

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

August 31, 2021

Abstract

We show that sticky prices exacerbate household hoarding of storable goods. When stores are slow to adjust prices following a cost shock, households have an incentive to stockpile just as in a typical retail sale. This incentive is present even in the absence of traditional panic or precautionary motives for hoarding. Using detailed US supermarket scanner data covering the 2008 global rice crisis—a shock triggered by an Indian rice export ban—we find that household hoarding anticipated retail price adjustments. We construct forecast tests relating the cross-section of product or store-level price adjustments to the expectations implied by consumer purchases. Bias and efficiency tests reject panic/precautionary motives in favor of a sticky-price view.

*This paper subsumes our earlier work on hoarding and commodity bubbles. We thank seminar participants at 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 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

Household hoarding of staples—the accumulation of inventories during disasters or other supply disruptions—has long been a concern for policymakers. Roman Emperor Julian blamed hoarding for artificial shortages and famine in Antioch as far back 362 A.D (Ó Gráda, 2008). More recently, the COVID-19 pandemic generated a widely documented run on staple foods, hand sanitizer and masks. Household hoarding is also thought to have exacerbated the 2000s commodities boom and the 2008 global rice crisis—the focus of this paper.

Discussions of hoarding often emphasize consumer panic or its close cousin, precaution. In these narratives, behavioral considerations or a desire to hedge market uncertainty lead households to increase demand by more than what is justified by current prices (or by risk-neutral rational expectations of future prices). Indeed, the prevailing diagnosis of the 2008 global rice crisis by experts and the media centers on panic buying:

The psychology of hoarding behaviour is important in explaining why rice prices suddenly shot up...decisions by millions of households...sparked a sudden surge in demand for rice and changed the gradual increase in rice prices from 2002 to 2007 into an explosion...panicked hoarding caused the rice price spike.¹

Similar explanations are common for other episodes, including the energy crisis of the 1970s and the COVID-19 pandemic.²

In this paper we highlight an alternative mechanism for household hoarding—sticky store prices—and show that it meaningfully drove excess purchases during the 2008 global rice crisis. A large literature documents retailers' tendencies to delay price adjustments after shocks (Nakamura & Steinsson, 2008), including disasters such as earthquakes, hurricanes or snowstorms that cause costs to rise (Cavallo *et al.*, 2014; Gagnon & Lopez-Salido, 2015). Heightened consumer outrage and the imposition of anti-price gauging laws during crises may further amplify retailers' motivations for pursuing sticky price strategies, which include concerns over customer anger (Anderson

¹The Rice Crisis: Markets, Policies and Food Security, D. Dawe (ed.), 2010, FAO, p.42. See also, "How Fear Turned A Surplus into Scarcity," *National Public Radio*, November 4, 2011.

²On 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." On the pandemic, see, e.g., "Coronavirus: The psychology of panic buying," *BBC*, March 4, 2020.

& Simester, 2010; Anderson *et al.*, 2017; Rotemberg, 2005) in addition to menu costs and costly attention or information gathering (Mankiw & Reis, 2002; Woodford, 2009).³ In other words, theory and evidence suggest that retail prices are likely to be sticky during the market disruptions that often precipitate hoarding episodes.

When these disruptions involve a cost shock—or any other shock expected to influence prices in the longer term—sticky prices provide a direct incentive for hoarding. This is the case even in absence of consumer panic or any precautionary motive. Recognizing that prices will rise in the future, households will be driven to shift demand intertemporally and stockpile storable goods. This incentive—which we term the *implicit promotion* effect—mimics the one provided by a standard retail sale (Gönül & Srinivasan, 1996; Erdem *et al.*, 2003; Sun *et al.*, 2003; Sun, 2005; Hendel & Nevo, 2006b).⁴

A major impediment to studying the drivers of household hoarding has been a lack of sufficiently detailed micro-data. To investigate the relationship between sticky prices and hoarding, it is necessary to go beyond aggregates and wholesale prices and observe the precise *shelf* prices consumers face (as well as the quantities they purchase). This poses a difficulty, as many large-scale hoarding episodes are historic or in developing countries where detailed micro-data is difficult to come by.

We overcome this limitation by using US supermarket scanner data covering the 2008 global rice crisis. These data provide sufficient granularity (e.g. prices and quantities at the product-store-week level) to capture the prices consumers saw on the shelf. While no doubt the availability of scanner data covering the COVID-19 pandemic and other recent episodes will permit detailed future research, we believe the incident we study is the first instance of large scale hoarding for

³The first US state law directed at price gouging was enacted in New York in 1979, during a period of commodity market instability. Since then anti-gouging laws have proven popular—the majority of states have enacted some type of regulation (Davis *et al.*, 2008). These measures are typically adopted during or following commodity shortages (see, e.g., Zwolinski, 2008; Giberson, 2011).

⁴This mechanism also echos the theory of Benabou (1989), which notes the related possibility of retail arbitrageurs, agents that take advantage of the constraints on reputable retailers by stockpiling and reselling at a later date through informal channels (i.e. online or through smaller retailers that are less bound by reputation norms). Such retail arbitrage strategies are akin to speculative or resale motive discussed in commodity markets (Scheinkman & Schechtman, 1983; Deaton & Laroque, 1992) that apply even in flexible price settings. In general, it is possible for there to be simultaneously sticky prices and bounded speculative storage in a game between consumers and firms when there are menu costs. Benabou (1989) extends classic menu cost models (i.e. fixed costs of price adjustment following Barro, 1972; Sheshinski & Weiss, 1977), which feature firm’s nominal costs rising due to inflation, to allow for storeable goods. The well-known sticky pricing strategy, or (S,s) rule, derived under non-storeability is shown to hold when there are only moderate amounts of speculative storage.

which comparably fine-grained consumer and retailer data are available.

The initial part of our paper uses these data to make two broad points. First, despite a sharp increase in wholesale prices, shelf prices for rice products were uniformly sticky during the 2008 global crisis. The crisis itself was sparked by a relatively pure 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. On the other hand, shelf prices largely remained constant through the most intense periods of the crisis in April and May. It was not until later that they gradually rose to a higher level (in line with the permanent increase in wholesale prices).⁵ A look at long-run retail prices suggests that that sticky prices generated an implicit promotion of roughly 22 percent. In other words, shelf prices were 22 percent below their long-run level during the crisis.⁶

Second, substantial household hoarding took place precisely as shelf prices stagnated relative to wholesale prices. Extreme excess purchases coincided with or slightly lagged the increase in commodities prices: retail store sales were 40 percent above normal levels in late April and early May of 2008. Elevated Google search volumes (Choi & Varian, 2012; Da *et al.*, 2011) for the term "rice price" were present alongside these abnormally high store sales, pointing to consumer awareness of the cost shock (Goel *et al.*, 2010). These searches were likely prompted by significant media coverage of rice markets in mid-April (Fang & Peress, 2009; Engelberg & Parsons, 2011). Naively comparing wholesale prices to retail sales, one might easily conclude that consumers were panicking in the face of rising expenses. However, the most severe hoarding *preceded* any increase in retail prices, suggesting that consumers responded to the wholesale price shock by purchasing rice in bulk at relatively low cost before shelf prices rose.

The timing of purchases is consistent with an *implicit promotion* motive driving household

⁵There was no corresponding shock to public health or safety that might trigger a simultaneous change in demand (as in a pandemic or natural disaster). The lack of retail price response in the face of a clear cost or *supply* shock highlights a key contribution of our study. Discussions of pricing during hoarding episodes typically focus on the dangers (or lack thereof) of gouging by retailers—with price-gouging usually referring to sharp increases in prices to match increased hoarding-driven *demand* (see <https://www.nytimes.com/2020/03/27/us/coronavirus-price-gouging-hand-sanitizer-masks-wipes.html>, <https://hbr.org/2013/07/the-problem-with-price-gouging-laws>, or <https://www.ftc.gov/public-statements/2006/02/moneyball-and-price-gouging>). Hesitancy to increase prices even to match cost changes observed in our data suggests that concerns over gouging (or, alternatively, concerns over the restrictiveness of anti-gouging regulations) may be misplaced, at least for the relatively reputable retailers observed in our data.

⁶The long-run increase corresponds to an approximately 40 percent passthrough from wholesale prices to retailers. This is similar to what studies have found for other storable goods such as coffee (Nakamura & Zerom, 2010).

hoarding: households shifting demand dynamically to take advantage of temporarily low (sticky) prices during the crisis. Indeed, a non-trivial portion of observed excess purchases can be explained as a normal household response to a 22 percent sale. Heterogeneity in the degree of hoarding across households with different ex-ante inventories further supports the implicit promotion motive. Just as households with larger inventories are less responsive to standard retail sales (Hendel & Nevo, 2006b), households with larger inventories hoarded less intensely during the episode.

