Sales and Price Spikes in Retail Scanner Data

Ryan A. Chahrour*

November 30, 2010

Abstract

Price filters used by macroeconomists place strong restrictions on "regular price" series. I propose a new price filter and show that, while many pricing facts are robust to filter specification, implications for price duration depend on the choice of filter.

JEL Codes: E31, L10

Keywords: Scanner Data, Price Setting, Price Stickiness

Forthcoming in *Economics Letters*, DOI: http://dx.doi.org/10.1016/j.econlet.2010.11.008.

*Columbia University, Economics Department, International Affairs Building, 420 West 118th Street, New York NY, 10027. Contact at rc2374@columbia.edu or 774-212-3854. Website www.columbia.edu/∼rc2374.

1 Introduction

Since the groundbreaking paper by Bils and Klenow (2004), macroeconomists have sought to reconcile extremely frequent price changes in micro-level data with stickiness in aggregate price indexes. In addressing this disconnect, many authors seek to filter out temporary "sales" from micro-level price series. Since sale price changes are temporary, it is argued, they will have a much smaller impact on measures of aggregate price stickiness. Using this method on retail data, authors (e.g. Nakamura and Steinsson (2008)) find much higher levels of price stickiness than first reported by Bils and Klenow (2004).

More recently, Eichenbaum, Jaimovich and Rebelo (2010) (EJR) argue that prices are characterized by fluctuations both below and *above* a more stable underlying series (reference price), which they identify as the most frequently observed price in a given quarter. The reference price view of pricing is significant because 1) it suggests a different pattern of pricing behavior than the standard sale/non-sale price story and 2) it excludes an even larger portion of price changes than sale filters, yielding more persistent series and suggesting a stronger role for nominal rigidities.

In this letter, I compare the implications of different price filters using the popular Dominick's scanner data set. I find that many pricing "facts" do not depend crucially on the specification of the price filter. Among these facts: sales are large and frequent; regular price increases are more common than price decreases; prices spend a large majority of time at the regular level; regular prices vary significantly, but less than observed weekly prices.

Despite the general "robustness" result, however, I find that price spikes have important implications for estimates of regular price persistence. The EJR filter restricts regular prices to change only on certain dates, and therefore may identify spurious sales and price spikes. Sales filters, in contrast, constrain filtered series to have *no* price spikes. I propose a filter designed to relax these restrictions. Using this filter, I show that price spikes are much smaller and less frequent than suggested by EJR, but that, despite this, removing price spikes greatly increases estimates of price persistence.

2 Defining Price Filters

The price filters examined here all presume that price series can be decomposed into two components, so that $p_{i,t} = x_{i,t} + \epsilon_{i,t}$. Because only $p_{i,t}$ is observed, restrictions are needed to identify these series separately. It is typically assumed that $x_{i,t}$ is relatively persistent, with a high probability that $x_{i,t} = x_{i,t-1}$, while $\epsilon_{i,t}$ is relatively fast moving, with a high probability that $\epsilon_{i,t} = 0$. Price filters aim to extract the more persistent $x_{i,t}$ because authors believe that this series relates most closely to measures of aggregate price stickiness. I refer to $x_{i,t}$ as an "attractor price."

Visually, price filtering is described as removing price changes that are quickly reversed. However, authors do not agree on the specifics of this procedure. For example, what constitutes a price reverse? The AC Neilsen filter, used by Kehoe and Midrigan (2007), records a sale if a price decrease is followed by *any* price increase shortly thereafter. In contrast, Nakamura and Steinsson (2008) consider, among others, a sale filter that records a sale only when a price decrease is followed by a return to the price in effect just before the decrease. Since this requirement is more stringent, the latter procedure will typically identify fewer sales and more frequent price changes.

The concern I highlight, however, is the assumption made by all sale filters that an observed price is *never above* the attractor price. In the decomposition above, this is the restriction that $\epsilon_{i,t} \leq 0$. For this reason, I refer to sale filters as "one-sided."[1] The validity of this restriction depends on the pricing behavior of firms. However, without a strong theoretical argument to support it, it seems unnecessary: price increases that are indeed changes in the attractor price, $x_{i,t}$, should only rarely be excluded by a good two-sided filter. By imposing the restriction, we risk missing important features of the data.

