

Individual Price Adjustment along the Extensive Margin*

Etienne Gagnon
Federal Reserve Board

David López-Salido
Federal Reserve Board

Nicolas Vincent
HEC - Montréal

April 2, 2012

Abstract

As the recent literature makes clear, firms employ a rich variety of price-setting strategies with often diverging implications for aggregate price dynamics. This heterogeneity poses a challenge for macroeconomists interested in bridging micro and macro price stickiness. In responding to this challenge, we follow Caballero and Engel (2007) by expressing the initial macro price response to shocks in terms of micro price adjustment along an intensive and an extensive margin. The intensive margin captures the response of price adjustments determined ahead of shocks and is closely related to the frequency of price changes. The extensive margin captures price adjustments that are triggered or cancelled by aggregate shocks. We use variation in the shape of the distribution of consumer price changes to show that adjustment along the extensive margin was key to the price level response to several macroeconomic shocks. We then use a variety of micro datasets to present evidence that items with large deviations from their optimum are more likely than others to adjust their price. Our evidence points to the extensive margin playing an economically important role in macro price adjustment.

*Prepared for the NBER's 27th Macro Annual Conference, April 20 – 21, 2012. This research was conducted with restricted access to Bureau of Labor Statistics (BLS) data. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System, any other person associated with the Federal Reserve System, or the BLS. We thank Ben Malin and Randal Verbrugge for their support with BLS micro data, as well as Mike Woodford, Jeff Campbell, and Ed Nelson for their insightful comments. Christine Garnier and Andrew Giffin provided superb research assistance. We are grateful to the SymphonyIRI Group for the scanner data. As a condition of use, SymphonyIRI reviews all papers using their data to check that the data are not described in a misleading fashion. However, all analyses in this paper based on SymphonyIRI Group, Inc. data reflect the work and conclusions of the authors, not SymphonyIRI Group, Inc. Comments and suggestions can be directed to etienne.gagnon@frb.gov, david.j.lopez-salido@frb.gov, nicolas.vincent@hec.ca.

1 Introduction

Over the past decade, economists have devoted substantial efforts to documenting basic facts about consumer micro price behavior, an endeavor made possible by the increased availability of large datasets of goods and services prices.¹ A salient finding is that firms employ a rich variety of pricing strategies. In some sectors, such as energy, air travel, or fresh produce, firms adjust prices frequently, whereas in others, such as newspapers, health services, or maintenance activities, firms adjust prices rather infrequently. Even at the level of universal product codes (UPC), researchers have found variation across firms in the magnitude, timing, and frequency of price adjustments.² And while sales and promotions are a defining feature of the way items are marketed to consumers in retail trade, temporary price discounts are uncommon in several other sectors.

Many macroeconomists had hoped that the new micro evidence would shorten the list of pricing mechanisms used in macroeconomic applications by revealing which mechanisms have empirical support and which ones do not. The avalanche and variety of new micro facts have instead stimulated researchers to introduce several new mechanisms and refine existing ones.³ With hindsight, this outcome was probably unavoidable. As has long been recognized in the field of industrial organization, product markets differ along several dimensions that influence firms' pricing strategies. These dimensions include the number of buyers and sellers, the degree of product homogeneity, the durability of items, the presence of long-term relationships between buyers and sellers, the substitutability or complementarity with other products, the role of advertisement and information, government regulation, the frequency of consumer purchases, and the ability to hold inventories.⁴ This heterogeneity would be unproblematic if aggregate shocks had similar effects across pricing mechanisms. Unfortunately, the choice of a particular mechanism is often consequential in macroeconomic applications. Notably, models matching the same average frequency and size of price changes can vary greatly in the speed of aggregate shock pass-through. (See Golosov and Lucas [2007] for an illustration.) The diversity of facts and pricing mechanisms thus underscores the importance of identifying key features of micro price behavior that macro models should aim to reproduce.

In this paper, we seek to connect micro and macro price stickiness by distinguishing between price changes that are *determined ahead* of shocks and those that are either *triggered or cancelled*

¹Initial empirical work using sectoral data by Carlton (1986), Cecchetti (1986), Lach and Tsiddon (1992), Kashyap (1995), and the survey-based analysis in Blinder (1994) have been followed by the seminal work on consumer prices of Bils and Klenow (2004), Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008) for the United States, Baharad and Eden (2004) for Israel, and several euro-area national studies summarized in Dhyne *et al.* (2005). For recent literature surveys, see Mackowiak and Smets (2008) and Klenow and Malin (2010).

²For an illustration, see Boivin *et al.* (2012), who report sharp differences in the frequency and size of individual price adjustments across U.S. on-line book retailers and across borders for retailers with international operations.

³The list of price-setting strategies and relevant frictions includes menu-cost models (e.g., Barro [1972] and Sheshinski and Weiss [1977]), the Calvo [1983] model, fixed-duration contracts (Taylor [1980]), infrequent information (e.g., Mankiw and Reis [2002]), rational inattention (e.g., Sims [2003]), uncertain and sequential trade (e.g., Eden [1994]), fair pricing (Rotemberg [2002]), price points (e.g., Levy *et al.* [2011]), search models (e.g., Head *et al.* [forthcoming]), price plans (Burstein [2006] and Alvarez *et al.* [2011]), reference prices (Eichenbaum, Jaimovich, and Rebelo [2011]), and best prices (Chevalier and Kashyap [2011]).

⁴This heterogeneity is the subject of an insightful survey paper by Carlton (1989).

by shocks. Our analysis builds on a decomposition, introduced by Caballero and Engel (2007, henceforth “CE”) in the context of generalized Ss models, that expresses the *initial* inflation response to an aggregate shock as the sum of an intensive margin and an extensive margin. The intensive margin captures the response of predetermined price adjustments and is connected to the observed frequency of price changes. The extensive margin captures the response of price adjustments triggered or cancelled by the shock and is at the heart of many debates on how item-level nominal price stickiness translates into aggregate price stickiness. The main goals of our paper are to explore how the extensive margin can be inferred from observed micro price behavior and then to assess its empirical importance. In doing so, we go beyond documenting unconditional moments of micro prices by focusing instead on pricing decisions in response to shocks.