Of course, these patterns—based in the time series dynamics of price and quantity changes—do not conclusively show that hoarding was driven by an implicit promotion motive rather than consumer panic or precaution. It could be, for example, that panic drove consumers to buy with no concern about future retail prices, and that prices coincidentally rose after hoarding subsided.

In the latter part of our paper we develop a series of forecast tests to distinguish the implicit promotion motive for hoarding from panic or precautionary motives. These tests exploit the cross-section of price changes and consumer purchases at the product and store level. We document considerable dispersion in post-crisis price changes, reflecting heterogeneity in the exposure of different products to the aggregate cost shock, different ex-ante pricing policies, and more. The logic of our tests is that—under an implicit promotion motive—excess purchases during the crisis should reflect latent rational expectations about coming price increases. Both panic and precautionary motives have distinct implications for the degree of hoarding relative to price changes.⁷

Our approach occurs in three steps. First, we recover the cross-sectional price expectations implied by excess purchases during the crisis. To do so, we estimate household elasticities to standard promotions outside of the hoarding period. In line with earlier work on promotions for other storeable goods, we see evidence of frequent short term promotional sales in rice across different products and stores and observe meaningful consumer responses to these promotions. We find sale elasticities on the order of 1.5, implying that a 10 percent sale generates a 15 percent increase in quantity purchased. We then use these elasticities to infer consumer forecasts of future price changes from excess purchases at the product and store level (under the strong assumption that excess purchases are driven by price expectations).⁸ This provides a full cross-section of bench-

⁷Our tests are in the spirit of Mincer & Zarnowitz (1969); Nordhaus (1987) style rational forecast test but utilize cross-sectional variation as opposed to time-series variation.

⁸For example, if we observe a 15 percent increase in purchases for a given store×UPC during the hoarding period,

mark price forecasts that we can compare against the realized path of prices.

Our second step is a set of forecast bias tests that allow us to reject panic and precaution as key drivers of hoarding in the crisis. If consumers are unbiased, and excess purchases are indeed driven by price expectations, the forecasted price changes we infer in our first step should be equal to realized price changes, on average. Both panic based and precautionary models of consumer hoarding give an alternative prediction: that forecasts should be higher than realized price changes. The logic of this prediction is simple: if hoarding is driven by behavioral overreaction (panic) or risk aversion (precaution), consumers must—on average—be purchasing more than would be predicted by a rational, risk neutral benchmark. We find, if anything, that forecasts fall below realizations on average. This result is inconsistent with a panic or precautionary motive, holds across levels of aggregation (at the product, store, and brand level), and is not driven by truncation due to stockouts.

In the third step, we implement forecast efficiency tests that provide support for an implicit promotion (sticky price) motive. The key intuition is that implicit promotions should lead households to increase purchases the most for the products and stores where they expect to see the largest future price increases. This suggests that the cross-section of excess purchases should be predictive of subsequent price increases. We find that this is the case. Across products and stores, the degree of hoarding during the crisis is robustly correlated with post-hoarding price increases. Formal efficiency tests that regress realizations on forecasts show highly significant positive coefficients of over 0.1. We consider and reject a number of alternative explanations for these patterns, including reverse causality, the degree of hoarding as a signal of a long run demand shift, and consumers targeting low-cost products when hoarding.

Taken together, these bias and efficiency tests (and our time-series evidence) reject panic or precautionary models in favor of a sticky-price view. The timing of purchases and the cross-sectional relationships between hoarding and later price increases suggest that households anticipated slow retail price updates. They did not overreact, on average, relative to the realized path of prices. This evidence suggests that firm pricing policies may lead to consumer stockpiling during crises, and that common policy approaches—like anti-price-gouging regulations—may further exacerbate hoarding. Moreover, our tests can be used to evaluate the importance of sticky prices versus

we infer that consumers expected that prices were 10 percent below the long run level.

other motives for hoarding during the COVID-19 pandemic and other episodes.

The remainder of the paper is organized as follows: Section 2 presents our data, Section 3 provides background on the rice crisis and evidence on the timing of purchases relative to price increases, Section 4 presents our cross-sectional forecast tests, and Section 5 concludes.

2 Data

Our primary sources of retailer and household 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.⁹

Retailer Data

For our store data we consider food retail channels only. This leaves us with 10,561 unique stores. These data contain weekly store-UPC 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 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 43953 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).

⁹UPCs with product module code 1319.

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 Universal Product Code (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 2008.¹⁰ 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.

¹⁰See the U.S. Census Bureau, Current Population Survey, 2008 Annual Social and Economic Supplement.

3 Stylized Facts: Sticky Store Prices and Hoarding During the Crisis

3.1 Overview of the Rice Crisis and Commodity Price Dynamics

Rice commodities 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 in a number of exporters.¹¹ 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-2008.¹² 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 shows a proxy for the global price of rice on the commodities markets from the IMF, highlighting the crisis and associated events.¹³ 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 the Japan agreement. Despite 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 commodities 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

¹¹See, 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.

¹²A 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.

¹³This line presents the rice series from the IMF's 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>.

fraction of stores that updated retail prices in the wake of the shock to commodities 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.¹⁴ 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 gradual and standard price increases relative to 2007.

Most 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 the hoarding period (shown in gray).¹⁵ 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.¹⁶ 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 storable goods (Leibtag, 2007). These patterns are consistent with a large literature in macroeconomics: the retail or supermarket prices that consumers face are sticky

¹⁴While this threshold is somewhat arbitrary, we get similar patterns when we consider different cut-offs, e.g. 110 percent.

¹⁵The 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

¹⁶We 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.

and tend to lag changes in commodities prices.¹⁷

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.¹⁸

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 levels similar to 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

¹⁷Although [McShane et al. \(2016\)](#) show that a larger fraction of positive wholesale price changes are eventually passed on to consumers.

¹⁸See, e.g. <https://www.reuters.com/article/uk-philippines-rice/philippines-arroyo-leads-crackdown-on-rice-hoarding-idUKMAN1898020080508>.

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, given 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.¹⁹

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.

¹⁹Futures 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.

3.5 The Implicit Promotion 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.²⁰ Classic narratives along these lines have consumers panicking just as prices skyrocket, leading them to purchase in large quantities at high prices.²¹ 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, in practice, 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 implicit promotion 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.

A first question is whether the observed magnitude of excess purchases is even plausibly explained by an implicit promotion motive. Retail prices in the hoarding period were approximately 22 percent below the post-hoarding average price, while purchases were roughly 40 percent above average. If consumers were forward looking and interpreted prices during the hoarding period as a 22 percent discount, a promotional elasticity of just above 2 would be sufficient to generate the entirety of observed hoarding (even in the absence of any panic or precautionary motive). Such an elasticity is not entirely out of line with the literature,²² and is in fact close to the elasticities we estimate in Section 4. This suggests that consumer hoarding in our episode is at least feasibly the result of the implicit promotion driven by sticky prices.

²⁰Much 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).

²¹See, e.g., <https://spectrumlocalnews.com/nc/charlotte/news/2021/05/11/higher-prices--panic-buying--what-the-colonial-pipeline-hack-means-for-north-carolina>

²²Bergtold *et al.* (2004) estimate unconditional price elasticities of roughly -1 that do not factor in the additional dynamic incentive provided by temporary sales.

3.6 Role of Inventories

An implicit promotion effect is the result of anticipatory demand: households shifting purchases forward in time to build up inventories that can be consumed later. This is analogous to the inventory models invoked in [Hendel & Nevo \(2006b,a\)](#) when studying retail sales. This mechanism—under standard assumptions of costly or depreciating inventory—implies a testable relationship between consumer inventory and the degree of hoarding during the crisis. Specifically, consumers who have built up greater inventories just before the crisis begins will be less likely to hoard (or will hoard less excessively). The prediction of a panic/precaution narrative regarding the role of inventories is ambiguous.²³

To test this, we develop a proxy for inventory in our household data following [Hendel & Nevo \(2006b\)](#). For each household, we first estimate weekly consumption as the average weekly purchase of rice in 2007. We then construct household inventories as the cumulative weekly difference between consumption and purchases. To ensure that our measure of inventory stays between 0 and 1, we subtract the minimum inventory level and divide by the maximum level for each household. We construct this measure both at the household level (measuring total rice inventories), and at the household-brand level (measuring inventories of each specific brand). Given our definition, we only consider households (or household-brand combinations) for which we see purchases in 2007.

With this measure in hand, we estimate the following regression specification to assess whether household inventories dampened the extent of hoarding:

$$y_{it} = \alpha \text{Hoarding Period}_t + \beta (\text{Hoarding Period}_t \times \text{Pre-Hoarding Inventory}_i) + \gamma_i + \varepsilon_{it}. \quad (1)$$

Here y_{it} is a measure of consumer purchases (either an indicator equal to one for any purchase, or a total quantity purchased in ounces) for household i in week t . Hoarding Period_t is an indicator equal to one in the hoarding period, and $\text{Pre-Hoarding Inventory}_i$ is our inventory measure in the week prior to the hoarding period (the week of April 12). γ_i is a household fixed effect.