In contrast, the simple EJR filter - selecting the most commonly observed price that quarter - relaxes the restriction that $\epsilon_{i,t} \leq 0$. In return, however, it restricts $x_{i,t}$ to change only at arbitrary quarter cutoff dates. This means that reference prices must last for whole-quarter periods of time. Even if reference price decisions are made at a quarterly frequency, it seems unlikely they will line up with any particular calendar definition of the quarters.

[1] The definitions of "one-sided" and "two-sided" used here are unrelated to the terms' standard meaning in timeseries filtering.

Figure 1, plotting actual and filtered prices for County Line Colby Jack Cheese, shows why this restriction is important. The second update in the EJR reference price occurs five weeks before the (visually obvious) adjustment in the attractor price. Over the period from the change in the reference price to the change in the actual attractor price, EJR will observe a price that is above the reference price. I argue that this type of misspecification explains a large portion of their observations above reference price. Figure 2, for Miller Lite Beer, shows how severely this problem can affect an individual series, even when the true attractor series is uncontroversial.

2.1 Proposed Filter

I propose a price filter to address these specification concerns.[2] This procedure determines regular prices based on a 13-week (t plus or minus 6 weeks) rolling window, rather than the fixed, non-overlapping quarter periods of EJR.[3] This allows for reference prices to last for fractional quarters as necessary. Under my specification, a filtered price must be the most commonly observed price *in some* thirteen-week period. This change goes a long way to addressing the concerns above.

The simple rolling window mode, however, does not always select the obvious price at transition points. At some transitions, the new and old modal prices appear in overlapping periods. My procedure selects the price that appears most often during the period of overlap. At other times, prices vary among non-attractor levels just at the point of transition, obscuring its actual date. My procedure places the transition to the new attractor price in the first week that the observed price is closer to the new attractor price.

The appendix provides a step-by-step description of the proposed filter. For the algorithmic details of the other price filters used in this paper, I refer readers to Nakamura and Steinsson (2008), Kehoe and Midrigan (2007) and EJR. For consistency across methods, I parameterize all filters so that "temporary" price movements must last fewer than 7 weeks.

[2]In calibrating their menu-cost model, Kehoe and Midrigan (2008) have adopted a two-sided filter largely in the spirit I propose. They do not address my central concern here, however, which is how the choice of filter affects price statistics and, most notably, measures of price persistence.

[3]Of course, the window is a parameter, which can be adjusted.

3 Results

To compare price filters, I use the Dominick's Finer Food dataset maintained by the James M. Kilts Center, Chicago Graduate School of Business. The data contain weekly scanner price observations for products sold in all of Dominick's locations over the nine-year period, 1989 - 1997. I include only products which have six or more consecutive quarters with at least twelve valid price observations (out of a possible thirteen.) For cases where there are more than one such sequence, I take the longest.

The data are separated into 29 product categories. To generate aggregate statistics, I calculate the statistic for each product, take the median within each product category, and then take the median across product categories. This aggregation procedure is roughly consistent with the median-median approach of EJR. Other authors use different aggregation procedures. Thus, although my results are *consistent* with others' findings, they generally are not strictly comparable.

Tables 1 and 2 show some of the basic statistics that are reported in other studies. Here, my results are in line with previous research. First, a large majority of prices are at their "attractor levels." Correspondingly, most expenditures occur at these prices. Second, for all filtered series, about two-thirds of price changes are increases. Third, filtered series are slightly less than two-thirds as variable as the observed price series. Given the important differences in specification, the similarity across filtering methods is striking.

Despite the similarities, some key differences emerge. First, while the EJR procedure finds that 21% of non-attractor prices are above the corresponding attractor price, my procedure yields a much smaller 12%. For the sales filters, this percentage is zero by definition. Note that prices spend the smallest portion of time at the EJR reference price series, while spending substantially more time at each of the other three filtered series. With approximately 30% fewer observations away from the filtered price, my procedure yields only about 40% ($.7 * \frac{12}{21} = .4$) as many spike periods as indicated by the EJR procedure.

