDING PRICE CHANGES

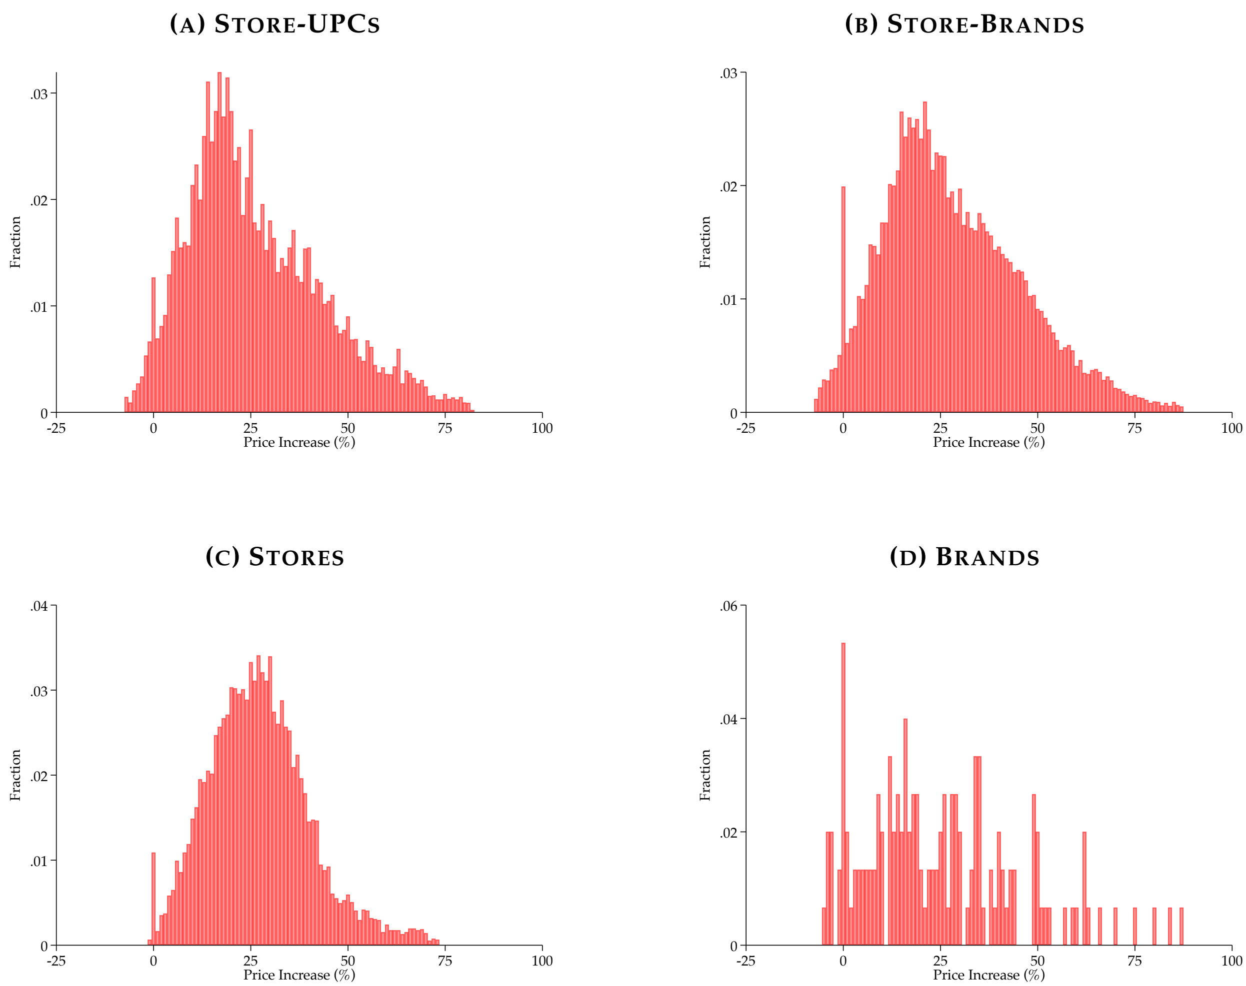

Notes: Histogram of realized post-hoarding price changes, trimmed at the 1 percent level. Letting $\bar{p}^h$ denote the average unit price in the hoarding period and $\bar{p}^a$ denote the average unit price after the hoarding period, it is defined as $100\times\left(\frac{\bar{p}^a-\bar{p}^h}{\bar{p}^h}\right)$. Store-UPC includes all store-UPC pairs with at least 5 units sold on average per week in 2007. Store-Brand includes all store-brand pairs with at least 5 units sold on average per week in 2007. Store refers to all stores in the sample while Brand refers to all rice brands with at least 5 units sold on average per week in 2007. The sample period covers 2007-2009, and the hoarding period is defined as the weeks of April 19th-May 10th, 2008.

Figure 5: Heterogeneity in UPC-Level Price Changes Correlated With Hoarding

High Post-Hoarding Price Increase

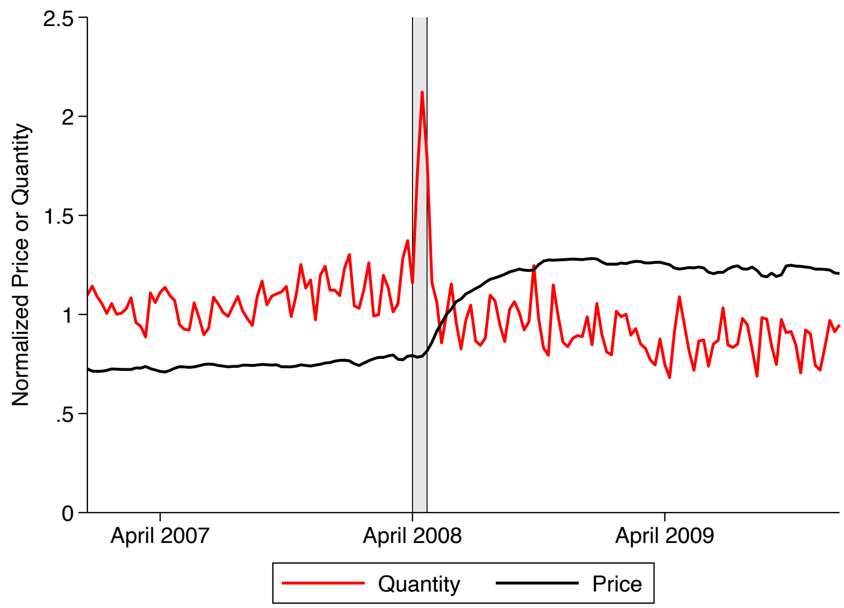

Low Post-Hoarding Price Increase

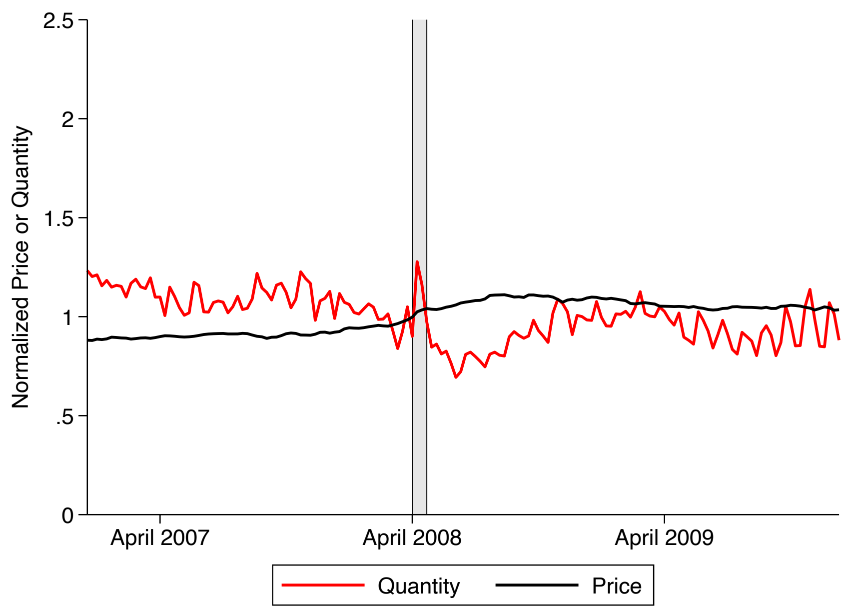

Notes: Average shelf price-per-ounce and ounces sold for UPC×store pairs with the highest and lowest post- hoarding period price changes. Across products, prices are measured as per ounce. For a given UPC×store pair, the post-hoarding price increase is defined as the ratio of the average price in the post-hoarding portion of our sample to the average price in the hoarding period itself (where the period is defined as the weeks of April 19th-May 10th). High and low post-hoarding price increases are the UPC-store pairs in the top or bottom quartile, respectively. Prices and quantities are normalized by the average over the sample period for the set of UPC×stores used in each plot. Grey region denotes our designated hoarding period.