²³On the one hand, households with existing inventories might feel sufficiently hedged against future price uncertainty and not accumulate as much, consistent with our findings. But if households' inventories were not meant as a buffer against the price uncertainty, households with inventories might end up accumulating even more to the extent they are more risk averse.

Consistent with an inventory model, we find that households with higher ex-ante inventories were less likely to hoard during the crisis. We present our results in Table 3. Columns 1 and 3 show results omitting the interaction term from Equation 1 and confirm that, on average, households were more likely to purchase rice, and purchased more rice, during the hoarding period. The key results in this table are the significant negative coefficients on the interaction terms in columns 2 and 4: higher ex-ante inventories decreased the likelihood households purchased rice (and the quantity of rice purchased). Columns 5-8 repeat this analysis at the household-brand level, and display the same results. Within a given brand, households with higher ex-ante inventories displayed more muted hoarding behavior. As a whole, the relationship between consumer inventory and hoarding behavior is consistent with an implicit promotion motive.²⁴

4 Forecast Tests: Hoarding Intensity in the Cross-Section

The stylized facts shown in Section 3 are consistent with an implicit promotion driving household hoarding. Households purchased excessively when retail prices were low relative to wholesale prices and had amassed considerable inventories by the time prices rose. Of course, these patterns are not conclusive. It is possible that consumers were driven by an entirely different set of motives—panic or precaution—and the path of retail prices during and after the hoarding period was merely coincidental.

In this final section of our paper we develop a series of tests using the cross-section of excess purchases and later price changes for different rice products, brands, and stores to distinguish an implicit promotion motive caused by sticky prices from panic or precaution. The logic of our tests is that—under an implicit promotion motive—excess purchases should reflect latent expectations about coming price changes. Similarly, both panic and precaution have direct implications for the degree of hoarding relative to later price increases.

4.1 Dispersion in Price Discounts

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

²⁴In general, inventory levels are negatively correlated with purchases, in line with [Hendel & Nevo \(2006b\)](#).

this increase varied significantly. This cross-sectional dispersion is the result of a number of factors, 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, presenting histograms at various levels of aggregation. Given our focus on the implicit promotion motive, we present these price changes not in terms of percentage increases after the hoarding period, but rather as percentage discounts *during* the hoarding period. How far below the post-crisis level were prices during the hoarding period? Put differently, how big of an implicit discount was generated by sticky prices? 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

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

For example, if a price was \$8 during the hoarding period and \$10 in the post-hoarding period, this would translate to a realized discount of 20 percent.²⁵

In Panel A, we report a histogram of realized discounts for store-UPC pairs in our sample. The median discount is roughly 20% and nearly all store-UPCs saw a positive discount (i.e. a post-hoarding increase). There is considerable dispersion: the 25th percentile is close to 5 percent while the 75th percentile is close to 30 percent. The same conclusions apply to 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 is the basis of the forecast tests conducted later in this section.

4.2 Retrieving Latent Consumer Price Expectations

The basis of our forecast tests is a comparison of realized price changes against the expectations implied by consumer purchases under a rational, risk neutral, benchmark. To retrieve these expect-

²⁵While we focus on this discount metric throughout our tests to frame our results in terms of an implicit promotion, our findings do not qualitatively change if we consider price increases directly. The only difference is scaling the change in price by the hoarding-period price as opposed to the post-hoarding price.

tations, we first estimate consumer’s elasticities to standard retail promotions. We then use these elasticities to construct our benchmark from observed excess purchases during the crisis.

4.2.1 Consumer responsiveness to promotions in non-hoarding periods.

We begin by estimating consumer responsiveness to the discounts provided by typical retail sales in our sample, excluding the hoarding period. Our primary definition of a retail sales or promotions is constructed at the store-UPC level. Following the approach in [Hendel & Nevo \(2006b\)](#) 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).²⁶ 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 [A.VI](#) 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 household-brand-store, a store-UPC, a store-brand, a store, or a brand depending on the regression):

$$y_{it} = \delta \text{Sale}_{it} + \gamma_i + \eta_t + \varepsilon_{it}. \quad (2)$$

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

²⁶Specifically, we define a $\text{Sale} = \max\{\frac{\text{Modal Price} - \text{Price}}{\text{Modal Price}}, 0\}$.

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 4. 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 10 percent discount leads households to purchase roughly 0.17 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 sales 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 $\log(\text{ounces sold})$.²⁷ The coefficient in column 4 is 0.019, suggesting that a 10 percent discount generates a roughly 19 percent increase in quantity sold. The magnitude drops slightly, to 0.014, 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 4 show the specification in Equation 2 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

²⁷We use levels in our household regressions due to the large number of 0s and a substantially smaller degree of dispersion.

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.

4.2.2 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 A.I repeats the analysis in Table 4, but includes only data from 2007. We see similar, if marginally larger, estimates across all specifications.

Alternative sale definition: Another concern is that our specific definition of a sale might drive the estimated results. To address this, Appendix Table A.II repeats the analysis in Table 4, but defines a sale based upon the price relative to the modal price in a rolling 32 week window (the 16 weeks on either side of the week in question). Results are again similar to our main specification.

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, be 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 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.²⁸ This measure addresses concerns

²⁸The 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.

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.²⁹

Appendix Table A.III repeats the analysis in Table 4, 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).³⁰ 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 A.VII 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 by changes in relative prices across products (either contemporaneous 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 A.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 2 estimated at the household-brand-store (i.e. the specifications shown in Columns 1-3 of Table 4). However, as a dependent variable, these regressions include total ounces purchased by the household *at all other*

²⁹Note 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.

³⁰In 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.

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, and considering only temporary sales (two weeks or less).

4.2.3 Recovering cross-sectional expectations of price increases in the hoarding period

To recover our benchmark consumer expectations about the sticky-price generated implicit discount during the rice crisis for each product or store i we require two objects: (i) the percentage increase in quantities (relative to some baseline) for unit i during the hoarding period (Quantity Growth $_i$), and (ii) an estimate of $\hat{\delta}$ from Equation 2, which measures responsiveness to typical promotions. With this in hand, we define the forecasted implicit price discount as:

$$\text{Forecasted Discount}_i = \frac{\text{Quantity Growth}_i}{\hat{\delta}}.$$

The logic behind these forecasts is that, if excess purchases are driven by an implicit discount, scaling the growth in purchases by the sale elasticity estimated above should recover consumer forecasts of the discount during the crisis. This will be the case if (i) the sole driver of excess purchases during the hoarding period are expectations of future price increases and (ii) consumers respond to the implicit discount created by future price discounts just as they do the standard retail sales we consider when estimating $\hat{\delta}$.

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.³¹

³¹An 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

Specifically, if \bar{q}_i^h is average weekly quantity sold for unit i in the hoarding period, and \bar{q}_i^b is the average weekly quantity sold in the pre-hoarding period, we define: Quantity Growth $_i = \frac{\bar{q}_i^h - \bar{q}_i^b}{\bar{q}_i^b}$. We construct this, and the forecasted discount, at the store-UPC, store-brand, store, and brand levels. For each, we use $\hat{\delta}$ estimated at the corresponding level of aggregation.

4.3 Bias Tests

With estimates of both the realized and forecasted price discounts, we now turn to our two forms of forecast tests. We begin with a relatively standard bias test, which allow us to assess the predictions of panic or precautionary models. These tests consider the mean of cross-sectional forecast errors (i.e. the mean of the difference between the realized price discount and the forecasted price discount):

$$\text{Realized Price Discount}_i - \text{Forecasted Price Discount}_i = \zeta + u_i. \quad (3)$$

If consumers have rational expectations, are risk neutral, and excess purchases are motivated by the implicit discount, then we should expect the forecast error to be 0 on average (i.e. $\zeta = 0$). Both panic and precautionary models give a distinct prediction:

Panic: $\zeta < 0$. If hoarding is driven by behavioral overreaction, consumers must, on average, be purchasing more than would be predicted by a rational, risk neutral benchmark. If this is the case, our inferred forecasts will exceed the realized price discounts.

Precaution: $\zeta < 0$. If hoarding is driven by a hedging motive or risk aversion, consumers again must, on average, be purchasing more than would be predicted by a rational, risk neutral benchmark. If this is the case, our inferred forecasts will exceed the realized price discounts. This prediction holds if consumers are hedging against potential stockouts as well as price fluctuations.

In Table 5 we present the results of our bias tests at the store-UPC, store-brand, store, and brand level. Because computing the forecasted price discount requires estimating $\hat{\delta}$ before estimating the mean $\hat{\zeta}$, we bootstrap the procedure 1000 times to compute standard errors, drawing clusters at the unit level.

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.

Across specifications, we uniformly reject $\zeta < 0$. We instead find relatively small positive and significant estimates. For example, the mean forecast error across stores and UPCs is just under 10.7 and highly statistically significant. This suggests that our inferred consumer forecasts were roughly 11 percentage points lower than the actual realized price discount. The mean forecast error is slightly larger for at the brand level (11.3) and smaller but still positive and significant at the store-brand level (2.3).