The extensive margin depends on what we call the *individual reset price*; that is, the posted price that a firm would set if granted a *one-time* opportunity to adjust it for free, with all constraints otherwise remaining in place (including the possibility that the posted price will not change for some time). If individual reset prices were observed, then one could go a long way toward measuring the extensive margin by studying how deviations of posted prices from individual reset prices are distributed and how these deviations influence price adjustment decisions. A central insight of our work is that partial identification is possible by noting that, in a broad class of economic models, firms *reveal* their individual reset price whenever they adjust their posted price. The history of price adjustments can therefore provide a benchmark for the path of individual reset prices. Individual price changes can also reveal the amount of price pressure having accumulated between adjustment periods, making the distribution of price changes an informative object about underlying inflation. That said, separating price pressure originating from idiosyncratic, sectoral, and aggregate sources is a delicate task. We illustrate this challenge by showing the limitations of a method proposed recently by Bils, Klenow, and Malin’s (forthcoming, henceforth “BKM”) to recover a time series of the common component of innovations to individual reset prices.

We next show how one can use variation in the shape of the distribution of nonzero price changes to derive bounds on the importance of adjustment along the extensive margin in response to aggregate shocks. Our procedure compares the shape of the distribution following a shock to an estimate of the distribution that would have prevailed in its absence. Any evidence that the two distributions differ by more than a lateral shift points to a role for the extensive margin. We apply our procedure to the study of four macroeconomic shocks whose timing and size are rather well identified: the sudden devaluation of the Mexican peso in late 1994, the April 1995 and January 2010 hikes in the Mexican value-added tax (VAT), and the collapse of oil prices in the fall of 2008. Our results indicate that a substantial share of the long-run response to these shocks was passed through upon impact, consistent with a relatively flexible aggregate price level. In the case of the devaluation and the VAT hikes, the extensive margin was the primary contributor to this price level flexibility. For the drop in world energy prices in 2008, however, the role played by the extensive margin in the transmission of this shock to retail energy prices was minor because a majority of these prices would have been updated absent the shock.

The role of the extensive margin in aggregate price adjustment can be especially important when there is a selection effect by which items whose posted prices are relatively far from their individual reset prices are more likely than others to have the timing of their price adjustments altered by shocks. Such a selection effect is highlighted in the work of Caplin and Spulber (1987) and Golosov and Lucas (2007), and clarified in CE, as the root of the lack of intrinsic persistence in standard menu-cost models. In the absence of a selection effect, the accumulation of aggregate price pressure should lead to a rise in the average size of individual price changes, whereas the accumulation of idiosyncratic price pressure should increase the dispersion of price changes. To check whether this is the case empirically, we first use the massive repricing of items created by the Mexican VAT hike in April 1995 to compute the distribution of consumer price changes in an environment dominated by aggregate reset price inflation. We find no drift in the average price change even as large amounts of aggregate reset price inflation accumulate. We then find no compelling evidence that the accumulation of idiosyncratic shocks in low-inflation environments leads to a rise in the dispersion of price changes in U.S. CPI data.

We also find support for a selection effect by studying how price adjustment decisions relate to deviations from the price of local competitors, which we use as a proxy for the individual reset price. We undertake this task using the IRI Marketing database, a very large (and relatively new to macroeconomists) dataset of weekly scanner prices from grocery stores and drugstores across the United States. The exceptional coverage of the dataset makes it possible to track prices of items with identical UPC across multiple outlets within a city. We follow Campbell and Eden (2010) in calculating the deviation from the average price of competitors at the UPC-market level. After accounting for permanent differences across stores, we find somewhat limited dispersion in the level of prices within local markets. Some of the observed dispersion is due to firms choosing prices away from those of their local competitors, which cautions against interpreting all forms of price dispersion as evidence of resource misallocations. We also find support for an adjustment probability that increases in our proxy of the (absolute) deviation. Like the limited dispersion, this fact is consistent with the presence of a selection effect and a role for adjustment along the extensive margin. These findings complement those of Eichenbaum *et al.*'s (2011), who document that movements in an item's replacement cost strongly influence its price adjustment probability for a chain of supermarkets. Finally, we use our estimates of the distribution of deviations and the probability of price change conditional on the deviation to compute the initial inflation response to shocks directly from the micro data. This exercise again points to an economically important role played by adjustment along the extensive margin in response to shocks.

After accounting for permanent differences across stores belonging to different chains, we find somewhat limited dispersion in the level of prices within local markets. Some of the observed dispersion is due to firms choosing prices away from those of their competitors, which cautions against interpreting all forms of price dispersion as resource misallocations. We also find evidence that the adjustment probability increases in our proxy of the deviation. Like the limited dispersion, this fact is consistent with the presence of a selection effect and a role for adjustment along the

extensive margin.

The paper is organized as follows. Section 2 presents CE’s decomposition and highlights the importance of individual reset prices in connecting micro and macro price adjustment. Section 3 discusses various approaches to inferring the importance of adjustment along the extensive margin using observed pricing behavior. Section 4 derives bounds on the importance of the extensive margin in the adjustment to four large macroeconomic shocks. Section 5 uncovers evidence of a selection effect in micro price adjustment by analyzing the distribution of price changes as price pressure builds up, and by estimating the adjustment hazard. The last section offers some concluding remarks.

2 Conceptual Framework

We begin this section by introducing the terminology and conceptual framework used throughout our paper to connect micro and macro price stickiness. We next discuss the main features of CE’s baseline decomposition and of its extensive margin in particular. We finally explore the implications of real rigidities and information frictions on the decomposition.

2.1 Economic Environment

We assume that time is discrete and that the length of time intervals matches that of price collection. The economy is populated by a continuum of firms indexed by $i \in [0, 1]$. Each firm produces a single item sold directly to consumers. Due to the presence of nominal rigidities, the *posted price*, $p_{i,t}$, may deviate from the “target” price, $p_{i,t}^*$, between price adjustment periods (all prices are in natural logs). The target price corresponds to the posted price chosen by a firm granted a *one-time* opportunity to adjust it freely, with all other constraints remaining in place. In choosing $p_{i,t}^*$, the firm takes into account how today’s posted price impacts profitability in the current and future periods; therefore, it generally differs from the price that maximizes current-period profits.⁵ We will refer to $p_{i,t}^*$ as the firm’s *individual reset price*. Our terminology seeks to highlight the connection between $p_{i,t}^*$ and the prices actually chosen by the firm when resetting them. It also ties in with the idea of aggregate *reset price inflation* studied by BKM, who aim to capture the component of innovations to individual reset prices that is common across items experiencing a price adjustment.