FIGURE 6: CROSS-SECTIONAL DIFFERENCES IN HOARDING PREDICT POST-HOARDING PRICE INCREASES

(A) STORE-UPCS

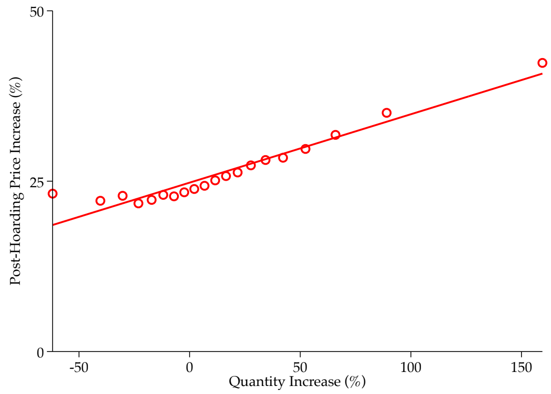

(B) STORE-BRANDS

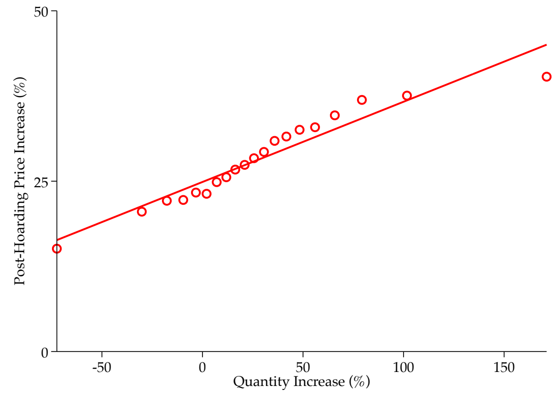

(C) STORES

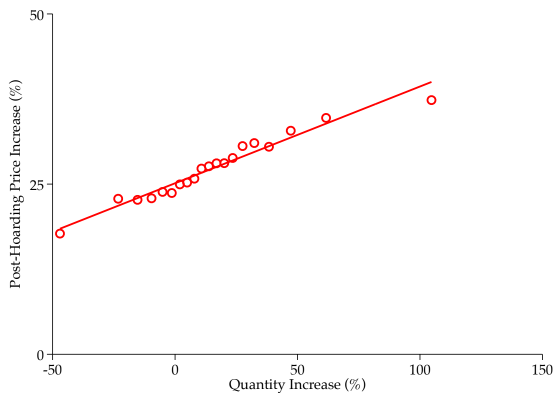

(D) BRANDS

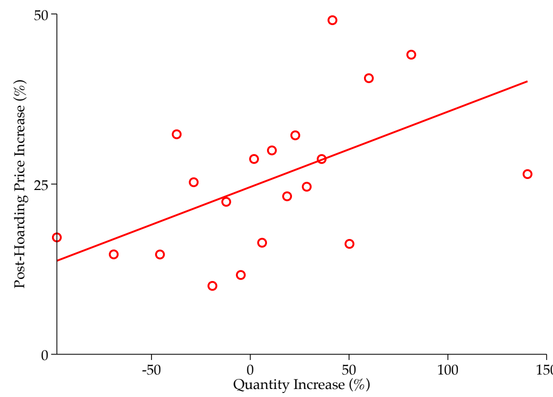

Notes: Binned scatterplots of post-hoarding prices increases on quantity increases during hoarding period. Letting $\bar{p}^h$ denote the average unit price in the hoarding period and $\bar{p}^a$ denote the average unit price after the hoarding period, the price increase is defined as $100 \times \left(\frac{\bar{p}^a - \bar{p}^h}{\bar{p}^h}\right)$. If $\bar{q}^h$ is the average quantity sold during the hoarding period, $\bar{q}^b$ is average quantity sold in the pre-period the quantity increase is $100 \times \left(\frac{\bar{q}^h - \bar{q}^b}{\bar{q}^b}\right)$. The sample period covers 2007-2009, and the hoarding period is defined as the weeks of April 19th-May 10th, 2008. Store-UPC includes all store-UPC pairs with at least 5 units sold on average per week in 2007. Store-Brand includes all store-brand pairs with at least 5 units sold on average per week in 2007. Store refers to all stores in the sample while Brand refers to all rice brands with at least 5 units sold on average per week in 2007. Points represent means within ventiles of predicted price discounts. Lines represent a linear fit through the underlying data. All variables are winsorized at the 1 percent level.

Figure 7: Sticky Prices and Hoarding During the Covid-19 Pandemic

Panel A: Prices Increased Following Covid-Driven Rice Hoarding

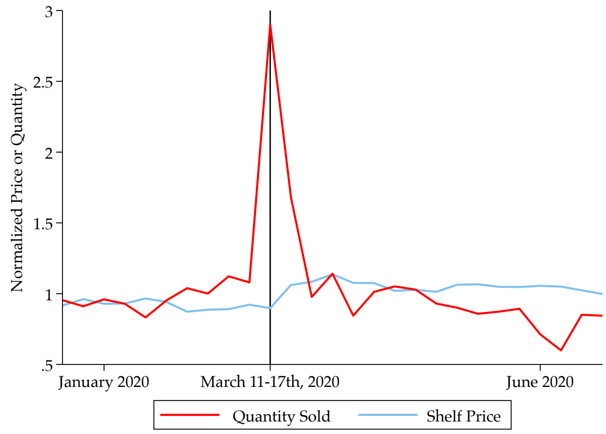

Panel B: Cross-sectional Differences in Hoarding Predict Price Increases

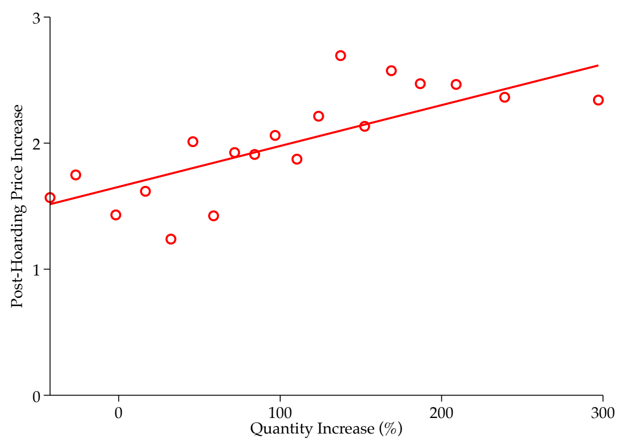

Notes: Panel A shows the dynamics of prices and quantities for rice in our sample of southern retailers in the first half of 2020. The blue line shows the sales weighted unit price and the red line shows total quantity sold. Both are normalized by their mean over the sample period. Panel B shows binned scatterplots of post-hoarding price growth against quantity growth at the store-upc level. Here, we let $\bar{p}^h$ denote the unit price in the week of March 11th-17th and $\bar{p}^a$ denote the average unit price across weeks after March 17th. Price growth is then defined as $100\times\left(\frac{\bar{p}^a-\bar{p}^h}{\bar{q}^h}\right)$. We let q_h be the total quantity sold in the week of March 11th-17th and q_b be the average weekly quantity sold in the weeks before March 11th. We define quantity growth as $\left(\frac{\bar{q}^b-\bar{q}^h}{\bar{q}^b}\right)$. Points represent means within ventiles of quantity growth.

TABLE 1: SUMMARY STATISTICS

	Mean	S.D.	1st Percentile	99th Percentile
	Panel A: Stores			
Quantity Sold (oz)	8068.5	18916.3	0	76776
Price (80oz)	5.39	1.61	2.34	9.67
Total Stores	10561			
	Panel B: UPC-Stores			
Ounces per Unit	52.6	64.3	14	320
Units Sold	11.5	18.9	0	69
Quantity Sold (oz)	695.9	3537.7	0	7920
Price (80oz)	5.58	3.42	1.75	13.9
Total UPCs	547			
UPC $\times$ Stores	71952			
	Panel C: Brand-Stores			
Units Sold	27.0	40.2	0	180
Quantity Sold (oz)	1688.8	5323.4	0	16512
Price (80oz)	6.22	4.17	1.82	19.1
Total Brands	154			
Brand $\times$ Stores	43953			
	Panel D: Households			
Volume (oz)	1.83	23.7	0	48
Any Purchase	0.025	0.16	0	1
Volume $\mid$ Purchase	71.9	130.5	12	480
Expenditure $\mid$ Purchase	3.55	4.61	0.50	19.0
Total Households	42441			

Summary statistics for weekly store and household data. Weekly prices are sales weighted averages within a store, store-upc or store-brand, normalized to 80 ounces.

TABLE 2: EVIDENCE OF HOUSEHOLD HOARDING AND STICKY PRICES IN APRIL-MAY OF 2008

	Commodity Prices and Sales Peak During Hoarding Period – Shelf Prices Do Not					
	IMF Commodities Price	US Wholesale Price	Quantity (HH)	Shelf Price (HH)	Quantity (Store)	Shelf Price (Store)
Hoarding Period	499.4^{***}	13.78^{***}	8.443^{***}	-0.738^{**}	3397.5^{***}	-0.662
	(112.4)	(4.960)	(0.689)	(0.366)	(381.0)	(0.409)
2007 Mean	332.4	19.3	7.57	4.73	8190.0	4.46
Observations	36	36	156	156	156	156

The first two columns show regressions of monthly time series from 2007-2009 on a dummy equal to one in April and May of 2008. IMF commodities price refers to the rice series from the IMFs Primary Commodity Price System. The series measures the Thailand nominal price quote for 5 percent broken milled white rice in USD per metric ton. US wholesale price refers to the average f.o.b. price for long grain rice at selected milling centers in Southwest Louisiana in USD per cwt (hundredweight). Provided by the USDA based on data from Agricultural Marketing Service, National Weekly Rice Summary. Columns 3-6 show regressions of weekly time series from 2007-2009 on a dummy equal to one during the hoarding period (the weeks of April 19th-May 10th). Quantity refers to the average quantity for households or stores in our data in ounces. Shelf price (store) refers to average shelf price across stores or households normalized to 80 ounces. Price variables are sales weighted within households or stores and equal weighted across stores or households. 2007 mean refers to the mean of the dependent variable in 2007. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 3: EXCESS HOARDING PURCHASES PREDICT POST-HOARDING SHELF PRICE GROWTH