Figure 5 presents histograms of these forecast errors. While there is unsurprisingly considerable dispersion, the mean and median of each distribution is positive, inconsistent with a panic or precautionary model. In other words, consumers were, if anything, under-reacting to the rice price shock (not buying enough relative to a rational risk-neutral benchmark), rather than overreacting.

4.3.1 Robustness for Bias Tests

We present robustness exercises in Panel A of Table 6 to address a series of potential concerns with our bias tests.

Downward biased forecasts due to stockouts: One potential concern is that some portion of the forecasts we recover are biased downward due to stockouts. If the supply of a particular product is exhausted, consumers may not be able to purchase their desired quantity at the going price. As a result, observed purchases may fail to reflect latent demand (and lead our procedure to underestimate consumer expectations).

There are at least two pieces of evidence indicating that stockouts are not meaningfully impacting the conclusions of our bias tests. The first is that the median forecast error—shown as a vertical line in each panel of 5—is positive at all levels of aggregation. Forecasts fell below realizations even at the median product or store. While there were media mentions of some stockouts during the crisis, no reports that we are aware of indicate shortages anywhere near the scale of 50 percent of store-UPC combinations. While stockouts may limit our ability to recover expectations in the tail of the distribution of excess purchases, consumer demand was likely fulfilled at the median.

The second piece of evidence is that our bias tests are virtually identical when considering only the products that we expect are least likely to face stockouts. Specifically, the first two columns of Panel A of Table 6 recreate our store-UPC level analysis, splitting the sample into products that

typically see a high degree of variation in sales versus those that typically see a low degree of variation in sales. We define high sales variation as store-UPCs with an above median coefficient of variation of weekly sales in 2007 (where the coefficient of variation is defined as the ratio of the standard deviation of weekly sales to the mean). The logic is that we expect products that typically see a high variability of sales (in percentage terms) were better able to withstand the increase in purchases that came during the hoarding period without experiencing stockouts. For both high and low variability products, we strongly reject the hypothesis that $\zeta < 0$. In both cases, point estimates are positive and similar in magnitude to the estimates in the full sample.

Shifts across products or stores: Another threat to our tests is the possibility that consumers shifted excess purchases to larger stores (or more commonly bought products) in a way that might bias our estimated forecasts downward. Because excess purchases are computed in percentage terms, the degree of hoarding might appear less stark at larger stores and, when averaging across stores, we might understate the total degree of hoarding. To address this, the third column of Table 6 repeats our store-UPC level analysis, but weights the regression by the level of sales in 2007. We again strongly reject the hypothesis that $\zeta < 0$.

Panic in specific communities: A final concern is the possibility that panic or precautionary motives were concentrated in particular communities for whom rice purchases were particularly salient. We consider two dimensions of heterogeneity in the last columns of Table 6. In columns 4 and 5, we split store-UPCs at the median according to Asian population (in percentage terms) at the county level. This split is motivated by contemporary news reports regarding the intensity of the hoarding episode in Asian communities.³² In the columns 6 and 7 we split store-UPCs at the median in terms of county level rice purchases (as observed in our sample in 2007) per capita. Across all specifications our estimates are similar to those in the full sample

4.4 Efficiency Tests

While our bias tests provide evidence against panic or precautionary models of hoarding in the rice crisis, they do not provide direct evidence in favor of a sticky price/implicit promotion view of hoarding. However, the implicit promotion view does make an explicit prediction about the cross-

³²See "A Run on Rice in Asian Communities," New York Times May 1, 2008.

sectional relationship between excess purchases and later price dynamics. Specifically, households should increase purchases the most for the products, brands, and stores in which they expect to see the largest future price increases. As long as consumers have some information about the coming cross-section of price changes—which might come from 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.

4.4.1 Hoarding Concentrated in Products With Large Post-Crisis Price Increases

We begin by showing that hoarding was indeed 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 6 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 were nearly double the average level among store-UPCs in the top quartile of post-hoarding price increases. Alternatively, there was very little excess purchasing for store-UPCs in the bottom quartile of post-hoarding price increases. This is consistent with consumer hoarding being driven by expectations about coming price increases, in line with an implicit promotion motive.

4.4.2 Formal Efficiency Tests

Our efficiency tests provide a 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, which echoes tests following [Mincer & Zarnowitz \(1969\)](#), by regressing the realized price discount (the discount of the hoarding period price relative to the long run price) on the forecasted price discounts we retrieved in Subsection 4.2:

$$\text{Realized Discount}_i = \alpha + \beta \text{Forecasted Discount}_i + u_i. \quad (4)$$

Because an estimate of $\hat{\delta}$ is necessary to compute $\text{Forecasted Discount}_i$, we once again bootstrap our standard errors, with clusters drawn at the unit level over 1000 repetitions. The key impli-

cation of a sticky price driven implicit promotion motive is that the degree of excess-purchases (captured by the forecasted discount) should predict actual price increases (captured by the realized discount):

Sticky-price driven implicit promotion: $\beta > 0$. If consumers are forward looking, have information regarding coming price changes, and are motivated by the implicit promotion caused by sticky prices, then forecasts should predict realized discounts.

We present results of our efficiency tests in Table 7 and Figure 7. This table and figure presents estimates of equation 4 at the store-UPC, store-brand, store, and brand levels. Column 1 of Table 7 shows a significant estimate of $\hat{\beta}$ of just above 0.1 at the store-UPC level. In the remaining columns, we see similar (if slightly larger) estimates at the store-brand and store level, and marginally smaller estimates (and slightly noisier, given substantially fewer observations) at the brand level. While a frictionless rational expectations model might predict even larger estimates (close to one), such a benchmark is likely unrealistic given measurement error in recovering Forecasted Discount_{*i*}, costly information acquisition for consumers, and other potential biases in beliefs. Overall, these consistently positive and significant coefficients provide support for the implicit promotion motive: excess purchases are robustly correlated with later price increases.

Figure 7 presents binned scatter plots of the same relationships, showing perhaps more striking evidence of this correlation. The relationship between forecasted discounts and realized discounts is roughly monotonic and close to linear for store-UPCs, store-brands and stores. At each point in the distribution, consumers purchased more of the products that later experienced greater price increases.

We interpret this pattern as evidence of the implicit promotion motive at work. Because prices were sticky, consumers had information during the hoarding period about price increases that were yet to come. 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 implicit promotion motive.

4.4.3 Robustness for Efficiency Tests

We begin by showing that the results in Table 7 are robust to various levels of fixed effects. To avoid the need for computationally intensive bootstrapping in these robustness exercises, we consider a slightly simplified version of our main specification. Specifically, for product i (which might represent a store-UPC, a brand-UPC, a brand, or a store) we consider cross-sectional regressions of the following form:

$$\text{Post-Hoarding Price Growth}_i = \gamma + \theta \text{Quantity Growth}_i + \varepsilon_i. \quad (5)$$

Post-Hoarding Price Growth $_i$ is a measure of the post-hoarding price increase, defined as the percentage increase in the post-hoarding period relative to the hoarding period itself. 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.³³ While the coefficients are not directly comparable to our formal efficiency tests above, the implicit promotion motive gives an analogous prediction: we should observe a positive relationship between hoarding and subsequent price increases, i.e. $\theta > 0$.

Table 8 presents the results of these regressions, which indicate 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 5 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 both store and UPC fixed effects. This suggests that the cross-sectional relationship between hoarding and price increases is not simply an artifact of hoarding at 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

³³We 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.

growth during the hoarding period, and post-hoarding price growth. The takeaway is straightforward and mirrors our efficiency tests. Consumers hoarded more with respect to the individual products, brands, and stores that later saw the greatest price increase.

Reverse causality

One concern is that the correlations shown in Table 7 and Table 8 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, Appendix Table A.V presents an alternative version of Table 8 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. However, Appendix Figure A.VIII shows that this does not appear to be the case. The figure displays the a binned scatter plot of the post-versus-pre increase in quantity sold at the store level against the change in the hoarding period relative to the pre-hoarding period.

This figure displays, if anything, a negative relationship between hoarding intensity and future demand.³⁴

Low prices during the crisis

A third possibility 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. Appendix Table A.VI shows that this is not the case. This table repeats the analysis in Table 8, but explicitly includes a control for the average price level 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.

These robustness exercises suggest that consumer anticipation was responsible for the cross-sectional correlation between excess purchases and later price increases—in other words, the implicit promotion motive drove consumer hoarding. On the whole, our forecast tests suggest that panic and precautionary models do not accurately describe cross-sectional patterns of consumer hoarding during the 2008 global rice crisis, and that sticky prices were responsible for a meaningful portion of excess purchases.