Because the individual reset price is a forward-looking object, it can be a complicated function of how the economy is expected to evolve over time, and how that evolution impacts future price adjustment decisions. In our benchmark specification, we assume that innovations to $p_{i,t}^*$ are described by

$$\Delta p_{i,t}^* = \pi_t^* + \nu_{i,t}, \tag{1}$$

⁵The notion of a price that maximizes current-period profits —either in partial equilibrium or in a frictionless general equilibrium environment— and the notion of a price that maximizes the present discounted stream of profits at times overlap in the work of Caballero and Engel (e.g., Caballero and Engel [1993b]). In a menu-cost model, the two concepts coincide when innovations to the frictionless optimal price follows a random walk with no drift and firms set prices under certainty equivalence. Departure from either of these assumptions breaks the equivalence between the two notions.

where π_t^* is aggregate reset price inflation and $\nu_{i,t}$ is a mean-zero idiosyncratic component that is *iid* across firms and over time. Our specification of $\nu_{i,t}$ introduces a random-walk element in individual reset prices.⁶ We make no particular assumption regarding the process describing the aggregate component of individual reset prices, except that its innovations, π_t^* , are common across items. For consistency, we also impose that firms share the same price-setting mechanism, to be described shortly. These simplifying assumptions are unlikely to hold in reality. For instance, differences in price stickiness, real rigidities, or item durability can lead to asymmetric responses to aggregate shocks.⁷ In our empirical implementation, we will partially address this concern by calculating our findings on more homogenous groups of products.

The decision to change the posted price is made after observing the shocks and aggregate variables in the period. We define the deviation from the individual reset price as $x_{i,t} = p_{i,t-1} - p_{i,t}^*$, where $p_{i,t-1}$ is the firm’s posted price inherited from the previous period. Following CE, we postulate that the probability of observing a price change is a time-invariant smooth function of the deviation from the individual reset price, $\Lambda(x_{i,t})$, called the “adjustment hazard function.” The model gives rise to infrequent and lumpy price adjustments, which is a central feature of consumer price data that we seek to reproduce.⁸

One appealing aspect of postulating a smooth adjustment hazard function is that forecasting individual price adjustments is typically a difficult task, as periods marked by frequent price changes are sometimes followed by long spells of inaction (and vice-versa) with no apparent changes in economic conditions. The adjustment hazard captures this randomness by leaving some uncertainty regarding the *timing* of adjustments. Its dependence on $x_{i,t}$ allows for the possibility that large deviations, which are suboptimal from the point of view of profit maximization, are more likely than small ones to trigger price adjustments. The function can also embed asymmetries in the response to positive and negative deviations that some authors argue could play a role in explaining apparent differences in the aggregate inflation response to positive and negative shocks (e.g., Caballero and Engel [1993b]). In what follows, we treat the adjustment hazard as a function of $x_{i,t}$ alone. All results carry through if we instead consider $\Lambda(x_{i,t}, \epsilon_{i,t})$, where $\epsilon_{i,t}$ is a vector of idiosyncratic states (e.g., month-specific and duration-specific dummies capturing seasonal patterns and duration dependence, respectively) that influence the adjustment probability but otherwise have no impact on how a firm’s individual reset price responds to an aggregate shock.⁹

⁶In standard time-dependent models with CES demand and constant returns to scale, one can express the individual reset price as a weighted sum of the current and future marginal costs (without linearizing the firm’s first-order condition). All else equal, random-walk innovations to the marginal cost thus enter $p_{i,t}^*$ additively.

⁷Barsky et al. (2007) and Gopinath and Itskhoki (2010) highlight this asymmetry in the presence of differences in durability and the degree of real rigidities, respectively.

⁸The model was extensively studied and developed in a series of papers by Caballero and Engel (1993a, 1993b, 1999).

⁹In theory, the adjustment hazard could depend on aggregate conditions beyond their influence on $x_{i,t}$. For example, Sheshinski and Weiss (1977) prove in a baseline menu-cost model that the width of the Ss band increases with steady-state inflation. The scant empirical evidence available suggests weak relationships at best. Gagnon (2009) reports that the average size of price increases and of price decreases was little impacted by the burst in inflation and that accompanied the Mexican Peso crisis. Wulfsberg (2009) uncovers no apparent change in the average absolute size of price adjustments in Norway as trend inflation fell from around 10 percent in the mid-1970s to about 2 percent

2.2 Price Adjustment along the Intensive and the Extensive Margins

Given the above assumptions, consumer price inflation can be expressed as

$$\pi_t = - \int x \Lambda(x) f_t(x) dx, \quad (2)$$

where $f_t(x)$ is the density of deviations from individual reset prices prevailing at the beginning of period t . Consider the impact on inflation of an aggregate shock to π_t^* taking place immediately before price adjustment decisions. For now, we assume that the shock, Δm , is passed through one-for-one to individual reset prices, thus ruling out real rigidities. We shall return to the implications of real rigidities shortly. The shock shifts the distribution of deviations from individual reset prices by $-\Delta m$, resulting in observed inflation

$$\pi_t(\Delta m) = - \int (x - \Delta m) \Lambda(x - \Delta m) f_t(x) dx. \quad (3)$$

Taking a first-order Taylor series expansion of $\pi_t(\Delta m)$ around $\Delta m = 0$, rearranging terms, and taking the limit as $\Delta m \rightarrow 0$, one obtains

$$F_t^{macro} = \lim_{\Delta m \rightarrow 0} \frac{\Delta \pi_t}{\Delta m} = \underbrace{\int \Lambda(x) f_t(x) dx}_{\text{intensive margin } (\mathcal{A}_t)} + \underbrace{\int x \Lambda'(x) f_t(x) dx}_{\text{extensive margin } (\mathcal{E}_t)}. \quad (4)$$

The above statistic captures the share of an (infinitesimal) aggregate shock to individual reset prices passed-through to posted prices upon impact. CE refer to this object as the *index of macroeconomic flexibility*, which they denote F_t^{macro} . This index has two components. The intensive margin, \mathcal{A}_t , captures the initial inflation response of items whose posted price would have been adjusted *absent the shock*. We refer to these price changes as being “predetermined.” The extensive margin, \mathcal{E}_t , captures the inflation contribution of items whose price adjustment is either *triggered or cancelled* by the occurrence of the shock, as hinted by the presence of $\Lambda'(x)$.

CE’s usage of the terms “intensive” and “extensive” margins differs from that popularized by Klenow and Kryvtsov (2008). The latter define the extensive margin as the frequency of price changes and the intensive margin as the average size of (nonzero) price changes. These definitions are motivated by Klenow and Kryvtsov’s investigation of how variation in the number of prices changes contributes to inflation dynamics. In contrast with these authors, who seek to explain the *level* of inflation, CE are only interested in the *boost* inflation due to the aggregate shock. And while the extensive margin in Klenow and Kryvtsov’s decomposition is a function of all price changes, the extensive margin in equation (4) depends only on the subset whose timing is altered by the shock. For example, a shock resulting in the simultaneous cancelling of a price decrease and triggering of a price increase would impact inflation solely through the intensive margin under Klenow and

in the mid-1990s. Both studies note a rise in the size of price adjustments since the late 1990s; however it is difficult to attribute this rise to changing aggregate conditions.