	Dependent Variable: Post-Hoarding Price Growth							
	Store-UPCs			Store-Brands			Stores	Brands
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Quantity Growth (%)	0.100^{***}	0.113^{***}	0.085^{***}	0.112^{***}	0.137^{***}	0.090^{***}	0.142^{***}	0.111^{***}
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.004)	(0.034)
Mean of Dep. Var.	26.5	26.5	26.5	28.9	28.9	28.9	27.4	26.3
Observations	69439	68662	68597	43065	42462	42447	9252	154
Store FE	No	Yes	Yes	No	Yes	Yes	No	No
UPC FE	No	No	Yes	No	No	No	No	No
Brand FE	No	No	No	No	No	Yes	No	No

Cross-sectional regressions of post-hoarding shelf price growth on quantity growth (excess purchases) at the store-UPC, store-brand, store and brand levels. Post-hoarding shelf price growth is defined based on a comparison of average price in the post-hoarding period (from May 10th to the end of 2009) to the average price during the hoarding period (the weeks of April 17th-May 10th). Quantity growth is defined based on a comparison of average purchases in the hoarding period to average purchases in the pre-hoarding period (from the beginning of 2007 to April 17th). Store-UPC includes all store-UPC pairs with at least 5 units sold on average per week in 2007. Store-Brand includes store-brand pairs with at least 5 units sold on average per week in 2007. Store refers to all stores in the sample and brand refers to all rice brands with at least 5 units sold on average per week in 2007. Both variables are winsorized at the 1 percent level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 4: EXCESS HOARDING PURCHASES PREDICT SHELF PRICE GROWTH: ROBUSTNESS

Panel A: Excess Hoarding Purchases Predict Shelf Price Growth Controlling for Price

	Store-UPCs			Store-Brands			Stores	Brands
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Quantity Growth (%)	0.103^{***}	0.113^{***}	0.059^{***}	0.073^{***}	0.096^{***}	0.068^{***}	0.090^{***}	0.075^{**}
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.004)	(0.033)
Price	-0.957^{***}	-0.637^{***}	-13.610^{***}	-1.737^{***}	-1.704^{***}	-3.773^{***}	-4.983^{***}	-0.724^{***}
	(0.030)	(0.036)	(0.119)	(0.021)	(0.024)	(0.057)	(0.113)	(0.159)
Mean of Dep. Var.	26.5	26.5	26.5	28.9	28.9	28.9	27.4	26.3
Observations	69439	68662	68597	43065	42462	42447	9252	154
Store FE	No	Yes	Yes	No	Yes	Yes	No	No
UPC FE	No	No	Yes	No	No	No	No	No
Brand FE	No	No	No	No	No	Yes	No	No

Panel B: Excess Hoarding Purchases Predict Long Run Shelf Price Growth

	Store-UPCs			Store-Brands			Stores	Brands
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Quantity Growth (%)	0.113^{***}	0.142^{***}	0.100^{***}	0.105^{***}	0.140^{***}	0.098^{***}	0.108^{***}	0.008
	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.008)	(0.051)
Mean of Dep. Var.	26.9	26.9	26.9	28.0	29.0	28.0	27.1	31.0
Observations	60453	59542	59480	41033	40396	40380	9252	154
Store FE	No	Yes	Yes	No	Yes	Yes	No	No
UPC FE	No	No	Yes	No	No	No	No	No
Brand FE	No	No	No	No	No	Yes	No	No

Panel A presents cross-sectional regressions of post-hoarding shelf price growth on quantity growth at the store-UPC, store-brand, store and brand levels, controlling for unit price-per-80 ounces. Panel B presents cross-sectional regressions of long run post-hoarding shelf price growth on quantity growth at the store-UPC, store-brand, store and brand levels. Long run shelf price growth is defined based on a comparison of prices in the last week of our sample (the last week of 2009) to average prices during the hoarding period (the weeks of April 17th-May 10th). Quantity growth is defined based on a comparison of average purchases in the hoarding period to average purchases in the pre-hoarding period (from the beginning of 2007 to April 17th). Store-UPC includes all store-UPC pairs with at least 5 units sold on average per week in 2007. Store-brand includes all stores with at least 5 units sold on average per week in 2007. Store refers to all stores in the sample and brand refers to all rice brands with at least 5 units sold on average per week in 2007. Both variables are winsorized at the 1 percent level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 5: CROSS-SECTIONAL EFFICIENCY TEST

	Store-UPC	Store-Brand	Store	Brand
	(1)	(2)	(3)	(4)
Predicted Price Increase	0.116***	0.126***	0.147***	0.040
	(0.001)	(0.002)	(0.006)	(0.067)
Observations	69439	43953	9252	154

This table presents forecast tests, slopes from regressions of realized post-hoarding shelf price growth on predicted price growth inferred from excess purchases. Post-hoarding shelf price growth is defined based on a comparison of average price in the post-hoarding period (from May 10th to the end of 2009) to the average price during the hoarding period (the weeks of April 17th-May 10th). Forecasted price growth is derived from observed quantity growth during the hoarding period. Specifically: if $\bar{q}^h$ is the average quantity sold during the hoarding period, $\bar{q}^b$ is average quantity sold in the pre-period, and $\frac{1}{\hat{c}}$ is an estimated elasticity to temporary sales at the corresponding level of observation from Appendix Table B.I, we define the forecast to be $\left(\frac{\bar{q}^b-\bar{q}^h}{\bar{q}^b}\right) \times \hat{c}$. The sample period covers 2007-2009, and the hoarding period is defined as the weeks of April 19th-May 10th, 2008. Store-UPC includes all store-UPC pairs with at least 5 units sold on average per week in 2007. Store-Brand includes all store-brand pairs with at least 5 units sold on average per week in 2007. Store refers to all stores in the sample and Brand refers to all rice brands with at least 5 units sold on average per week in 2007. Realized and predicted price discounts are winsorized at the 1 percent level. Standard errors are based on a clustered bootstrap at the cross-sectional unit level (e.g. Store-UPC) with 1,000 iterations. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Internet Appendix: For Online Publication

A Supplementary Tables and Figures

FIGURE A.I: NIELSEN PANEL DEMOGRAPHICS

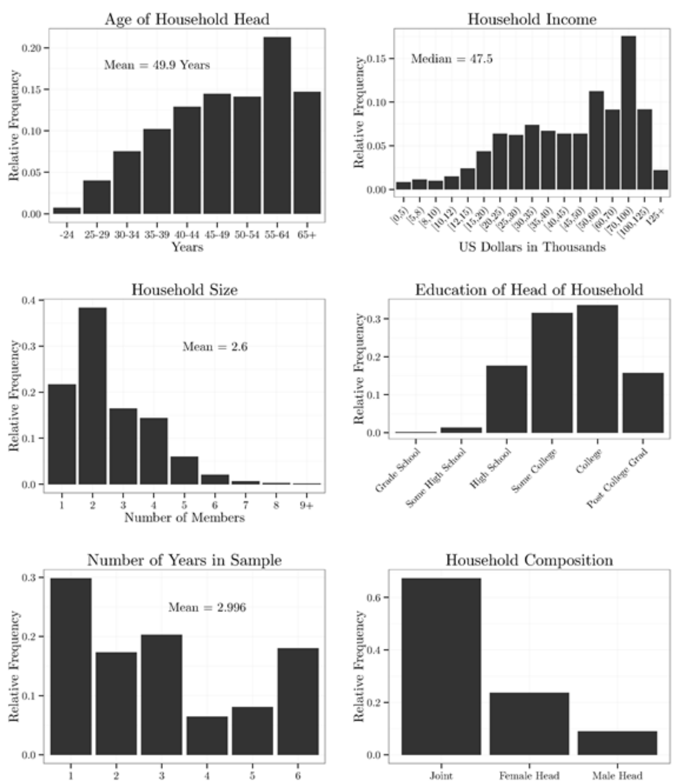

Notes: This figure plots the distribution of demographics of the overall Nielsen Panel.

FIGURE A.II: ALTERNATIVE PRICE MEASURES: FIXING PRODUCT CHARACTERISTICS WITHIN STORES (UPC)