5 Conclusion

We point to the neglected role of sticky store prices as a trigger for household hoarding using US supermarket scanner data covering the 2008 global rice crisis. Even absent precautionary or panic motives, sticky prices generate an incentive for households to hoard. Recognizing that prices are temporarily low, households are driven to shift demand intertemporally and stockpile as if they were facing a standard retail promotion. We show that, following US news coverage of global supply disruptions and a shock to commodities prices, US households sharply increased rice purchases. Crucially, these purchases came in advance of any meaningful change in *shelf* prices. Consumers built inventories at relatively low cost before prices rose.

³⁴The figure as shown presents the data winsorized at the 10 percent level. A similar negative slope is present without winsorization (or with 1 percent winsorization), but the binned scatter plot is then dominated by extreme cases.

The cross-sectional relationship between excess purchases and price changes at the product (or store) level suggests that household hoarding was driven by the implicit promotion generated by sticky prices, and not by panic or precaution. To show this, we recover the price expectations implied by consumer purchases during the crisis, and compare these to realized price changes post-crisis. While both panic and precautionary models of hoarding have consumers over-reacting (relative to a rational, risk-neutral benchmark) we find that, if anything, consumers under-reacted to realized price dynamics on average across products and stores.

Furthermore, we find strong evidence of a correlation between excess purchases and later price increases in the cross-section. Households purchased more of the products that subsequently saw the largest price increases, consistent with demand shifting intertemporally in response to sticky prices. Taken together, our evidence indicates that sticky prices were a key driver of household hoarding during the 2008 global rice crisis. This suggests that retailer pricing policies play an important role in hoarding episodes, and that anti-price gouging regulations and related policies may actually amplify consumer stockpiling.

References


- ANDERSON, ERIC, MALIN, BENJAMIN A, NAKAMURA, EMI, SIMESTER, DUNCAN, & STEINSSON, JÓN. 2017. Informational rigidities and the stickiness of temporary sales. *Journal of Monetary Economics*, **90**, 64–83.
- ANDERSON, ERIC T, & SIMESTER, DUNCAN I. 2010. Price stickiness and customer antagonism. *The quarterly journal of economics*, **125**(2), 729–765.
- 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.
- 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.
- CHOI, HYUNYOUNG, & VARIAN, HAL. 2012. Predicting the present with Google Trends. *Economic Record*, **88**, 2–9.
- DA, ZHI, ENGELBERG, JOSEPH, & GAO, PENGJIE. 2011. In search of attention. *The Journal of Finance*, **66**(5), 1461–1499.
- 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.
- ENGELBERG, JOSEPH E, & PARSONS, CHRISTOPHER A. 2011. The causal impact of media in financial markets. *The Journal of Finance*, **66**(1), 67–97.
- ERDEM, TÜLIN, IMAI, SUSUMU, & KEANE, MICHAEL P. 2003. Brand and quantity choice dynamics under price uncertainty. *Quantitative Marketing and economics*, **1**(1), 5–64.
- FANG, LILY, & PERESS, JOEL. 2009. Media coverage and the cross-section of stock returns. *The Journal of Finance*, **64**(5), 2023–2052.
- 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.
- GOEL, SHARAD, HOFMAN, JAKE M, LAHAIE, SÉBASTIEN, PENNOCK, DAVID M, & WATTS, DUNCAN J. 2010. Predicting consumer behavior with Web search. *Proceedings of the National academy of sciences*, **107**(41), 17486–17490.
- GÖNÜL, FÜSUN, & SRINIVASAN, KANNAN. 1996. Estimating the impact of consumer expectations of coupons on purchase behavior: A dynamic structural model. *Marketing science*, **15**(3), 262–279.
- 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. 2006a. Measuring the Implications of Sales and Consumer Inventory Behavior. *Econometrica*, **74**(6), 1637–1673.
- HENDEL, IGAL, & NEVO, AVIV. 2006b. Sales and consumer inventory. *The RAND Journal of Economics*, **37**(3), 543–561.

- LEIBTAG, EPHRAIM. 2007. *Cost pass-through in the US coffee industry*. USDA Economic Research Serv.
- MANKIW, N GREGORY, & REIS, RICARDO. 2002. Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve. *The Quarterly Journal of Economics*, **117**(4), 1295–1328.
- 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.
- NAKAMURA, EMI, & ZEROM, DAWIT. 2010. Accounting for incomplete pass-through. *The Review of Economic Studies*, **77**(3), 1192–1230.
- 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*.
- Ó GRÁDA, CORMAC. 2008. The ripple that drowns? Twentieth-century famines in China and India as economic history. *Economic History Review*, **61**(s1), 5–37.
- PRIEST, TYLER. 2012. The Dilemmas of Oil Empire. *Journal of American History*, **99**(1), 236–251.
- ROTEMBERG, JULIO J. 2005. Customer anger at price increases, changes in the frequency of price adjustment and monetary policy. *Journal of Monetary Economics*, **52**(4), 829–852.
- SCHEINKMAN, JOSE A, & SCHECHTMAN, JACK. 1983. A simple competitive model with production and storage. *The Review of Economic Studies*, **50**(3), 427–441.


- SHESHINSKI, EYTAN, & WEISS, YORAM. 1977. Inflation and costs of price adjustment. *The Review of Economic Studies*, **44**(2), 287–303.
- SLAYTON, TOM. 2009. *Rice crisis forensics: How Asian governments carelessly set the world rice market on fire*. Tech. rept.
- SUN, BAOHONG. 2005. Promotion effect on endogenous consumption. *Marketing science*, **24**(3), 430–443.
- SUN, BAOHONG, NESLIN, SCOTT A, & SRINIVASAN, KANNAN. 2003. Measuring the impact of promotions on brand switching when consumers are forward looking. *Journal of marketing research*, **40**(4), 389–405.
- WOODFORD, MICHAEL. 2009. Information-constrained state-dependent pricing. *Journal of Monetary Economics*, **56**, S100–S124.
- 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 IMF's 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



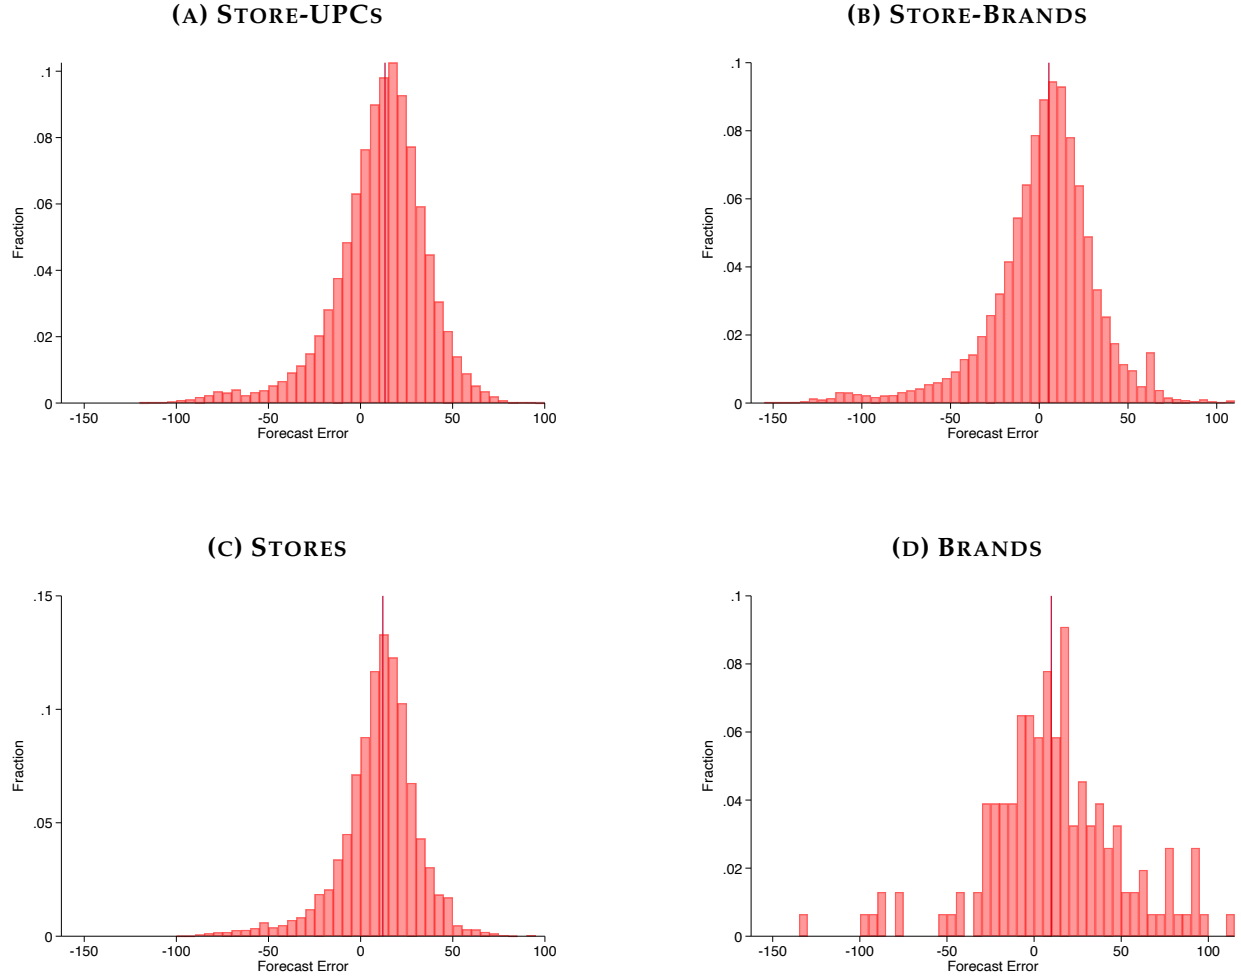
Notes: The black line displays a proxy for the US wholesale rice price at the monthly level. The proxy is based on average f.o.b. price for long grain rice at selected milling centers in Southwest Louisiana. Data provided by USDA, based on data from Agricultural Marketing Service, *National Weekly Rice 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 implicit price discounts generated by sticky prices, trimmed at the 1 percent level. The realized price discount is a measure of the post hoarding price increase. 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 $100 \times \left(\frac{\bar{p}^a - \bar{p}^h}{\bar{p}^a} \right)$. 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: DISTRIBUTIONS OF FORECAST ERRORS