Table 3 shows the average sizes of sales and price spikes. For all methods, the average price during a sale is around 13% less than the filtered attractor price. On the other hand, there is substantial heterogeneity across the methods in average price "spikes." By definition, the sale filters never capture spikes. For the EJR method, the average spike price is 6.4% above the reference price, while with my method this number is only 4.4%. With the EJR method, sales are significantly larger than price spikes. With my method the difference is even greater; sales are three times as large as spikes.

Table 4 shows the impact on the frequency of price change and implied attractor price duration (1/frequency of price change.) Both the Nakamura-Steinsson and the AC Neilsen/Kehoe-Midrigan filters give an estimated price-change frequency just under 5%, corresponding to price spells around 1.6 quarters. The EJR-based estimate of 2.3% is much smaller and implies average price spells of 3.3 quarters. My filter suggests that attractor prices are somewhat more flexible, but the estimated frequency of 3.0% is still much lower than the estimates based on the sale-only filters and the implied duration of 2.6 quarters is significantly longer. Overall, table 4 shows that two-sided filters have noticeably different implications for the frequency of attractor price changes.

Finally, I study the robustness of my results to the sampling frequency of the data. I construct pseudo-monthly data from the original weekly data by selecting observations at four-week intervals. I then run the same analysis, calibrating each filter so that sales last no longer than two months. Tables 5 and 6 show that the patterns in the key statistics are unchanged. Price spikes remain much smaller than sales, and they are smaller using my filter compared to the filter of EJR. Furthermore, the sales-only filters yield an implied price duration of between 1.9 and 2.2 quarters, compared to 2.8 quarters for my filter, with a longer duration of 3.8 quarters for the EJR filter.

4 Discussion

These results lead to my conclusion that temporary price spikes are far less common and significantly smaller than sales. One interpretation of these results is that they introduce a second "excess price flexibility" puzzle. The sale puzzle - temporary, large decreases in price - is well recognized by the authors cited in this paper. The second - temporary, small increases in price

- is novel. Among recent macroeconomic models, there are at least a few rationale for sales, including consumer heterogeneity in price elasticities (Guimaraes and Sheedy, 2010), good-specific habit in consumption (Nakamura and Steinsson, 2009), and firms that face lower menu costs for temporary price decreases (Kehoe and Midrigan, 2007). On the other hand, temporary and tiny price increases seem harder to justify, with Kehoe and Midrigan (2008) the only model (to my knowledge) that can match this feature of the data. To do so, however, their model requires consumer heterogeneity in price elasticities *and* a very particular configuration of demand shocks.

An alternative interpretation is that observed price spikes are not a fundamental feature of pricing data, and instead mostly stem from issues related to price reporting, typographical error, etc. Absent a compelling economic motivation for the spikes in the data, this is a sort of default view. If it is correct, macroeconomists might justifiably ignore spikes, including in their models of price-setting.

Regardless of which interpretation prevails, there are good reasons for macroeconomists to examine micro pricing data using a two-sided, rather than a sales-only filter. The arguments for excluding large and frequent sales (e.g. for model calibration) are even more compelling when applied to small, quickly reversed price spikes. As I show here, properly accounting for these spikes has a crucial impact on estimates of price rigidity.

Acknowledgemts

I thank Emi Nakamura, Stephanie Schmitt-Grohé, Jón Steinsson, Martín Uribe, and an anonymous referee for helpful comments.

References

Bils, Mark, and Perter J. Klenow, 2004, Some Evidence On the Importance of Sticky Prices, Journal of Political Economy 112, 947–985.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo, 2010, Reference Prices and Nominal Rigidities, American Economic Review. Forthcoming.

Guimaraes, Bernardo, and Kevin D. Sheedy, 2010, Sales and Monetary Policy, American Economic Review. Forthcoming.

Kehoe, Patrick J., and Virgiliu Midrigan, 2007, Sales and the Real Effects of Monetary Policy, *Working Paper 652*: Federal Reserve Bank of Minneapolis.

—— 2008, Temporary Price Changes and the Real Effects of Monetary Policy, *Research Department Staff Report*: Federal Reserve Bank of Minneapolis.

Nakamura, Emi, and Jón Steinsson, 2008, Five Facts About Prices: a Reevaluation of Menu Cost Models, Quarterly Journal of Economics 123(4), 1415 – 1464.