Kryvtsov’s decomposition because it leaves the number of price adjustments unchanged. The same shock would instead operate entirely through the extensive margin under CE’s decomposition. To avoid any confusion, our terminology follows exclusively that of CE.

We note several aspects of equation (4) that are useful in understanding the nature of the two margins and in assessing their empirical importance. First, the intensive and the extensive margins are functions of deviations from individual reset prices, which are typically unobserved. Under the assumptions made thus far, however, firms *reveal* their individual reset price whenever they adjust their posted price. For example, if a quarter of prices are adjusted during a period, then a quarter of all individual reset prices are observed. Moreover, the size of price changes reveals the amount of price pressure that has cumulated since the last price adjustment. We call these two implications of our framework the *revelation principle*. This principle will play a central role in our empirical strategy, which is laid out in section 3.

Second, the intensive margin has a strong connection to observables. To see this, notice that integrating the adjustment hazard over the distribution of deviations gives the observed frequency of price changes. In environments with low and stable inflation, the fraction of adjusting prices does not vary much in response to moderate movements in inflation due to offsetting variation in the number of price increases and price decreases (see Klenow and Kryvtsov [2008] and Gagnon [2009]). For these environments, the average frequency of price changes, \overline{fr} , offers a reasonable approximation of the intensive margin, \mathcal{A}_t . Appendix B provides estimates of the average frequency of price changes for the U.S. CPI and IRI Marketing scanner database. The extensive margin is more difficult to relate to observables because disentangling items whose price adjustment is triggered or cancelled by the shock from those for which it is predetermined is a challenging task. That said, in section 4, we will show how one can use changes in the *shape* of the distribution of price changes to bound its quantitative importance.

Third, the decomposition offers an intuitive way of distinguishing between time-dependent models, for which all price changes are predetermined, and state-dependent models, for which the timing of price changes can be impacted by shocks. In time-dependent models, such as the popular Calvo (1983) model and fixed-duration contracts (Taylor 1980), we have $F_t^{macro} = fr_t$. More generally, the frequency of price changes is a lower bound on the macroeconomic flexibility index because the extensive margin is typically positive. Put differently, the extensive margin creates a wedge between the observed frequency of price changes and the index of macroeconomic flexibility index.

Fourth, there need not be many price changes triggered or cancelled by a shock for the extensive margin to be a major contributor to the inflation response. In the presence of a *selection effect* by which items with much price pressure are especially likely to have their adjustment either triggered or cancelled by the shock, a small number of items revising the timing of their price adjustment can have large impact on inflation. For example, if a 1-percent shock to π_t^* induces a firm to raise its price by 10 percent rather than keep it constant (thus releasing pressure accumulated from a variety of sources since its last price adjustment), then the impact of that single price change on inflation will be as large as that of 10 predetermined price changes increasing by an extra 1 percentage point.

Caplin and Spulber (1987) and Golosov and Lucas (2007) trace the lack of intrinsic persistence in standard menu-cost models to this selection effect.

Fifth, the shape of the distribution of deviations, $f_t(x)$, also affects the magnitude of the selection effect in the period. If $f_t(x)$ is large in regions where the adjustment hazard is very steep and $|x|$ is large, then the extensive margin will be important. In periods of elevated macroeconomic instability, several shocks may be hitting the economy so that the distribution of deviations is more skewed in some direction than others. Seasonal patterns of price adjustments may also impact the distribution of deviations. Identifying the distribution of deviation is challenging due to its dependence on the history of shocks. To circumvent this difficulty, much of our analysis abstracts from characterizing the full dynamic process of inflation, focusing instead on instances when a “clear and large shock” can be easily identified.

Finally, expression (4) was derived for an infinitesimal shock Δm . In the data, large shocks are not only the easiest to identify but also the most likely to alter the timing of price changes. Inference about the relative importance for shock pass-through of price changes that are predetermined versus those that are triggered or cancelled is thus likely to depend on the magnitude of the shock. One can decompose the initial pass-through to a shock of arbitrary size as

$$\frac{\Delta\pi_t}{\Delta m} = \mathcal{A}_t + \mathcal{E}_t + \mathcal{O}_t(\|\Delta m\|), \quad (5)$$

where $\mathcal{O}_t(\|\Delta m\|)$ is the sum, scaled by Δm , of all terms of order two or higher in the Taylor-series expansion of $\pi_t(\Delta m)$ around $\Delta m = 0$. The intensive margin is now the frequency of price changes that would have been observed *absent the shock*. The object $\mathcal{O}_t(\|\Delta m\|)$ converges to zero as $\Delta m \rightarrow 0$. Importantly, it equal zero whenever all price adjustments are predetermined. For this reason, our strategy for bringing equation (5) to the data, which is presented in section 3.3 and implemented in section 4, lumps together \mathcal{E}_t and $\mathcal{O}_t(\|\Delta m\|)$. We will refer to any price movement that contributes to \mathcal{E}_t or $\mathcal{O}_t(\|\Delta m\|)$ as an adjustment *along* the extensive margin.

2.3 Real Rigidities

We have so far abstracted from the question of real rigidities, which dampened the response of individual reset prices to nominal shocks. Several authors have argued that such rigidities are essential for sticky-price models to generate the observed inertia in aggregate inflation and output data.¹⁰ Attempts to infer their importance from micro price behavior have yielded mixed results.¹¹

¹⁰See Blanchard and Fisher (1989) and Ball and Romer (1990) for early expositions. Woodford (2003a), Christiano, Eichenbaum, and Evans (2005), and Smets and Wouters (2007) argue, as do many others, that the inclusion of real rigidities in DSGE models improves their ability to account for observed aggregate economic dynamics.