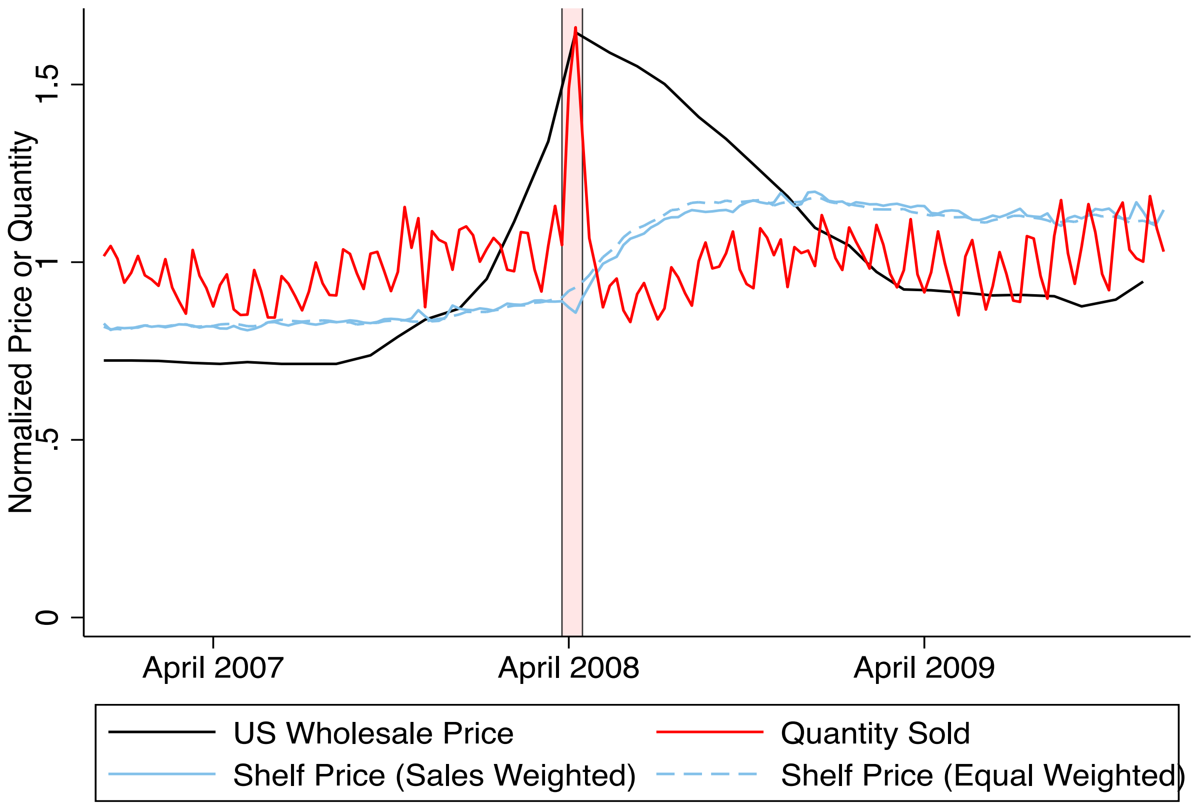

Notes: The black line displays a proxy for the US wholesale rice price at the monthly level. The proxy is based on average f.o.b. price for long grain rice at selected milling centers in Southwest Louisiana. Data provided by USDA, based on data from Agricultural Marketing Service, *National Weekly Rice Summary*. The red line displays average weekly sales at the store level, based on scanner data. The solid blue line displays the sales weighted weekly average shelf price. The dotted blue line shows the equal weighted unit price across all UPCs that appear consistently across all weeks in our sample period. All variables are normalized by the average over the period shown: 2007-2009. Shaded region denotes our designated hoarding period, the weeks of April 19th-May 10th.

FIGURE A.III: HOUSEHOLD INVENTORIES LEAD RETAIL PRICES

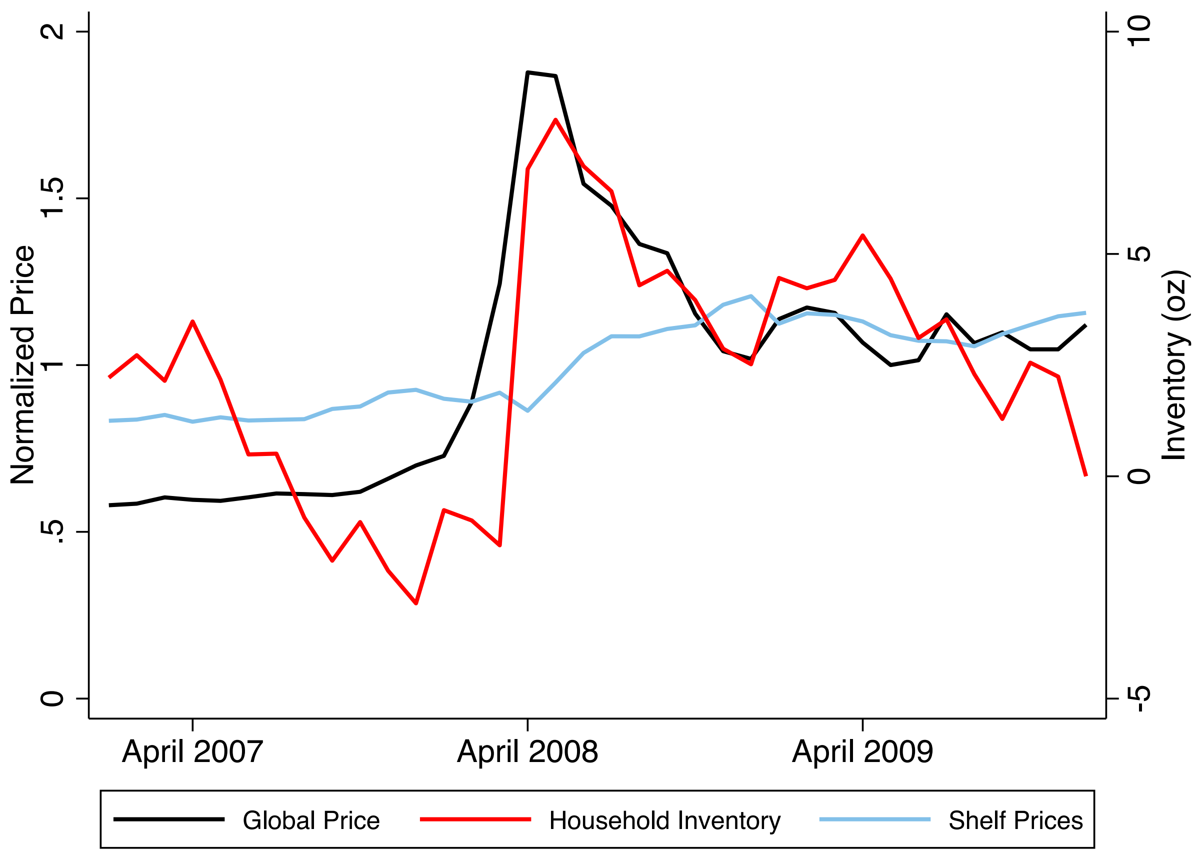

Notes: The black line displays the rice series from the IMFs Primary Commodity Price System. The series measures the Thailand nominal price quote for 5 percent broken milled white rice in USD per metric ton. The blue line shows shelf prices, calculated as the average unit price paid by households in our panel. Both price series are normalized to the mean over the sample period. The red line shows household inventories. Inventories are calculated following the procedure in Hendel & Nevo (2006). For each household, we estimate monthly consumption based on average purchases throughout our sample period. We then construct inventories in each month as the cumulative difference between purchases and consumption up to that month.

FIGURE A.IV: RICE FUTURES

FUTURES PRICES

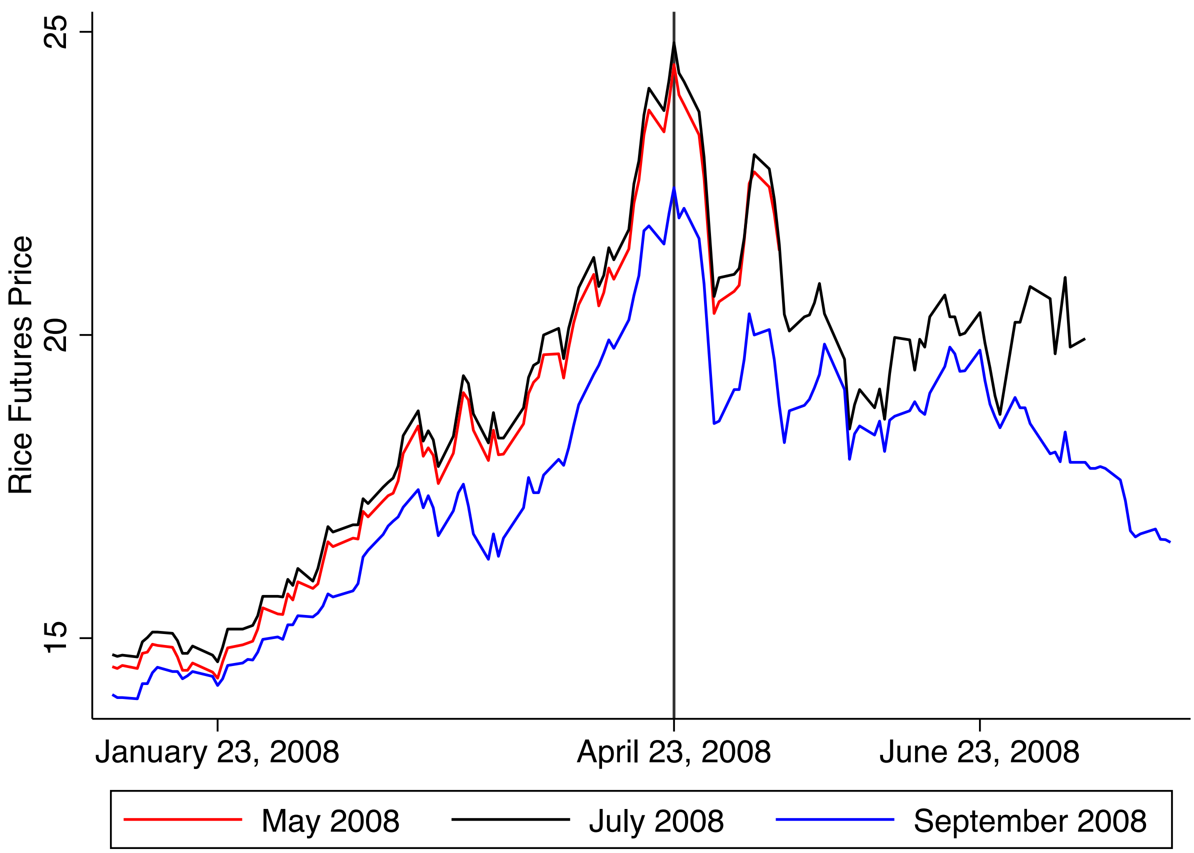

Notes: Figure plots daily close prices for rice futures with expiration in May 2008, July 2008 and September 2008 from the CME. The futures contract is for 2,000 cwt (hundredweight), which corresponds to about 200,000 pounds or circa 91 metric tons, of rough rice, no. 2 or better, and the price quote is in cents per hundredweight. Vertical line denotes April 23, 2008, the peak of prices for all 3 contracts.