—— 2009, Price Setting in Forward-Looking Customer Markets, *manuscript* Columbia University.

Table 1: Basic Statistics

	Nakamura-Steinsson	AC Neilsen/KM	EJR Reference	Chahrour
Frac of Price Obs. at Attractor Price	0.84	0.83	0.74	0.82
Frac of Expenditure at Attractor Price	0.77	0.76	0.68	0.76
Frac of Non-Attractor Prices above Attractor Price	0.00	0.00	0.21	0.12
Frac of Non-Attractor Prices below Attractor Price	1.00	1.00	0.79	0.88
Frac of Attractor Price Changes which are Increases	0.65	0.67	0.67	0.62
Frac of Attractor Price Changes which are Decreases	0.35	0.33	0.33	0.38

Table 2: Price Volatility

	Unfiltered Price	Nakamura-Steinsson	AC Neilsen/KM	EJR Reference	Chahrour
σ(log-level)	0.085	0.052	0.050	0.050	0.052
σ(log-change)	0.074	0.014	0.013	0.011	0.013

Table 3: Sales and Price Spikes

	Nakamura-Steinsson	AC Neilsen/KM	EJR Reference	Chahrour
Average Sale (%)	-13.5	-12.9	-11.7	-13.5
Average Spike (%)	0.0	0.0	6.4	4.4

Table 4: Price Persistence

	Unfiltered Price	Nakamura-Steinsson	AC Neilsen/KM	EJR Reference	Chahrour
Weekly Prob of Price Change (%)	24.3	4.8	4.6	2.3	3.0
Implied Duration (Quarters)	0.32	1.62	1.67	3.31	2.56