¹¹Bils and Klenow (2004) note that CPI inflation since the turn of the 1990s is far too volatile and transient to be explained by standard sticky-price models, even absent real rigidities. Subsequent work with co-authors (Klenow and Willis [2006], BKM) reinforced their initial conclusion. Burstein and Hellwig (2007) calibrate a menu-cost model with decreasing returns to scale to the dynamics of retail prices and market shares. Their preferred parametrization is consistent with moderate pricing complementarities but the implications for aggregate dynamics are modest overall. Gopinath and Itskhoki (2011) document that exchange rate movements are passed-through to trade prices over more than one price adjustment, but the strategic complementarities they consider (variable markups) have limited

We will not seek to reconcile the micro and macro evidence on real rigidities, but simply stress that the key properties of equation (4) highlighted above extend to environments where real rigidities are present with minor qualifications. To this end, we follow CE in assuming that aggregate reset price inflation can be described as

$$\pi_t^* = (1 - a) \Delta m_t + a \pi_t. \quad (6)$$

The parameter a controls the extent to which firms trade-off raising their price in line with the price of their competitors versus matching the rate of money growth, Δm_t . This money growth term should be interpreted as standing for a host of nominal shocks impacting π_t^* , such as nominal wages or the nominal exchange rate. The elasticity a is an index of real rigidities; the smaller is a , the faster the response of individual reset prices to money growth. The process described by equation 6 should not be seen as a general description of the path for π_t^* but rather as a local approximation taken at a particular point in time.¹²

As was the case earlier, a shock Δm shifts the entire distribution of deviations from individual reset prices by an amount that now depends on the extent of pricing complementarities. Taking a first-order Taylor series expansion of $\pi_t(\Delta m)$ around $\Delta m = 0$, rearranging terms, and letting $\Delta m \rightarrow 0$, one gets, as stated in CE,

$$F_t^{macro} = \lim_{\Delta m \rightarrow 0} \frac{\Delta \pi_t}{\Delta m} = \frac{(1 - a)(\mathcal{A}_t + \mathcal{E}_t)}{1 - a(\mathcal{A}_t + \mathcal{E}_t)}. \quad (7)$$

The presence of real rigidities makes the index of macroeconomic flexibility a nonlinear function of the sum of the intensive and the extensive margins. This function is strictly increasing in $\mathcal{A}_t + \mathcal{E}_t$, is bounded between 0 and 1 when $\mathcal{A}_t + \mathcal{E}_t \in (0, 1)$, and is decreasing in the extent of real rigidities (i.e., $\frac{\partial \Delta \pi_t}{\partial a} < 0$).¹³ As in the baseline case with no real rigidities, the extensive margin may add substantially to the flexibility of the price level. The revelation principle continues to hold: Individual reset prices are revealed through price adjustments and the size of price changes corresponds to the price pressure that has accumulated since the previous adjustment. Moreover, a shock Δm to π_t^* still results in a lateral shift of the distribution of deviations, but its magnitude,

$$\Delta \pi_t^* = \left(\frac{1 - a}{1 - a(\mathcal{A}_t + \mathcal{E}_t)} \right) \Delta m,$$

implies incomplete initial pass-through of Δm to π_t^* whenever $\mathcal{A}_t + \mathcal{E}_t$ is less than 1.

macroeconomic effects unless coupled with exogenous sources of persistence. Also, Eichenbaum, Jaimovich, and Rebelo (2011) find rapid pass-through of movements in a retailer's cost to prices.

¹²Burstein and Hellwig (2007) provide microfoundations for the reduced-form parameter a . In their environment, firms have a CES demand function with elasticity θ , use production functions with decreasing returns in labor of the form $y_i = z_i (l_i)^\alpha$, and face nominal wages determined by $W_t = (M_t)^{1-\gamma} (P_t)^\gamma$. Their linearized solution around the frictionless steady-state has $a = 1 - \frac{1-\alpha\gamma}{\alpha+\theta-\alpha\theta}$.

¹³While intuitive, this latter property does not hold in all environments. Dotsey and King (2005) present an example in which an increase in real rigidities increases the flexibility of the price level. In their state-dependent model, the negative impact on F_t^{macro} of an increase in a is more than offset by the positive impact of a rise in $\mathcal{A}_t + \mathcal{E}_t$ as more firms find it profitable to incur the fixed cost of adjusting prices.

2.4 Informational Frictions

The framework used thus far assumes that firms have full information at all times. There are contexts with information frictions for which the CE decomposition also holds. Information frictions alone are insufficient to generate *infrequent* price adjustments. Absent some fixed cost of adjusting posted prices, firms would change them continuously by an amount consistent with the information they have at hand.¹⁴

A popular approach to integrating information frictions into menu-cost models is to bundle the fixed cost of changing the posted price with a fixed cost of revealing the state of the economy. In these models, we can define the individual reset price as the posted price chosen by a firm granted a *one-time* opportunity to jointly reveal the state of the economy and adjust the posted price for free. Individual reset prices are not known with certainty between adjustment periods. In particular, some such models assume that firms receive no information between adjustment periods. They include Bonomo and Garcia (2001), Bonomo and Carvalho (2004), and variants of Bonomo *et al.* (2010). The optimal strategy is then to update the posted price at fixed time intervals, so that the adjustment hazard is independent of $x_{i,t}$ and the extensive margin is effectively shut down (i.e., $\Lambda(x_{i,t}, \epsilon_{i,t}) = \Lambda(\epsilon_{i,t})$, where $\epsilon_{i,t}$ is the time elapsed since the last information update). In other such models, firms receive signals that are informative about $x_{i,t}$. In Gertler and Leahy (2008), firms are notified when an infrequent but volatile idiosyncratic shock hits but are not informed of its size. The optimal strategy is to update the information set and then change the posted price under full information whenever the idiosyncratic shock hits. Because the arrival pattern of these shocks follows a Poisson process, the adjustment hazard is effectively constant, as in the Calvo (1983) model. In Woodford (2009), firms receive a noisy signal about the aggregate shock that can trigger price reviews. His model provides foundations for an adjustment hazard that is increasing in the absolute deviation from the individual reset price. Although the information friction models above differ in terms of the shape of the adjustment hazard, the revelation principle holds in all of them and they are amenable to the CE decomposition.

When the fixed costs of changing the price and of revealing the state of the economy are separate, firms may adjust prices using outdated information (see Alvarez *et al.* [2011] for a recent treatment). In such a situation, both the timing and size of some price changes are determined ahead of the shock. The revelation principle no longer applies to these situations as they break with the maintained assumption that price changes correspond to a reoptimization. In particular, it is no longer the case that all predetermined price changes respond to the current shock. These models thus lie outside our analytical framework.

¹⁴Many leading models of information frictions, such as Mankiw and Reis (2002), Sims (2003), Mackowiak and Wiederholt (2009), and Woodford (2003b), feature continuous rather than infrequent micro price adjustments. As such, they are not designed to explain nominal stickiness at the item level unless they are enriched with other frictions. In these models, all price changes are predetermined so that the adjustment hazard trivially equals 1. The information frictions act as a real rigidity that impedes the transmission of shocks to π_t^* . Recent advances in rational inattention models are consistent with firms selecting among a finite array of real prices (see Matejka and Sims [2011]). It is unclear to us whether these models could be amenable to explaining nominal price stickiness in the presence of nonzero aggregate inflation.