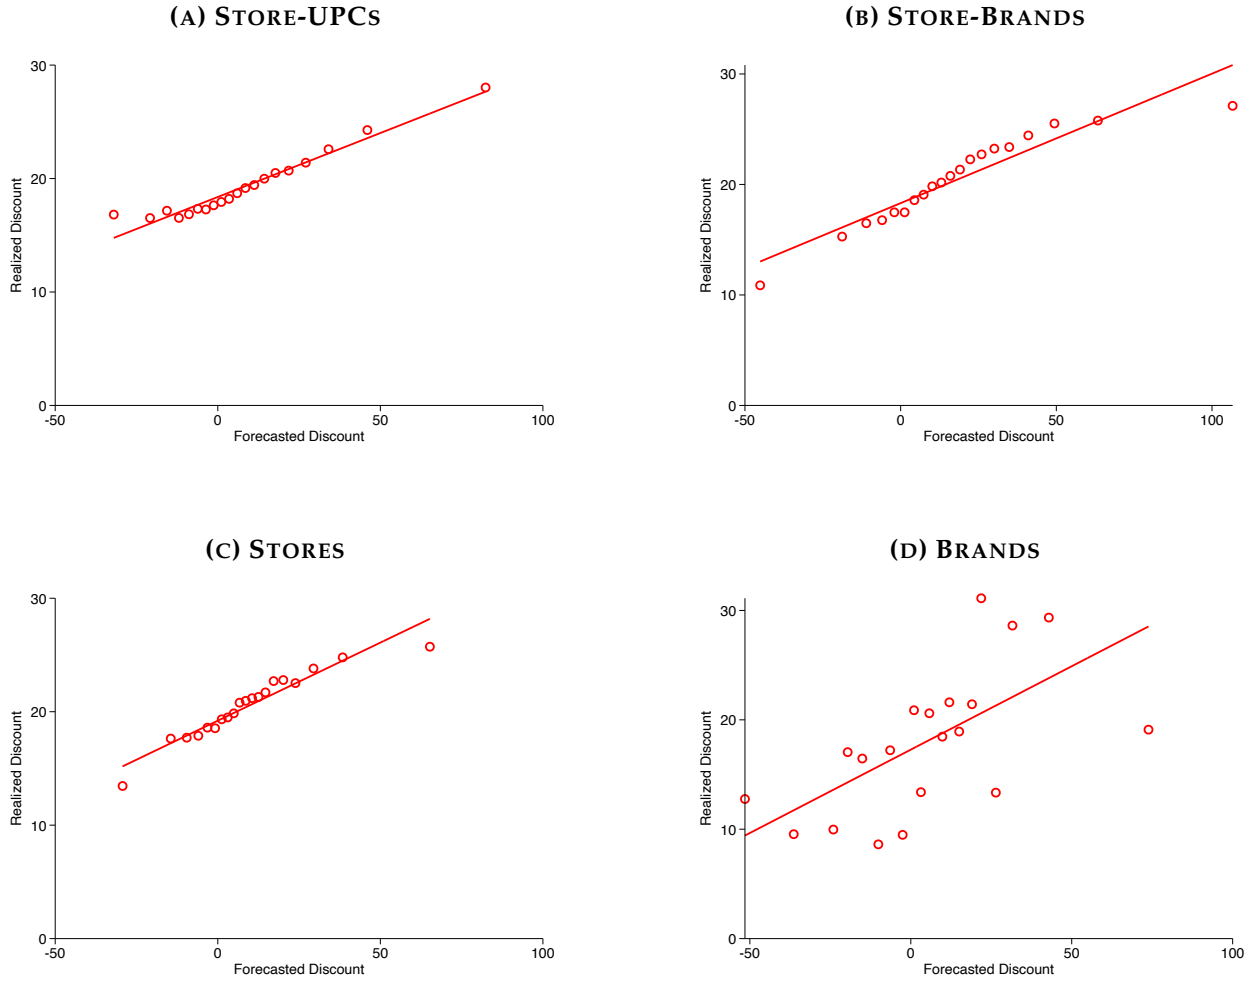
Notes: Histogram of forecast errors, defined as the difference between realized and predicted price discounts. The realized price discount is a measure of the post-hoarding price increase. Letting \bar{p}^h denote the average unit price in the hoarding period and \bar{p}^a denote the average unit price after the hoarding period, it is defined as $100 \times \left(\frac{\bar{p}^a - \bar{p}^h}{\bar{p}^a} \right)$. The predicted price discount 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 $\hat{\delta}$ is an estimated sale elasticity at the corresponding level of observation from Table 4, we define the forecasted discount to be $\left(\frac{\bar{q}^b - \bar{q}^h}{\bar{q}^b} \right) / \hat{\delta}$. 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. 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.

FIGURE 6: HETEROGENEITY IN UPC-LEVEL PRICE CHANGES CORRELATED WITH HOARDING



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 7: CROSS-SECTIONAL DIFFERENCES IN HOARDING PREDICT POST-HOARDING PRICE INCREASES



Notes: Binned scatterplots of realized implicit price discount during the hoarding period on forecasted price discounts. The realized price discount is a measure of the post-hoarding price increase. 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 $100 \times \left(\frac{\bar{p}^a - \bar{p}^h}{\bar{p}^a} \right)$. The predicted price discount 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 $\hat{\delta}$ is an estimated sale elasticity at the corresponding level of observation from Table 4, we define the forecasted discount to be $\left(\frac{\bar{q}^b - \bar{q}^h}{\bar{q}^b} \right) / \hat{\delta}$. 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. 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.

TABLE 1: SUMMARY STATISTICS

Panel A: Stores				
	Mean	S.D.	1st Percentile	99th Percentile
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 Purchase	71.9	130.5	12	480
Expenditure 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*** (112.4)	13.78*** (4.960)	8.443*** (0.689)	-0.738** (0.366)	3397.5*** (381.0)	-0.662 (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 IMF's 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: RICE PURCHASES DURING THE CRISIS BY INVENTORY LEVEL

	Household Level				Household-Store-Brand Level			
	Any Rice Purchase		Quantity Purchased (oz)		Any Rice Purchase		Quantity Purchased (oz)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Hoarding Period	0.013*** (0.000)	0.081*** (0.002)	1.999*** (0.098)	12.678*** (0.792)	0.005*** (0.000)	0.035*** (0.002)	0.482*** (0.034)	2.781*** (0.206)
Hoarding Period \times Inventory		-0.119*** (0.003)		-17.821*** (1.125)		-0.055*** (0.002)		-4.110*** (0.290)
2007 Mean	0.026	0.026	1.78	1.78	0.012	0.012	0.60	0.60
Observations	6663237	3702688	6663237	3702688	7725028	3916836	7725028	3916836
HH FE	Yes	Yes	Yes	Yes	No	No	No	No
HH-Store-Brand FE	No	No	No	No	Yes	Yes	Yes	Yes

Regressions of purchase behavior at the household or household-store-brand level on an indicator for the hoarding period (the weeks of April 19th-May 10th) and the interaction between this indicator and the pre-hoarding period household inventory level. To calculate inventory at the household-week and household-store-brand-week levels, we first estimate weekly consumption as average weekly rice purchases in 2007. We then calculate the level of inventory as the cumulative difference between weekly purchases and consumption up to the week before the hoarding period (The week of April 12, 2008). Regressions without inventory use a weekly balanced panel from 2007-2009 of all observed store-brand-household combinations with at least one purchase of rice (or one purchase of the brand in question). Regressions with inventory further restrict to the sample with observed purchases in 2007. $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 4: CONSUMER RESPONSES TO RETAIL SALES