Figure A.V: Global Rice Commodity Prices Rise Following India Export Ban

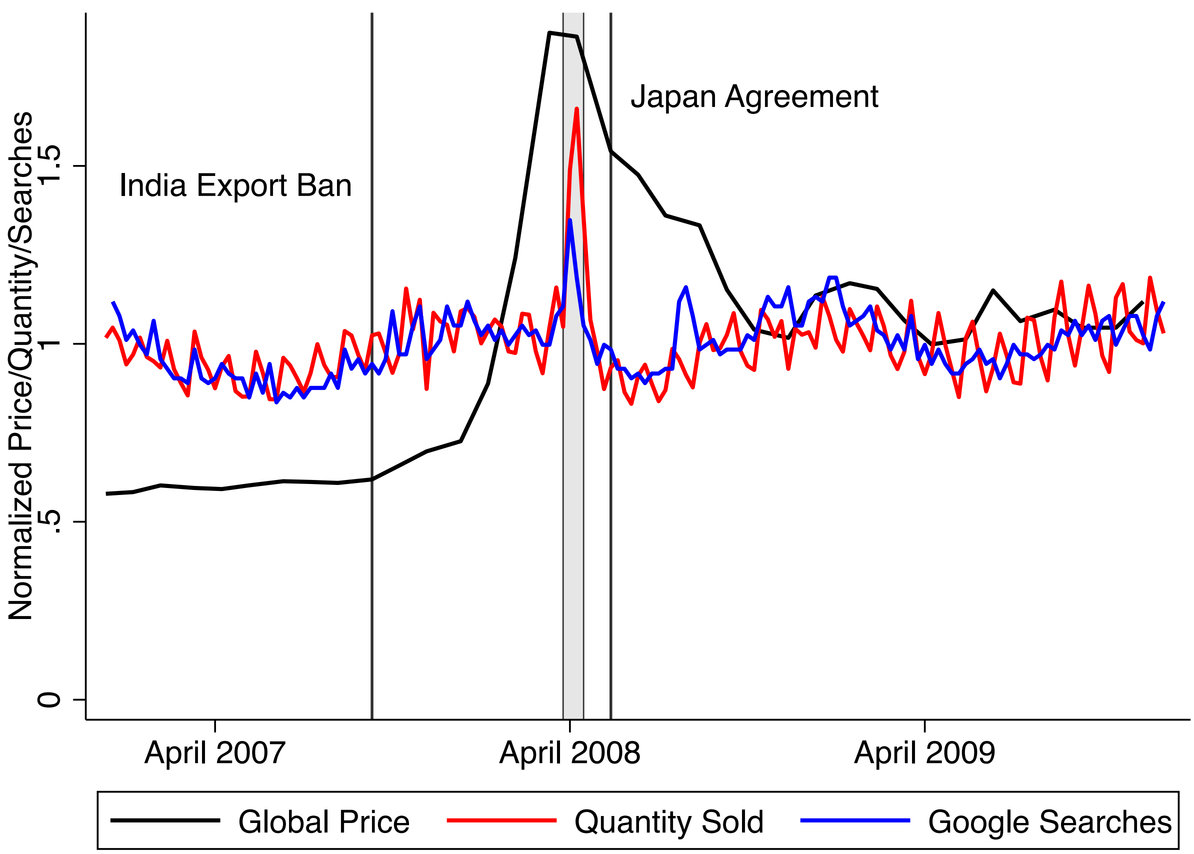

Notes: The black line displays the rice series from the IMFs Primary Commodity Price System. The series measures the Thailand nominal price quote for 5 percent broken milled white rice in USD per metric ton. The red line shows weekly total quantity sold for stores in our sample, and the blue line shows weekly Google search volume. All variables are normalized by the mean over the sample period. The first vertical line denotes India's ban on rice exports in October 2007, while the second vertical line denotes June 2008, to approximate the timing of Japan's agreement to release rice reserves.

FIGURE A.VI: GEOGRAPHIC DISTRIBUTION OF STORES IN REGIONAL DATA

Geographical locations of the stores and demographics of the Zip codes.

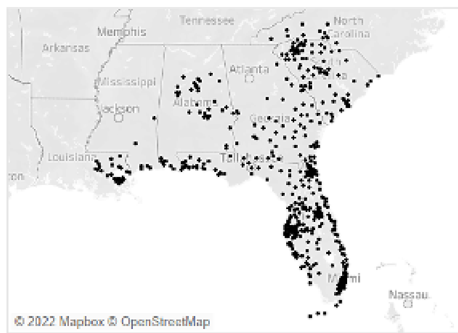

	mean	sd	Min	Median	Max	Histogram
Total.Population	27316.17	14257.13	2373	25542.00	76365	
% White	0.72	0.20	0.02	0.77	1.00	
% Black	0.21	0.20	0.00	0.14	0.95	
% Hispanic	0.14	0.18	0.00	0.08	0.96	
Median.Age	41.51	8.00	20.60	39.75	71.60	
% College Degree	0.57	0.13	0.18	0.56	0.93	
Median.Income	63051.10	19206.99	27380	60116.00	168656	

Notes: This figure presents the geographic location of stores in the data used to analyze the COVID-19 pandemic in Section 6 as well as demographic information for the zipcodes they are located in.

FIGURE A.VII: CROSS-SECTIONAL DIFFERENCES IN PREDICTED PRICE INCREASES VERSUS POST-HOARDING PRICE INCREASES

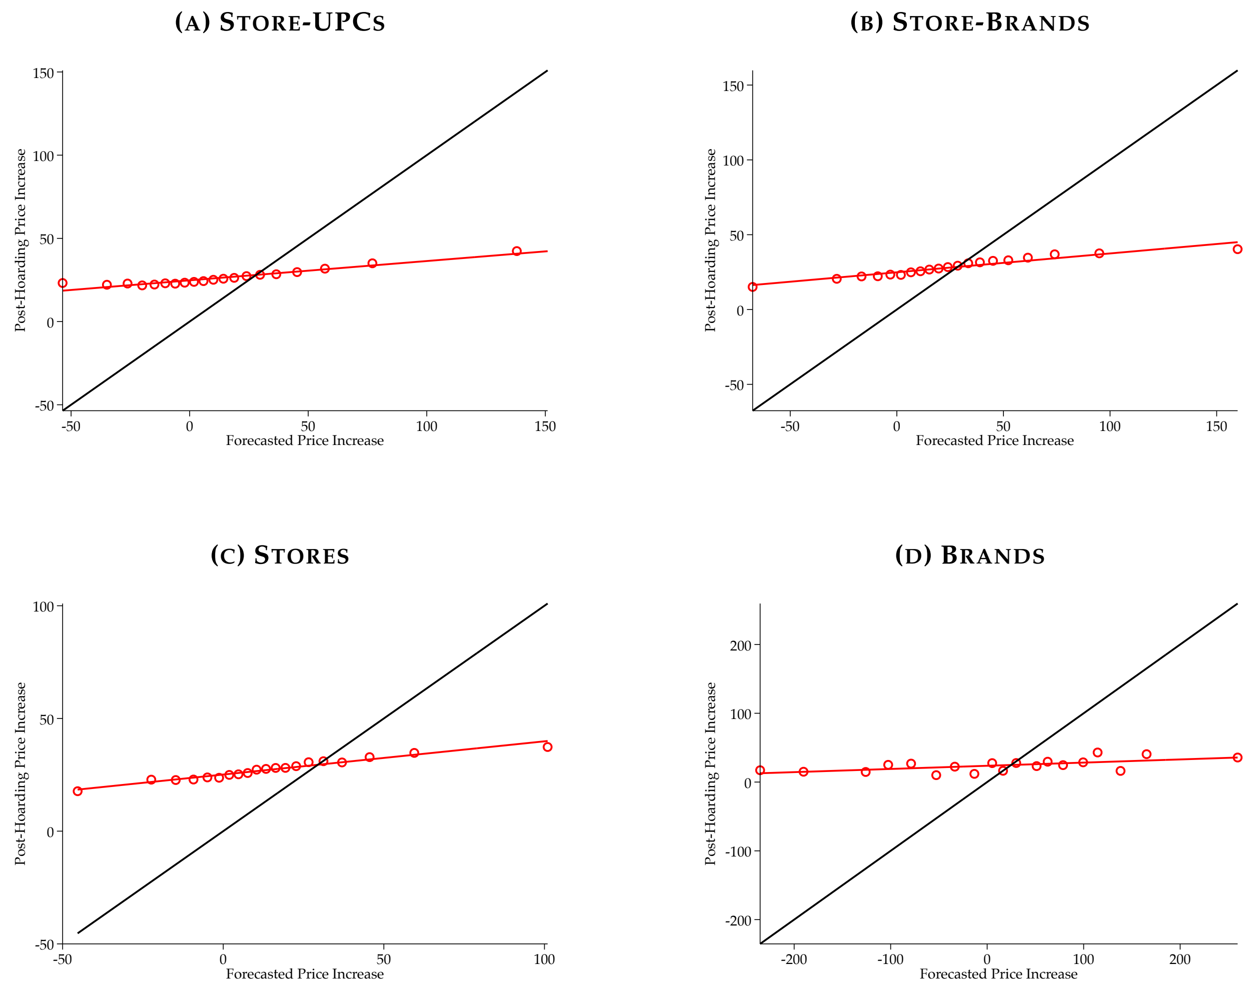

Notes: Binned scatterplots of post-hoarding prices growth on forecasted price growth. Letting $\bar{p}^h$ denote the average unit price in the hoarding period and $\bar{p}^a$ denote the average unit price after the hoarding period, the price increase is defined as $100\times\left(\frac{\bar{p}^a-\bar{p}^h}{\bar{p}^h}\right)$. The forecasted price increases are derived from observed quantity growth during the hoarding period. Specifically: if $\bar{q}^h$ is the average quantity sold during the hoarding period, $\bar{q}^b$ is average quantity sold in the pre-period, and $\frac{1}{\hat{c}}$ is an estimated sale elasticity at the corresponding level of observation from Appendix Table B.I, we define the forecast to be $\left(\frac{\bar{q}^b-\bar{q}^h}{\bar{q}^b}\right)\times\hat{c}$. The sample period covers 2007-2009, and the hoarding period is defined as the weeks of April 19th-May 10th, 2008. Store-UPC includes all store-UPC pairs with at least 5 units sold on average per week in 2007. Store-Brand includes all store-brand pairs with at least 5 units sold on average per week in 2007. Store refers to all stores in the sample while Brand refers to all rice brands with at least 5 units sold on average per week in 2007. Points represent means within ventiles of predicted price discounts. Lines represent a linear fit through the underlying data. All variables are winsorized at the 1 percent level, except the brand level plot which is winsorized at the 5 percent level to avoid stretching the x-axis.