Finally, there are other menu-cost models in which the state of the economy is *never* fully revealed (e.g., Gorodnichenko [2010]). One can still think of the inflationary impacts of an aggregate shock in terms of the contribution of items whose price change is predetermined versus items whose timing of adjustment is altered by the shock. A muted initial response of the price of adjusters to the shock would be akin to the real rigidities explored above. That said, such models typically assume that firms receive heterogeneous signals about the state of the economy so that their response to shocks differs, contrary to the assumption maintained so far.

3 Inferring Individual Reset Prices and their Effects

A key challenge for bringing equation (4) to the data is that deviations from individual reset prices are typically not observed. Estimating them is a challenging task because they may depend on more aggregate, sectoral, and idiosyncratic factors than the econometrician has at hand. In this section, we first argue that partial identification is possible using only micro price data by focusing on instances with price changes. We then use the CE decomposition to discuss a method proposed by BKM to infer individual reset prices between adjustment periods. We finally derive two sets of bounds on the share of initial pass-through attributable to adjustment along the extensive margin in response to well identified shocks.

3.1 The Revelation Principle and the Distribution of Price Changes

Under the revelation principle, individual reset prices are observed whenever posted prices are adjusted. As a concrete example, Bils and Klenow (2004) report an average frequency of price changes of 26 percent in the U.S. CPI in the mid-1990s. Applying the principle to these data implies that $p_{i,t}^*$ is observed about a quarter of the time. The revelation principle also leads to new interpretations of the distribution of price changes. For items experiencing a price adjustment, we have $\Delta p_{i,t} = -x_{i,t}$. The distribution of price changes thus maps into the distribution of deviations from individual reset prices that prevailed at the beginning of the period for price adjusters. Alternatively, $\Delta p_{i,t}$ equals the cumulative change in the individual reset price since the last nominal adjustment. For a firm changing its posted price after keeping it constant for τ periods, we have $\Delta p_{i,t} = p_{i,t}^* - p_{i,t-\tau}^*$. As such, the distribution of price changes can be used to extract information on how price pressure builds over time at the item level. For instance, one can condition this distribution on the number of periods since the last price adjustment (we use this approach in section 5). These observations suggests that the distribution of price changes is a richly informative object worthy of the attention received thus far in the empirical literature, and that replicating this distribution in macro models is a key step towards ensuring consistency with microeconomic facts.¹⁵

¹⁵See Eden (2001) for an early investigation of how the shape of the distribution of consumer price changes differs between low and high inflation. For U.S. evidence on consumer prices, see Klenow and Kryvtsov (2008), Klenow and Malin (2010), and Berger and Vavra (2011). Calibrations to the distribution of price changes have appeared recently in the literature (e.g., Woodford [2008], Midrigan [2011], and Costain and Nakov [2011])

These interpretations come with some qualifications. First, Chevalier and Kashyap (2011) present empirical evidence that the price of competing brands may be jointly determined by retailers, so that, at least in the retail trade sector, individual price decisions may not be seen in isolation from those of their close substitutes. Second, as hinted in section 2.4, there could be environments in which both the timing *and* size of price adjustments are pre-determined; for example if firms followed price plans or pre-announced price changes. Price adjustments that do not coincide with a reoptimization obscure the inference about the set of firms adjusting their posted prices in response to shocks and the revelation principle breaks down. That said, there is limited support for widespread use of pre-announcements, at least at the retail level, and of price plans by firms. Fabiani et al. (2006) report that price reviews are more common than price changes among European firms. Moreover, Alvarez *et al.* (2011) argue that adjusting prices without reoptimizing them is a suboptimal strategy in low-inflation environments. Finally, there could be situations in which the individual reset price is not unique; for example, when firms follow mixed pricing strategies. The revelation principle does not hold in this case because the size of price changes need not equal the amount of price pressure having accrued since the last adjustment.¹⁶ We leave the empirical relevance of multiple individual reset prices as an open issue.

3.2 The Extensive Margin and Bils-Klenow-Malin Reset Price Inflation

Imputing individual reset prices between adjustment periods is a delicate task. With only price information at hand, one does not know if an item’s price change is releasing pressure generated by single large shock or accumulated gradually. In a recent paper, BKM propose a method that uses the behavior of items changing their price to infer the amount of price pressure building up for other items in the sample. We briefly review their approach and then use the CE decomposition to illustrate its shortcomings when adjustment along the extensive margin is important.

3.2.1 BKM’s Identification Method

BKM estimate item i ’s individual reset price in period t as follows

$$p_{i,t}^{*BKM} = \begin{cases} p_{i,t} & \text{if } p_{i,t} \neq p_{i,t-1} \\ p_{i,t-1}^* + \pi_t^{*BKM} & \text{if } p_{i,t} = p_{i,t-1} \end{cases} . \quad (8)$$

They impute the individual reset price of items whose posted price is unchanged by incrementing their previous-period imputed value by an estimate of aggregate reset price inflation, π_t^{*BKM} . This estimate corresponds to the average change in $p_{i,t}^{*BKM}$ of items whose price is adjusted in period t . A time series for π_t^{*BKM} is constructed by initializing individual reset prices and then recursively computing π_t^{*BKM} . The dependence of π_t^{*BKM} on the initial conditions quickly fades as price

¹⁶ A multiplicity of optimal prices can arise when the profit function is very flat because firms tradeoff the unit markup and the number of consumer purchases. See Eden’s (1994) model of uncertain and sequential trade and the search model of Head *et al.* (forthcoming).

adjustments occur. The previous expression (8) can be rewritten as

$$p_{i,t}^{*BKM} = \begin{cases} p_{i,t} & \text{if } p_{i,t} \neq p_{i,t-1} \\ p_{i,t-\tau} + \sum_{s=0}^{\tau-1} \pi_{t-s}^{*BKM} & \text{if } p_{i,t} = p_{i,t-1} \end{cases},$$

where τ is the number of periods since item i 's last price adjustment. This data-generating process differs from that assumed in expression (1), which implied that item's i actual individual reset price evolves according to

$$p_{i,t}^* = \begin{cases} p_{i,t} & \text{if } p_{i,t} \neq p_{i,t-1} \\ p_{i,t-\tau} + \sum_{s=0}^{\tau-1} \pi_{t-s}^* + \sum_{s=0}^{\tau-1} v_{i,t-s} & \text{if } p_{i,t} = p_{i,t-1} \end{cases}. \quad (9)$$

The term $\sum_{s=0}^{\tau-1} v_{i,t-s}$ captures the accumulation of idiosyncratic shocks since the last price adjustment. Comparing the above two equations, it is readily apparent that BKM's procedure abstracts from the presence of idiosyncratic shocks. Provided that idiosyncratic shocks are mean zero and, importantly, do not influence the set of firms experiencing a price change in response to shocks, then π_t^{*BKM} is an unbiased estimator of π_t^* . As we now show, a bias occurs when idiosyncratic shocks are present and firms adjust along the extensive margin.