	Quantity Purchased (oz)			Log(Ounces Sold)					
	HH-Brand		(3)	Store-UPC		(6)	Store-Brand	Store	Brand
	(1)	(2)		(4)	(5)		(7)	(8)	(9)
Sale	0.017*** (0.003)	0.016*** (0.003)		0.019*** (0.001)	0.014*** (0.001)		0.016*** (0.002)	0.015*** (0.000)	0.013* (0.007)
Sale: 0-10 Percent			0.098*** (0.020)				0.171*** (0.008)		
Sale: 10-20 Percent			0.260*** (0.042)				0.291*** (0.018)		
Sale: 20-30 Percent			0.324*** (0.075)				0.461*** (0.022)		
Sale: 30+ Percent			0.493*** (0.172)				0.749*** (0.052)		
Unit Price		-0.022** (0.010)			-0.117*** (0.023)				
Mean of Dep. Var.	0.57	0.57	0.57	5.77	5.77	5.77	6.61	7.43	9.82
Observations	4466520	4466520	4466520	10465671	10465671	10465671	6614642	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 percentage sales. A sale is defined as the percentage discount relative to the modal price in corresponding half year period (January-June or July-December) for a given UPC in a given store. Sales 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 5: CROSS-SECTIONAL BIAS TEST

	Store-UPC	Store-Brand	Store	Brand
	(1)	(2)	(3)	(4)
Mean Forecast Error	10.689*** (0.084)	2.329*** (0.192)	9.914*** (0.261)	11.286 (378.701)
Observations	69439	43953	9252	154

Forecast errors are defined as the difference between realized and forecasted price discounts. The realized price discount is a measure of the post-hoarding price increase. 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 $100 \times \left(\frac{\bar{p}^a - \bar{p}^h}{\bar{p}^a} \right)$. The forecasted price discount 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 δ is an estimated sale elasticity at the corresponding level of observation from Table 4, we define the forecasted discount to be $\left(\frac{\bar{q}^h - \bar{q}^b}{\bar{q}^b} \right) / \delta$. 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. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 6: ROBUSTNESS FOR CROSS-SECTIONAL FORECAST TESTS

Panel A: Bias Tests							
	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)
Mean Forecast Error	9.807*** (0.134)	11.376*** (0.130)	5.766*** (0.388)	11.184*** (0.133)	10.996*** (0.125)	10.156*** (0.128)	11.987*** (0.127)
Observations	30907	30904	69439	31419	30381	32263	29537

Panel B: Efficiency Tests							
	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)
Predicted Price Discount	0.091*** (0.002)	0.131*** (0.002)	0.091*** (0.005)	0.100*** (0.002)	0.134*** (0.002)	0.109*** (0.002)	0.125*** (0.002)
Observations	30907	30904	69439	31419	30381	32263	29537

Forecast errors in Panel A are defined as the difference between realized and predicted price discounts. Panel B presents regressions of realized price discount during the hoarding period on the predicted price discount. All regressions are at the store-upc level. The realized price discount is a measure of the post-hoarding price increase. Letting \bar{p}^h denote the average unit price in the hoarding period and \bar{p}^s denote the average unit price after the hoarding period, it is defined as $100 \times \left(\frac{\bar{p}^s - \bar{p}^h}{\bar{p}^h} \right)$. The predicted price discount 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}^s is average quantity sold in the pre-period, and $\hat{\delta}$ is an estimated sale elasticity at the corresponding level of observation from Table 4, we define the forecasted discount to be $\left(\frac{\bar{q}^s - \bar{q}^h}{\bar{q}^h} \right) / \hat{\delta}$. 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. 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. Low vs. high sales variance are store-UPC pairs 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 7: CROSS-SECTIONAL EFFICIENCY TESTS

	Dependent Variable: Realized Price Discount During Hoarding Period			
	Store-UPC	Store-Brand	Store	Brand
	(1)	(2)	(3)	(4)
Predicted Price Discount	0.113*** (0.001)	0.117*** (0.002)	0.125*** (0.004)	0.101* (0.061)
Observations	69439	43953	9252	154

Regressions of realized price discount during the hoarding period on the predicted price discount. The realized price discount is a measure of the post-hoarding price increase. 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 $100 \times \left(\frac{\bar{p}^a - \bar{p}^h}{\bar{p}^a} \right)$. The predicted price discount is derived from observed quantity growth during the hoarding period. Specifically: if \hat{q}^h is the average quantity sold during the hoarding period, \hat{q}^a is average quantity sold in the pre-period, and $\hat{\delta}$ is an estimated sale elasticity at the corresponding level of observation from Table 4, we define the forecasted discount to be $\left(\frac{\hat{q}^h - \hat{q}^a}{\hat{q}^a} \right) / \hat{\delta}$. 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. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.


TABLE 8: 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*** (0.001)	0.113*** (0.001)	0.085*** (0.001)	0.112*** (0.002)	0.137*** (0.002)	0.090*** (0.002)	0.142*** (0.004)	0.111*** (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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

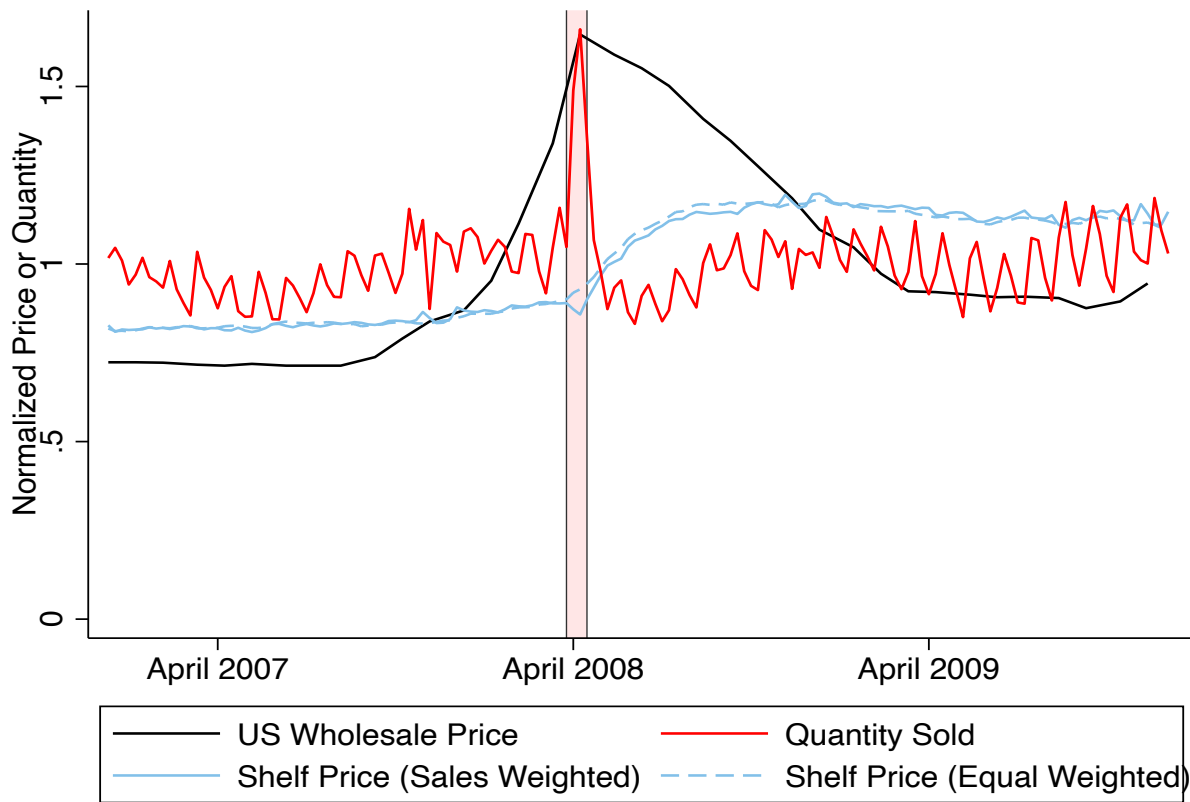
Internet Appendix: For Online Publication

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 IMF's 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 \(2006b\)](#). 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




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 IMF's 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: EXAMPLE OF RETAIL SALES