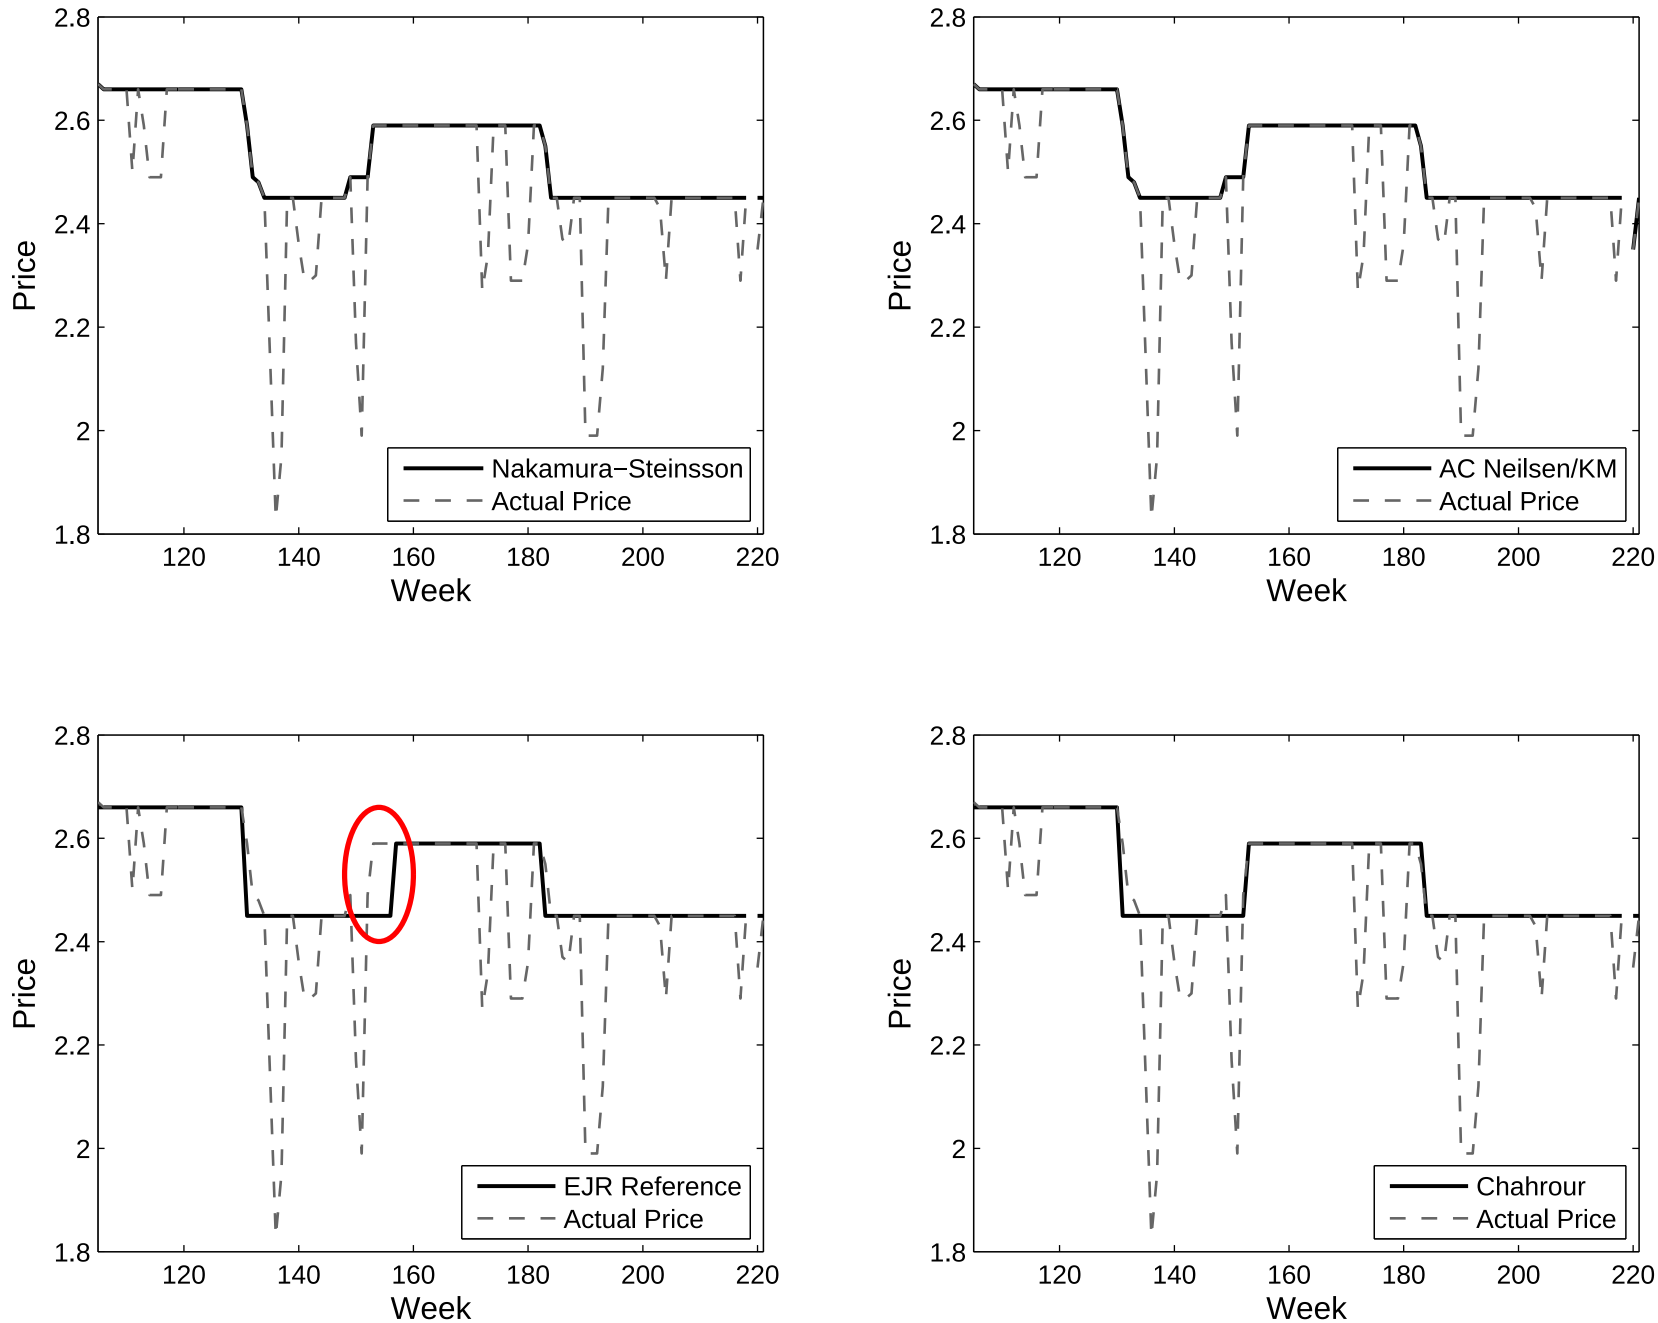

Figure 1: County Line Colby Jack Cheese

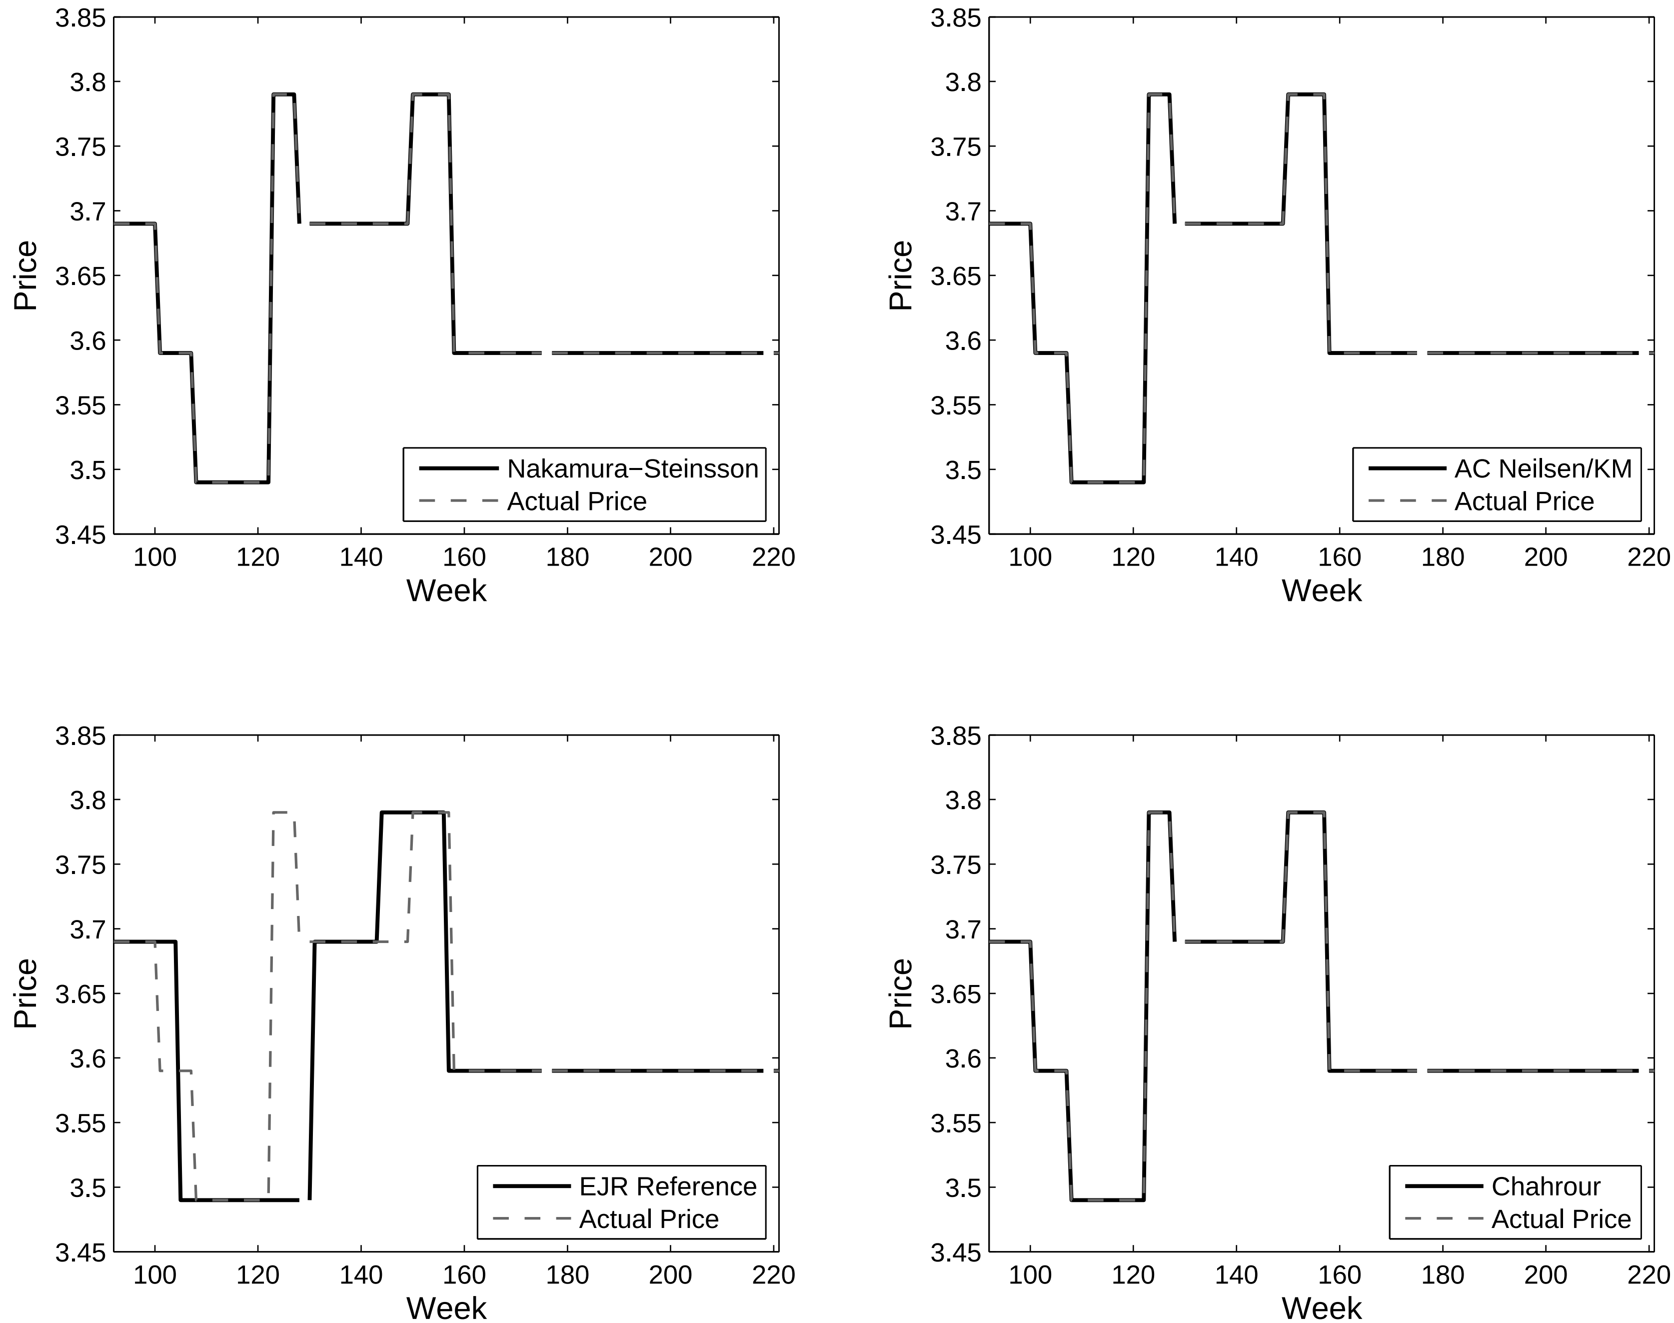

Figure 2: Miller Genuine Draft Beer

Table 5: Sales and Price Spikes (Pseudo-monthly Data)

	Nakamura-Steinsson	**AC Neilsen/KM**	**EJR Reference**	**Chahrour**
Average Sale (%)	-11.4	-12.5	-11.1	-10.9
Average Spike (%)	0.0	0.0	6.5	5.4

Table 6: Price Persistence (Pseudo-monthly Data)