3.2.2 BKM's Identification and the CE Decomposition

For ease of exposition, consider an economy that starts in a steady state with $\pi_t^{*BKM} = \pi_t^* = \bar{\pi}$ and in which actual individual reset prices evolve according to equation (9). Suppose that all individual reset prices are perturbed by a one-time impulse Δm prior to price adjustment decisions. Let $fr(\Delta m)$ and $dp(\Delta m)$ denote the initial response of the average frequency and size of (nonzero) price changes, respectively. The inflation innovation attributable to the shock can be written as

$$\Delta\pi = fr(\Delta m) \cdot dp(\Delta m) - fr(0) \cdot dp(0),$$

where $fr(0)$ and $dp(0)$ represents the pre-shock values of these variables. Klenow and Kryvtsov (2008) show that nearly all the variation in U.S. CPI inflation over the past two decades is attributable to movements in $dp(\cdot)$ with movements in $fr(\cdot)$ being relatively small and weakly correlated with inflation. Using this observation to approximate $fr(\Delta m)$ and $fr(0)$ by the average frequency, \overline{fr} , we get

$$\Delta\pi \approx \overline{fr} \cdot (dp(\Delta m) - dp(0)) \approx \overline{fr} \cdot \Delta\pi_t^{*BKM}.$$

The last expression uses the fact that BKM's procedure imputes the average size of nonzero price changes to a movement in aggregate reset price inflation in the initial period. Finally, using expression (4) and $\mathcal{A}_t \approx \overline{fr}$, we obtain

$$\Delta\pi_t^{*BKM} \approx \left(1 + \frac{\mathcal{E}_t}{\mathcal{A}_t}\right) \cdot \Delta m. \quad (10)$$

In short, BKM’s method overestimates the innovation to reset price inflation by a factor that depends on the relative importance of the extensive margin. Absent adjustment along the extensive margin ($\mathcal{E}_t = 0$), π_t^{*BKM} would be an unbiased estimator of π_t^* . One implication is that π_t^{*BKM} is overly volatile relative to π_t^* whenever $\mathcal{E}_t/\mathcal{A}_t > 0$. Another implication is that the extra $(\mathcal{E}_t/\mathcal{A}_t) \Delta m$ wrongly imputed to the individual reset price of nonadjusters in period t needs to be offset in subsequent periods by increments to π_{t+i}^{*BKM} summing up to negative $(\mathcal{E}_t/\mathcal{A}_t) \Delta m$. A rapid offsetting could be a contributor to BKM’s estimate of a negatively persistent aggregate reset price inflation series. Figure 1 illustrates these effects in a baseline menu-cost model à la Golosov and Lucas (2007) in which $\mathcal{E}_t/\mathcal{A}_t$ is about 2.¹⁷ Upon impact of a 1-percent shock to individual reset prices, the BKM procedure imputes a change in individual reset prices of nonadjusters of about 3 percent. The procedure quickly corrects this overestimation in subsequent periods as firms responding to the shock with a delay release less price pressure than initially imputed.

We stress that the bias in BKM’s method is not due to the extensive margin in itself but rather to the *joint* presence of idiosyncratic shocks and the asymmetric effect of the shock on the probability of price change of items with positive and negative deviations.¹⁸ A positive aggregate shock triggers the release of price pressure due to idiosyncratic factors through the selection effect, but BKM’s procedure is not designed to disentangle that pressure from that accruing to the common component of individual reset prices.

BKM are well aware that their estimator may be biased when the data are state dependent. In their 2009 NBER working paper, they deal with this issue by applying their procedure to reset price inflation series generated from baseline Calvo and menu-cost models calibrated to match aggregate CPI inflation dynamics. They then compare the resulting aggregate reset price inflation series to that estimated on CPI micro data. The estimated series is highly volatile and exhibits either no or negative persistence. These features are most closely replicated by their baseline menu-cost model with no strategic complementarities. This model entails a major role for adjustment along the extensive margin in explaining U.S. inflation dynamics.

However, at least three factors other than the extensive margin could create spurious volatility and negative persistence in π_t^{*BKM} . As with the selection effect discussed above, these factors imply that the behavior of price adjusters is not representative of that of nonadjusters. First, the method is sensitive to sample size because it leverages the behavior of a minority of price adjusters

¹⁷The model assumes a symmetric Ss band, a process for individual reset prices as in equation (9), and normally distributed idiosyncratic shocks. Aggregate reset price inflation is set to the average monthly U.S. CPI inflation excluding shelter from 1988 to 2007. The width of the Ss band and the variance of idiosyncratic innovations are calibrated to match an average frequency and average absolute magnitude of nonzero price changes of 25 percent and 10 percent, respectively.

¹⁸In particular, BKM’s procedure is consistent when the data are generated from the Caplin and Spulber (1987) model for which all price adjustments occur along the extensive margin through the selection effect. That model has no idiosyncratic shocks, so there is no confusion that price adjustments reflect solely price pressure attributable to aggregate shocks. The method is also consistent when the data are generated from a Calvo model with idiosyncratic shocks for which the exogenous probability of price changes varies in response to aggregate conditions. Although this model has no selection effect, it still has an extensive margin due to variation in the frequency of price changes. In both of these examples, the frequency of price changes is sensitive to the size of the shock. The approximation of the frequency of price changes by a constant used to derive equation (10) would not be appealing in these environments.

to infer the price pressure accruing to all items. Any sampling uncertainty or coding mistakes in the subset of adjusters will thus be imputed to nonadjusters. Second, sectoral shocks may be confused with aggregate shocks. BKM deal with this issue by estimating separate reset price inflation series for 64 groups of products, but we see a risk of accentuating biases due to sample size. Third, the response of individual reset prices to aggregate shocks may be asymmetric across items in ways that correlate with the frequency of price changes. For instance, Gopinath and Itskhoki (2010) show that frequent adjusters pass-through a larger share of exchange rate movements than infrequent adjusters in a sample of U.S. trade prices. Kara (2011) provides a theoretical example in which frequent adjusters are more responsive to transitory shocks than infrequent adjusters.

Our discussion highlights the importance of incorporating firm-level information into the estimation of individual reset prices between adjustment periods (we will return to this task in section 5.2). Rather than inferring the importance of adjustment along the extensive margin from the dynamic properties of inflation, we concentrate for now on the more modest objective of estimating the role of the extensive margin for the initial response of price changes to clearly identifiable aggregate shocks.