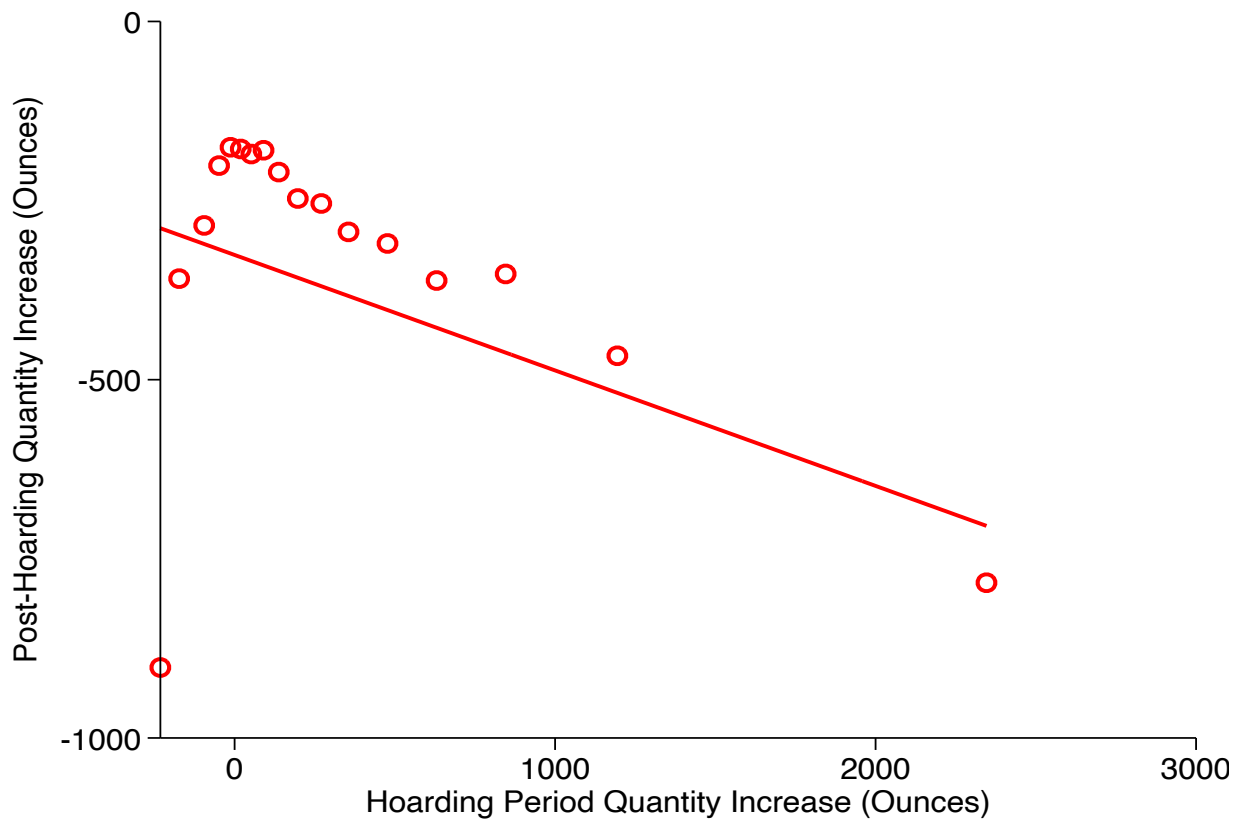
Notes: Example price path for a UPC in our sample in the pre-hoarding period (January 1st 2007-April 19th, 2008). Red portions indicate sales as identified by our algorithm.

FIGURE A.VII: 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.

FIGURE A.VIII: HOARDING INTENSITY DOES NOT SIGNAL FUTURE DEMAND



Notes: Binned scatterplot of pre-versus-post hoarding difference in average weekly sales (post-hoarding period quantity increase) at the store level in ounces against the difference between the hoarding period and the pre-hoarding period (hoarding period quantity increase). Each point represents a ventile of the hoarding period quantity increase distribution. Solid line shows the OLS fit through the raw data. Both variables are winsorized at the 10 percent level.

TABLE A.I: 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	0.020*** (0.003)	0.018*** (0.003)		0.021*** (0.001)	0.016*** (0.002)		0.017*** (0.002)	0.018*** (0.000)	0.016*** (0.006)
Sale: 0-10 Percent			0.098*** (0.029)			0.167*** (0.011)			
Sale: 10-20 Percent			0.215*** (0.025)			0.336*** (0.020)			
Sale: 20-30 Percent			0.415*** (0.116)			0.519*** (0.027)			
Sale: 30+ Percent			0.931*** (0.291)			0.750*** (0.046)			
Unit Price		-0.024** (0.012)			-0.105*** (0.040)				
Mean of Dep. Var.	0.59	0.59	0.59	5.81	5.81	5.81	6.63	7.51	9.92
Observations	1528020	1528020	1528020	3638961	3638961	3638961	2247245	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 the percentage discount relative to the rolling modal price in a 32 week window centered on the week in question. Sales 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 A.II: CONSUMER RESPONSES TO RETAIL SALES (ROLLING MODAL PRICE)

	Quantity Purchased (oz)			Log(Ounces Sold)					
	HH-Brand		(3)	Store-UPC		Store-Brand	Store	Brand	
	(1)	(2)		(4)	(5)	(6)	(7)	(8)	(9)
Sale	0.020*** (0.003)	0.019*** (0.003)		0.020*** (0.001)	0.014*** (0.001)		0.017*** (0.002)	0.015*** (0.000)	0.018*** (0.006)
Sale: 0-10 Percent			0.100*** (0.021)			0.191*** (0.008)			
Sale: 10-20 Percent			0.286*** (0.048)			0.301*** (0.015)			
Sale: 20-30 Percent			0.438*** (0.091)			0.473*** (0.024)			
Sale: 30+ Percent			0.682*** (0.200)			0.744*** (0.053)			
Unit Price		-0.022** (0.010)			-0.120*** (0.023)				
Mean of Dep. Var.	0.57	0.57	0.57	5.77	5.77	5.77	6.61	7.43	9.82
Observations	4466520	4466520	4466520	10465671	10465671	10465671	6614642	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 percentage sales. A sale is defined as the percentage discount relative to the modal price in corresponding half year period (January-June or July-December) for a given UPC in a given store. Sales 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 A.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	0.019*** (0.004)	0.014*** (0.004)		0.023*** (0.001)	0.017*** (0.002)		0.022*** (0.002)	0.012*** (0.000)
Sale: 0-10 Percent			0.080*** (0.020)			0.070*** (0.007)		
Sale: 10-20 Percent			0.293*** (0.048)			0.331*** (0.018)		
Sale: 20-30 Percent			0.379*** (0.084)			0.538*** (0.029)		
Sale: 30+ Percent			0.316 (0.207)			0.868*** (0.072)		
Unit Price		-0.070*** (0.015)			-0.121*** (0.023)			
Mean of Dep. Var.	0.61	0.61	0.61	5.77	5.77	5.77	6.61	7.43
Observations	3552848	3552848	3552848	10417267	10417267	10417267	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 the percentage discount relative to the modal price in corresponding half year period (January-June or July-December) for a given UPC in a given store. Sales 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 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 A.IV: CROSS-SUBSTITUTION IN RESPONSE TO
RETAIL SALES**

	Quantity Purchased (oz)		
	HH-Brand		
	(1)	(2)	(3)
Sale	0.000 (0.003)	0.000 (0.004)	
Sale: 0-10 Percent			0.031 (0.033)
Sale: 10-20 Percent			0.035 (0.076)
Sale: 20-30 Percent			-0.142 (0.122)
Sale: 30+ Percent			-0.083 (0.199)
Unit Price		-0.000 (0.005)	
Mean of Dep. Var.	2.00	2.00	2.00
Observations	4466520	4466520	4466520
HH-Store-Brand FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes

Regressions of total weekly ounces purchased by a household excluding a given brand-store pair on percentage sales for that brand. A sale is defined as the percentage discount relative to the modal price in corresponding half year period (January-June or July-December) for a given UPC in a given store. Sales 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A.V: EXCESS HOARDING PURCHASES PREDICT LONG RUN SHELF PRICE GROWTH

	Dependent Variable: Long Run Post-Hoarding Price Growth							
	Store-UPCs			Store-Brands			Stores	Brands
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Quantity Growth (%)	0.113*** (0.002)	0.142*** (0.002)	0.100*** (0.002)	0.105*** (0.003)	0.140*** (0.003)	0.098*** (0.003)	0.108*** (0.008)	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

Cross-sectional regressions of long run post-hoarding shelf price growth on quantity growth (excess purchases) 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 A.VI: EXCESS HOARDING PURCHASES PREDICT SHELF PRICE GROWTH CONTROLLING FOR PRICE

	Dependent Variable: Post-Hoarding Price Growth							
	Store-UPCs			Store-Brands			Stores	Brands
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Quantity Growth (%)	0.103*** (0.001)	0.113*** (0.001)	0.059*** (0.001)	0.073*** (0.002)	0.096*** (0.002)	0.068*** (0.002)	0.090*** (0.004)	0.075** (0.033)
Price	-0.957*** (0.030)	-0.637*** (0.036)	-13.610*** (0.119)	-1.737*** (0.021)	-1.704*** (0.024)	-3.773*** (0.057)	-4.983*** (0.113)	-0.724*** (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

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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. We additionally control for the unit price-per-80 ounces. For levels of aggregation above store-UPC, these prices are sales weighted within weeks.