TABLE A.I: ROBUSTNESS FOR CROSS-SECTIONAL RELATIONSHIP BETWEEN QUANTITIES AND PRICES

	Dependent Variable: Post-Hoarding Price Growth						
	Low Sales Variance	High Sales Variance	Sales Weighted	Low Asian Pop.	High Asian Pop.	Low Rice Consumption	High Rice Consumption
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Quantity Growth (%)	0.097^{***}	0.105^{***}	0.081^{***}	0.092^{***}	0.115^{***}	0.103^{***}	0.104^{***}
	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)
Observations	30907	30904	69439	31419	30381	32263	29537

All regressions are at the store-upc level Letting $\bar{p}^h$ denote the average unit price in the hoarding period and $\bar{p}^a$ denote the average unit price after the hoarding period, post-hoarding price growth is defined as $100\times\left(\frac{\bar{p}^a-\bar{p}^h}{\bar{p}^a}\right)$. If $\bar{q}^h$ is the average quantity sold during the hoarding period and $\bar{q}^b$ is average quantity sold in the pre-period we define quantity growth to be $\left(\frac{\bar{q}^b-\bar{q}^h}{\bar{q}^b}\right)\Big/\hat{\delta}$. The sample period covers 2007-2009, and the hoarding period is defined as the weeks of April 19th-May 10th, 2008. Store-UPC includes all store-UPC pairs with at least 5 units sold on average per week in 2007. Price and quantity growth are winsorized at the 1% level. Low vs. high sales variance are store-UPC pars with below vs. above median coefficient of variation of weekly quantity sold in 2007. Sales weighted refers to means and regressions weighted by pre-hoarding period average sales. Low and high Asian population refers to above vs. below median stores in terms of county level Asian population. High and low rice consumption refers to above vs. below median stores in terms of total average weekly rice purchases in 2007 at the county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B Estimating Consumer Responsiveness to Sales

A key input of our forecast test is an estimate of c, the parameter governing storage costs. Our strategy follows from the first order condition shown in Equation 2. We assume that during a typical retail sale or promotion, there is no uncertainty over the size of the sale (and hence the normal price that the product will return to in the future, after the sale concludes). This allows us to set $\sigma_p^2 = 0$ and implies that during a retail sale:

$$I_i^* = \frac{1}{c}(p_{1i} - p_{0i}).$$

Given this, we can estimate the parameter $\delta = \frac{1}{c}$ using data from standard retail sales for rice. In our implementation, this effectively captures the elasticity of demand with respect to the value of a retail sale.

Our primary definition of retail sales or promotions is constructed at the store-UPC level. Following the approach in Hendel & Nevo (2006) we define sales relative to the modal price. Specifically, we consider a store-UPC to be on sale if the price is below the modal price in the corresponding half year period (e.g. January-June 2007 or July-December 2007). Because we simultaneously consider different rice products and sizes, we normalize price changes by p_{0i} and focus the value of sales in percentage terms as our measure of $p_{1i} - p_{0i}$.[27] When aggregating to the store-brand level, store level, or brand level, we take the equal weighted average across all UPCs. Our results are not sensitive to alternative definitions, for example, comparing the current price to the modal price in the preceding 6 months or excluding any promotions that last longer than four weeks.

UPC level prices are equal to the modal price in the corresponding half year period most of the time—just over 70 percent of all store-UPC-weeks observed in our full retail sample. Appendix Figure B.I shows an example of our definition for a single UPC in the pre-hoarding period. The black line denotes periods in which there is not a retail sale while the red line denotes periods with a retail sale. Most UPCs display similar patterns, with sharp and temporary deviations from a relatively stable base price (at varying frequencies).

Our basic specification takes the following form for unit i in week t (where i represents a

[27] Specifically, we define a Sale Value=$\max\{\frac{\text{Modal Price-Price}}{\text{Price}}, 0\}$. This captures the percentage increase in price a consumer expects to occur when the sale ends. We refer to this as the *value* of the sale.

household-brand-store, a store-UPC, a store-brand, a store, or a brand depending on the regression):

$$y_{it} = \delta \text{Sale Value}_{it} + \gamma_i + \eta_t + \varepsilon_{it}. \tag{5}$$

Here, y_{it} represents the quantity of rice. γ_i and η_t represent unit and week fixed effects, respectively. Unit fixed effects are intended to capture cross-sectional differences in the price level, while our sale definition—relative to a time-varying modal price—is intended to capture high frequency time series variation in the price. Ideally, this would leave us with an estimated elasticity $\hat{\delta}$ that largely captures the intertemporal shift in purchases generated by a temporary sale. However, because the response during sales may also be a real consumption response to a lower price, we also explicitly control for the price level in some specifications. We cluster our standard errors at the unit level throughout. We include all weeks from 2007-2009 except for those within the hoarding period itself.

We present results from regressions of this form in Table B.I. Columns 1-3 take the household-brand-store as a unit. We consider all household-brand-store combinations that satisfy the following two criteria: (i) the household purchases that brand in that store at least once during our sample and (ii) the store-brand combination appears in our store level data. As noted in Section 2, this leaves us with nearly 18,000 households purchasing 168 brands at 8,194 stores. The dependent variable is the number of ounces of rice purchased.

The results indicate that households indeed purchase more when rice is on sale. Column 1 suggests that a sale value of 10 percent leads households to purchase 0.12 more ounces for any brand they are observed purchasing in the sample period. This is roughly 30 percent of the mean. Column 2 shows that this sale-responsiveness is effectively unchanged when conditioning on the price level, suggesting that these estimates capture the response to a temporary discount, and not a more general consumption elasticity. Column 3 shows a more flexible specification of the relationship, including dummies for sale vales up to 10 percent, between 10-20 percent, between 20-30 percent, and over 30 percent in place of the linear term, with similar results.

When considering store-UPC level data in Columns 4-6 we see a similar pattern: quantity sold increases when rice is on-sale. For these and all remaining columns in the table, y_{it} represents log(ounces sold).[28] The coefficient in column 4 is 0.012, suggesting that a sale value of 10 percent generates a roughly 12 percent increase in quantity sold. The magnitude drops slightly, to 0.008, when we include a control for the price level, suggesting that a small portion of the estimate in column 4 captures a consumption response and not simply an intertemporal storage motive. Column 6 shows a more flexible specification, mirroring column 3.

The remaining columns of Table B.I show the specification in Equation 5 with the unit defined to be the store-brand, store, or brand. The coefficients are generally consistent with our store-UPC level results, although the point estimates are slightly smaller, particularly at the store-level. We use the results in columns 4, 7, 8 and 9 as inputs when computing forecasts at the store-UPC, store-brand, store, or brand level below.

To summarize, at the Store-UPC level, our estimate $\hat{\delta}$ is 0.012, indicated an elasticity of 1.2. This suggests that if consumers expect a price to increase by 10% when a sale ends, they will increase quantity purchased by 12%. Our estimate $\hat{c}$ is simply $\frac{1}{\hat{\delta}}$

B.0.1 Robustness for estimates of sales elasticity

We next present a series of robustness exercises to support the reliability of our estimates of the sale elasticity δ.

Considering only the pre-crisis period: One concern is that price dynamics surrounding the crisis and hoarding episode might contaminate our estimates of $\hat{\delta}$. To address this, Appendix Table B.II repeats the analysis in Table B.I, but includes only data from 2007. We see similar, if marginally larger, estimates across all specifications.

Using chain level pricing policies to address endogenous sales: A further concern is that our estimates are biased due to classic endogeneity concerns. The sales we identify, might, in principle, by driven by changes in demand, or both might be driven by some omitted factor. While we generally believe our OLS approach to be the simplest and most transparent way of estimating δ, and find the consistency of our results across aggregation levels (and given the rich set of fixed effects) to be reassuring, we conduct robustness exercises here to address a particular form of

[28]We use levels in our household regressions due to the large number of 0s and a substantially smaller degree of dispersion.

endogeneity. Specifically, we consider the possibility that the sales identified by our algorithm are endogenous responses to temporary and local changes in demand. For example, because a sudden drop in consumer demand led a particular store to reduce prices. To address this, we construct a proxy in the spirit of Hausman (1996) that exploits the uniformity of pricing policies within supermarket chains (as documented in DellaVigna & Gentzkow, 2017). Our proxy is a leave-store-out measure of sales at the supermarket chain level.[29] This measure addresses concerns that store-specific or other local demand shocks are driving pricing, but may of course still suffer from similar endogeneity issues if demand shocks are correlated across different stores within the same chain.[30]

Appendix Table B.III repeats the analysis in Table B.I, but includes our leave-store-out proxy in place of the sale variable. The results are very similar to those when using our sale variable, likely due to the uniformity of pricing policies within supermarket chains (although we have fewer observations as we are unable to identify the chain for all stores in our sample).[31] While this approach does not resolve all potential endogeneity concerns, the consistency with our main specifications provides reassurance that focusing on OLS estimates is reasonable.