	Unfiltered Price	**Nakamura-Steinsson**	**AC Neilsen/KM**	**EJR Reference**	**Chahrour**
Monthly Prob of Price Change (%)	40.7	15.9	13.8	8.2	10.9
Implied Duration (Quarters)	0.76	1.93	2.23	3.77	2.81

A Proposed Price Filter

My reference price filter loops over successive 2*K + 1 week windows, centered at time t. To describe the procedure, it is convenient to define the sets $\Omega_t = \{p_{t-K}, p_{t-K+1}, ..., p_t, ..., p_{t+K}\}$ and $\omega_t = \{t-K, t-K+1, ..., t, ..., t+K\}$, where p_t is the price *observed* at time t. These sets are truncated appropriately at sample beginning and end. Let $\hat{r}_t$ be a "candidate" reference price, and r_t be a "final" reference price. Then the procedure precedes as follows:

1. Suppose that we have already selected a sequence of candidate attractor prices, $\hat{r}_0, ..., \hat{r}_{t-1}$. Define $\hat{r}_t$ to be the most frequently occurring price in Ω_t. Ties (though rare) are broken by order of appearance.

2. **Case 1**: If $\hat{r}_t = \hat{r}_{t-1}$, index t by one period, and repeat.

   Case 2: If $\hat{r}_t \neq \hat{r}_{t-1}$, define $\tau^a \in \omega_t$ as the last period such that $p_{\tau^a} = \hat{r}_{t-1}$. Similarly, define $\tau^b \in \omega_t$ as the first period such that $p_{\tau^b} = \hat{r}_t$.

   - **Subcase A**: If $\tau^a \geq t$ and $\tau^b > t$, assign $\hat{r}_t = \hat{r}_{t-1}$. This is a case where the "naive" choice for $\hat{r}_t$ has jumped too soon; the old attractor price appears again at time t or later, and the new price does not first appear (within the window) until after t.

   - **Subcase B**: Otherwise, if $\tau^a < t$, $\tau^b < t$, and $\tau^a < \tau^b$, then let $\hat{r}_{\tau^a+1}, ...\hat{r}_{t-1} = \hat{r}_t$. The algorithm has jumped too late; the new reference price appears before time t, and after the last appearance of the old reference price.

   - **Subcase C**: Otherwise, if $\tau^a \geq \tau^b$, let $\hat{r}_t$ be the most frequently occurring price in the set $\{p_{\tau^a}, ..., p_{\tau^b}\}$. This is the case of "overlap"; the first occurrence of the new attractor candidate is prior to the last occurrence of the old attractor price.

   - **Subcase D**: Otherwise, if $\tau^a < \tau^b$, let $\tilde{\tau} \in \omega_t$ be the first period where $|p_{\tilde{\tau}} - \hat{r}_t| < |p_{\tilde{\tau}} - \hat{r}_{t-1}|$. Then let $\hat{r}_{\tau^a+1}...\hat{r}_{\tilde{\tau}-1} = \hat{r}_{t-1}$ and let $\hat{r}_{\tilde{\tau}}...\hat{r}_{\tau^b} = \hat{r}_t$. This is case where there is a "gap" between the last occurrence of the reference price and first occurrence of the new price.

   - **Subcase E**: Otherwise, leave $\hat{r}_t$ at its original value.

Note that, in principle, $\hat{r}_\tau$ is subject to change until the procedure has progressed far enough that τ has fallen out of the set ω_t. It is at this point that $\hat{r}_t$ becomes r_t.

Table A.1 summarizes how the algorithm handles different cases for relationship between τ_a, τ_b and t.

Table A.1

		Ta <t		Ta = t		Ta >t	
		Ta<Tb	Ta>Tb	Ta<Tb	Ta>Tb	Ta<Tb	Ta>Tb
Tb < t	Ta<Tb	**B**	-	-	-	-	-
	Ta>Tb	-	**C**	-	**C**	-	**C**
Tb = t	Ta<Tb	**E**	-	-	-	-	-
	Ta>Tb	-	-	-	-	-	**C**
Tb >t	Ta<Tb	**D**	-	**A**	-	**A**	-
	Ta>Tb	-	-	-	-	-	**A**

A: Too Soon
B: Too Late
C: Overlap
D: Gap
E: No Change
-: Impossible Case