Non-parametric estimates of sale elasticities:. Our primary specification imposes a relatively restrictive linear model. Of course, it is possible that the true underlying relationship is highly non-linear, particular as sales become large or in the region close to 0 sale. Appendix Figure B.II shows that linearity appears to be a reasonable approximation in our context. This figure presents a binned scatter plot of log(ounces) sold against sales at the store-UPC level, with each dot representing the mean within a percentile.

Cross-price effects: When converting excess purchases into forecasts, one potential issue is distortions due-to cross-price effects. For any given product, excess purchases during the hoarding episode might be driven not just by expectations of coming price changes for that product, but also

[29] The leave out mean for a given UPC is defined as the average sale on that UPC in that week for other stores in the same chain as the store in question. This captures the degree to which a UPC was on sale at other stores within the same chain. The average of this leave out mean is 2.5 percent, with a standard deviation of 5.8 percent. When aggregating to the Store-Brand or Store, we take the equal weighted average leave out mean sale across UPCs.

[30] Note that our results are similar if we impose geographic constraints on our leave-out proxy, for example, considering only stores in the same chain located in other counties or states.

[31] In fact, IV specifications that instrument for our sale variable with this leave-store-out proxy give nearly identical results, with first stage coefficients very close to one.

by changes in relative prices across products (either contemporanous relative prices or expected future relative prices). If promotions drive meaningful cross-product substitution, this would constrain our ability to recover expectations about price changes from observed quantity changes without fully estimating the demand system.

Appendix Table B.IV suggests that there are negligible cross-price impacts of promotional sales in our sample. This table presents versions of the specification shown in Equation 5 estimated at the household-brand-store (i.e. the specifications shown in Columns 1-3 of Table B.I). However, as a dependent variable, these regressions include total ounces purchased by the household *at all other brands*, excluding the brand in question. This captures the relationship between a sale in a given brand and purchases of all other brands. Across specifications, we see fairly tightly estimated 0 effects, indicating limited impacts of sales for one product on purchases of another at the household level. One potentially explanation for this is that our estimates largely reflect intertemporal substitution of purchases, rather than meaningful changes in consumption. Regardless, while it is still possible that there are meaningful cross-product impacts of promotions, these estimates suggest that choosing to set such effects aside as a first approximation is reasonable.

Other specifications: Our results are also robust to a number of alternative specifications not shown here, including eliminating or varying the level of fixed effects, controlling for consumer inventories, considering only temporary sales (two weeks or less), and using a rolling or backwards looking window to define the modal price.

FIGURE B.I: EXAMPLE OF RETAIL SALES

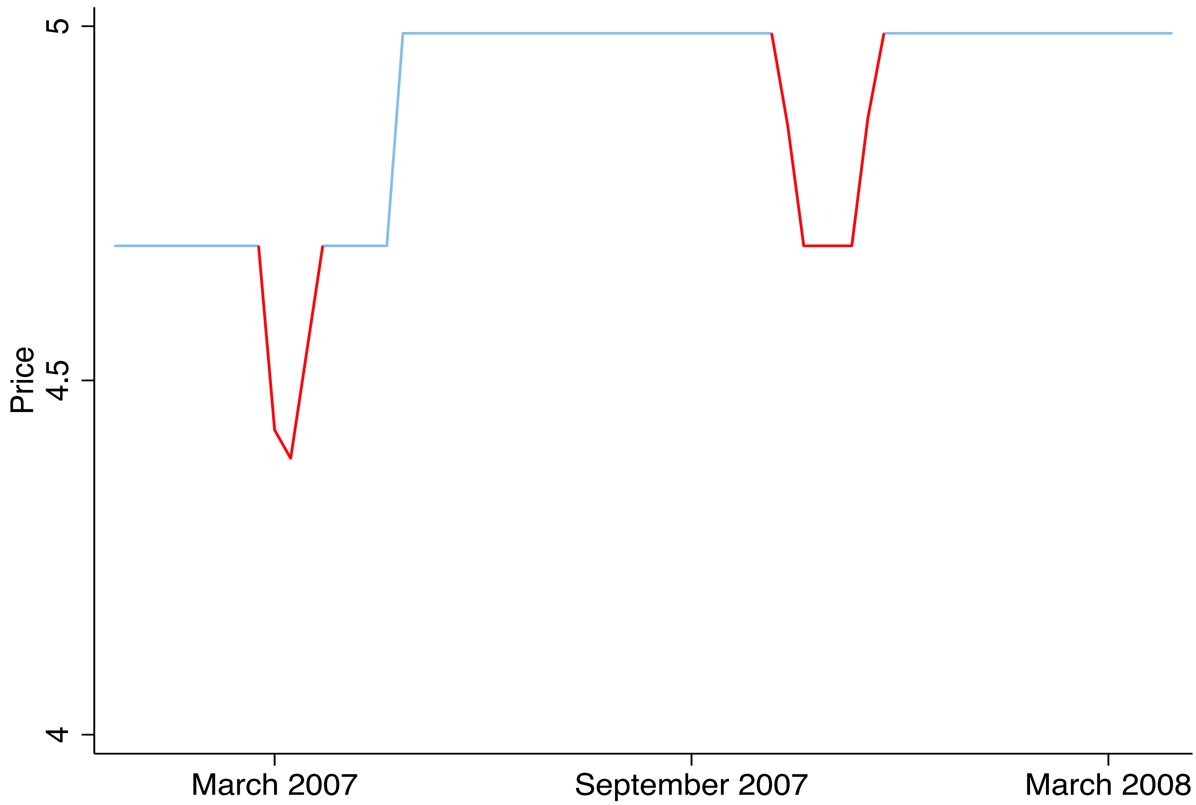

Notes: Example price path for a UPC in our sample in the pre-hoarding period (January 1st 2007-April 19th, 2008). Red portions indicate sales as identified by our algorithm.

FIGURE B.II: LINEAR APPROXIMATION TO SALE RESPONSIVENESS

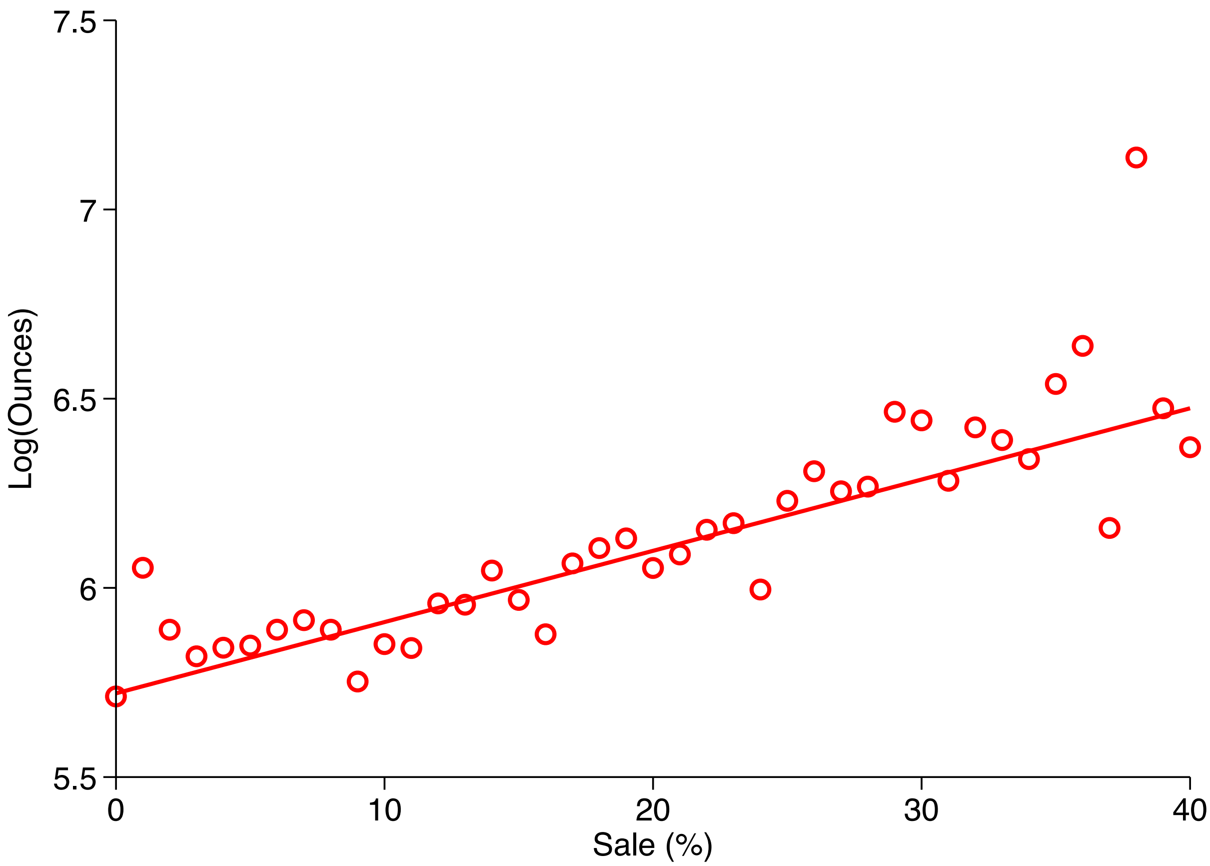

Notes: Binned scatter plot of store-UPC level log ounces of rice sold on percentage sale. Sample is a weekly balanced panel from 2007-2009 of all observed store-upc pairs with an average of more than 5 units sold per week in 2007. Percentage sale is the percentage discount of the weekly price relative to the modal price in the corresponding half year period (January-June vs. July-December). Scatter plot shows each percentage sale between 1-50. Red line shows a linear fit through the raw data.

TABLE B.I: CONSUMER RESPONSES TO RETAIL SALES

	Quantity Purchased (oz)			Log(Ounces Sold)					
	HH-Brand			Store-UPC			Store-Brand	Store	Brand
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sale Value	0.012^{***} (0.002)	0.011^{***} (0.002)		0.012^{***} (0.001)	0.008^{***} (0.001)		0.011^{***} (0.001)	0.010^{***} (0.000)	0.004 (0.005)
Sale Value: 0-10 Percent			0.099^{***} (0.019)			0.168^{***} (0.008)			
Sale Value: 10-20 Percent			0.216^{***} (0.039)			0.255^{***} (0.016)			
Sale Value: 20-30 Percent			0.284^{***} (0.078)			0.403^{***} (0.017)			
Sale Value: 30+ Percent			0.418^{***} (0.104)			0.631^{***} (0.036)			
Unit Price		-0.022^{**} (0.011)			-0.123^{***} (0.024)				
Mean of Dep. Var.	0.57	0.57	0.57	5.77	5.77	5.77	6.61	7.43	9.82
Observations	4466506	4466506	4466520	10465608	10465608	10465671	6614629	1420020	22290
HH-Store-Brand FE	Yes	Yes	Yes	No	No	No	No	No	No
Store-UPC FE	No	No	No	Yes	Yes	Yes	No	No	No
Store-Brand FE	No	No	No	No	No	No	Yes	No	No
Store FE	No	No	No	No	No	No	No	Yes	No
Brand FE	No	No	No	No	No	No	No	No	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Regressions of weekly ounces purchased (at the household-brand-store level) or log quantity sold (at the Store-UPC, Store-Brand, Store and Brand levels) on the sale value in percentage term. A sale is defined as a price below the modal price in corresponding half year period (January-June or July-December) for a given UPC in a given store. If p_m is the modal price, and p_s is the sale price, the sale value is defined as $\frac{p_m - p_s}{p_s}$. Sale values are set to 0 if the price is at or above the modal price. When aggregating to the Store-Brand, Store, or Brand, we take the equal weighted average sale across UPCs. Samples are weekly balanced panels from 2007-2009 omitting the hoarding period. The HH-Brand Sample consists of all household-brand-store combinations that satisfy the following two criteria: (i) the household purchases that brand in that store at least once during our sample and (ii) the store-brand combination appears in our store level data. Store-UPC includes all store-UPC pairs with at least 5 units sold on average per week in 2007. Store-Brand includes all store-brand pairs with at least 5 units sold on average per week in 2007. Store refers to all stores in the sample and Brand refers to all rice brands with at least 5 units sold on average per week in 2007. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE B.II: CONSUMER RESPONSES TO RETAIL SALES (2007 ONLY)

	Quantity Purchased (oz)			Log(Ounces Sold)					
	HH-Brand			Store-UPC			Store-Brand	Store	Brand
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sale Value	0.015^{***} (0.002)	0.013^{***} (0.002)		0.013^{***} (0.001)	0.009^{***} (0.002)		0.012^{***} (0.001)	0.013^{***} (0.000)	0.012^{**} (0.005)
Sale Value: 0-10 Percent			0.096^{***} (0.029)			0.163^{***} (0.010)			
Sale Value: 10-20 Percent			0.201^{***} (0.023)			0.297^{***} (0.017)			
Sale Value: 20-30 Percent			0.274^{**} (0.131)			0.453^{***} (0.025)			
Sale Value: 30+ Percent			0.737^{***} (0.177)			0.651^{***} (0.035)			
Unit Price		-0.025^{**} (0.012)			-0.121^{***} (0.044)				
Mean of Dep. Var.	0.59	0.59	0.59	5.81	5.81	5.81	6.63	7.51	9.92
Observations	1528019	1528019	1528020	3638944	3638944	3638961	2247243	477294	7775
HH-Store-Brand FE	Yes	Yes	Yes	No	No	No	No	No	No
Store-UPC FE	No	No	No	Yes	Yes	Yes	No	No	No
Store-Brand FE	No	No	No	No	No	No	Yes	No	No
Store FE	No	No	No	No	No	No	No	Yes	No
Brand FE	No	No	No	No	No	No	No	No	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Regressions of weekly ounces purchased (at the household-brand-store level) or log quantity sold (at the Store-UPC, Store-Brand, Store and Brand levels) on percentage sales. A sale is defined as a price below the modal price in corresponding half year period (January-June or July-December) for a given UPC in a given store. If p_m is the modal price, and p_s is the sale price, the sale value is defined as $\frac{p_m - p_s}{p_s}$. Sale values are set to 0 if the price is at or above the modal price. When aggregating to the Store-Brand, Store, or Brand, we take the equal weighted average sale across UPCs. Samples are weekly balanced panels for 2007. The HH-Brand Sample consists of all household-brand-store combinations that satisfy the following two criteria: (i) the household purchases that brand in that store at least once during our sample and (ii) the store-brand combination appears in our store level data. Store-UPC includes all store-UPC pairs with at least 5 units sold on average per week in 2007. Store-Brand includes all store-brand pairs with at least 5 units sold on average per week in 2007. Store refers to all stores in the sample and Brand refers to all rice brands with at least 5 units sold on average per week in 2007. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE B.III: SALES AND UNIT PRICE RESPONSE TO QUANTITY OF RICE PURCHASED (LEAVE OUT MEAN OF SALES)

	Quantity Purchased (oz)			Log(Ounces Sold)				
	HH-Brand			Store-UPC			Store-Brand	Store
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sale Value	0.014^{***}	0.013^{***}		0.016^{***}	0.011^{***}		0.017^{***}	0.010^{***}
	(0.003)	(0.003)		(0.001)	(0.001)		(0.002)	(0.000)
Sale Value: 0-10 Percent			0.079^{***}			0.065^{***}		
			(0.020)			(0.006)		
Sale Value: 10-20 Percent			0.256^{***}			0.299^{***}		
			(0.032)			(0.019)		
Sale Value: 20-30 Percent			0.295^{***}			0.440^{***}		
			(0.089)			(0.017)		
Sale Value: 30+ Percent			0.372^{***}			0.739^{***}		
			(0.139)			(0.048)		
Unit Price		-0.026^{**}			-0.123^{***}			
		(0.011)			(0.023)			
Mean of Dep. Var.	0.57	0.57	0.57	5.77	5.77	5.77	6.61	7.43
Observations	4466520	4466520	4466520	10417265	10417265	10417265	6595963	1420020
HH-Store-Brand FE	Yes	Yes	Yes	No	No	No	No	No
Store-UPC FE	No	No	No	Yes	Yes	Yes	No	No
Store-Brand FE	No	No	No	No	No	No	Yes	No
Store FE	No	No	No	No	No	No	No	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Regressions of weekly ounces purchased (at the household-brand-store level) or log quantity sold (at the Store-UPC, Store-Brand and Store levels) on the leave-out-mean percentage sales at other stores in the same chain as the store in question. A sale is defined as a price below the modal price in corresponding half year period (January-June or July-December) for a given UPC in a given store. If p_m is the modal price, and p_s is the sale price, the sale value is defined as $\frac{p_m - p_s}{p_s}$. Sale values are set to 0 if the price is at or above the modal price. The leave out mean for a given UPC is defined as the average sale value on that UPC in that week for other stores in the same chain as the store in question. When aggregating to the Store-Brand or Store, we take the equal weighted average leave-out sale across UPCs. Samples are weekly balanced panels from 2007-2009 omitting the hoarding period. The HH-Brand Sample consists of household-brand-store combinations that satisfy the following two criteria: (i) the household purchases that brand in that store at least once during our sample and (ii) the store-brand combination appears in our store level data. Store-UPC includes store-UPC pairs with at least 5 units sold on average per week in 2007. Store-Brand includes store-brand pairs with at least 5 units sold on average per week in 2007. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE B.IV: CROSS-SUBSTITUTION IN RESPONSE TO RETAIL SALES

	Quantity Purchased (oz)	
	HH-Brand	
	(1)	(2)
Sale Value	-0.000	-0.000
	(0.003)	(0.003)
Unit Price		-0.000
		(0.005)
Mean of Dep. Var.	2.00	2.00
Observations	4466506	4466506
HH-Store-Brand FE	Yes	Yes
Week FE	Yes	Yes

Regressions of total weekly ounces purchased by a household excluding a given brand-store pair on the sale value for that brand. A sale is defined as a price below the modal price in corresponding half year period (January-June or July-December) for a given UPC in a given store. If p_m is the modal price, and p_s is the sale price, the sale value is defined as $\frac{p_m - p_s}{p_s}$. Sale values are set to 0 if the price is at or above the modal price. When aggregating to the brand level, we take the equal weighted average sale across UPCs. Samples is a weekly balanced panel 2007-2009 omitting the hoarding period. We include all household-brand-store combinations that satisfy the following two criteria: (i) the household purchases that brand in that store at least once during our sample and (ii) the store-brand combination appears in our store level data. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.