3.3 Inference from the Distribution of Price Changes

For discrete shocks to π_t^* , we can relate price adjustment along the extensive margin (i.e., $\mathcal{E}_t + \mathcal{O}_t(\|\Delta m\|)$ in equation [5]) to the observed distribution of price changes, and to how that distribution differs from the one that would have been observed *absent* the shock (the “counterfactual distribution”). Although this counterfactual distribution is not observed, one can often infer its shape by considering the distribution that prevailed prior to the shock.

To introduce these ideas formally, let $\mu_t(\cdot)$ and $\tilde{\mu}_t(\cdot)$ be measures representing the observed and counterfactual distributions of price changes, respectively. These measures have a mass point at 0 corresponding to the fraction of items whose price is unchanged. They are otherwise equal to the density of individual price changes. Given a shock Δm raising π_t^* by $\Delta\pi_t^*$, we can rewrite equation (3) as

$$\pi_t(\Delta m) = - \int (x - \Delta\pi_t^*) \Lambda(x - \Delta\pi_t^*) f_t(x) dx = \int \Delta p_{i,t} d\mu_t(\Delta p_{i,t}).$$

The last equality uses the revelation principle to replace $x_{i,t} - \Delta\pi_t^*$ by $-\Delta p_{i,t}$ and then notes that $\Lambda(-\Delta p_{i,t}) f(-\Delta p_{i,t}) d(\Delta p_{i,t})$ equals $d\mu_t(\Delta p_{i,t})$. Absent the shock, inflation would have been

$$\pi_t(0) = \int \Delta p_{i,t} d\tilde{\mu}_t(\Delta p_{i,t}).$$

The difference between $\pi(\Delta m)$ and $\pi(0)$ is the initial boost to inflation attributable to the shock, $\Delta\pi_t(\Delta m)$. Adding and subtracting the inflation contribution of predetermined price changes,

$\int \Delta\pi_t^* d\tilde{\mu}(\Delta p_{i,t}) = \widetilde{f}r_t \Delta\pi_t^*$, performing a change of variable, and reorganizing terms, we get

$$\Delta\pi_t = \widetilde{f}r_t \Delta\pi_t^* + \int \Delta p_{i,t} (d\mu(\Delta p_{i,t}) - d\tilde{\mu}(\Delta p_{i,t} - \Delta\pi_t^*)). \quad (11)$$

The first and second terms on the right-hand side capture the contribution to $\Delta\pi_t$ of adjustments along the intensive and the extensive margin, respectively. The shock shifts the distribution of predetermined price changes by $\Delta\pi_t^*$, resulting in a boost to inflation of $\widetilde{f}r_t \Delta\pi_t^*$. The shock may also trigger or cancel price adjustments, in which case the observed distribution of price changes will differ from the counterfactual distribution shifted by $\Delta\pi_t^*$. Our decomposition above attributes any such difference to adjustment along the extensive margin.

If we observed $\Delta\pi_t^*$, then computing the respective contribution to $\Delta\pi_t$ of adjustment along either margin would be straightforward given an estimate of the counterfactual distribution of price changes. Absent knowledge of $\Delta\pi_t^*$, we can nevertheless compute bounds on the relative importance of adjustment along the extensive margin. We pursue two strategies. In our “full-pass-through” approach, we compute the adjustment along the extensive margin that is consistent full immediate pass-through of the shock to individual reset prices (and thus no real rigidities). To see how this approach provides a lower bound, note that, given $\Delta\pi_t$, the larger the initial response of individual reset prices, $\Delta\pi_t^*$, the larger the contribution to $\Delta\pi_t$ of adjustment along the intensive margin, $\widetilde{f}r_t \Delta\pi_t^*$. Assuming full immediate pass-through of the shock maximizes $\Delta\pi_t^*$ and thus minimizes the importance of adjustment along the extensive margin. In our second approach, we assume that there is no selection effect in the choice of items whose adjustment is either triggered or cancelled by the shock, only variation in the total number of price changes. This situation is similar to what would happen in a Calvo model in which the probability of price adjustment, while exogenous to the firm, responds to the state of the economy. For this reason, we will refer to this approach as “Calvo+.” A selection effect would tend to increase the contribution to $\Delta\pi_t$ of adjustment along the extensive margin by selecting items with relatively large deviations from their individual reset price. This lower bound is akin to the contribution of the extensive margin in a Calvo model in which the exogenous probability of price adjustment varies with the state of the economy. Without the shock, inflation would be given by $\tilde{\pi}_t = \widetilde{f}r_t \widetilde{d}p_t$, where $\widetilde{d}p_t$ is the average size of nonzero price changes. Following a shock Δm , and absent real rigidities, the average nonzero price change would increase by Δm (i.e., $\Delta\pi_t^* = \Delta m = dp_t - \widetilde{d}p_t$) and the frequency of price changes would rise to fr , so that inflation would equal $\pi_t = fr_t (\widetilde{d}p_t + \Delta m)$. The boost to inflation due to the shock can be written as

$$\Delta\pi_t = \widetilde{f}r_t \Delta m + (fr_t - \widetilde{f}r_t) (\widetilde{d}p_t + \Delta m). \quad (12)$$

Given the observed and the counterfactual distribution of price changes (i.e., with knowledge of fr_t , $\widetilde{f}r_t$, $\widetilde{d}p_t$, and $\Delta\pi_t$), one can recover Δm and the respective contributions to $\Delta\pi_t$ of adjustment along the intensive and the extensive margins. Absent the selection effect, the initial impact of the shock is to rescale the distribution of price changes by a factor $fr_t/\widetilde{f}r_t$ and translate it by $\Delta\pi_t^*$. Any evidence of that the observed and the counterfactual distributions of price changes differ by

more than a scaling factor (e.g., different dispersion or skewness) could suggest that a selection effect is at work.

4 Extensive Margin Adjustment to Large Shocks

We apply the bound analysis laid out above to the study of four macroeconomic shocks whose timing and impact on individual reset prices are relatively well identified. These shocks are the sudden collapse of the Mexican peso in late 1994, the April 1995 and January 2010 hikes in the Mexican VAT, and the drop in world energy prices surrounding the bankruptcy of Lehman Brothers in 2008.

4.1 The Late-1994 Peso Devaluation

Prior to the collapse of the peso in late 1994, Mexico was operating under a crawling peg system in which the exchange rate acted as the economy's *de facto* nominal anchor. Facing mounting pressure on its foreign exchange reserves, the government announced a 15-percent devaluation of the peso on December 19, 1994. The measure proved insufficient and was abandoned three days later in favor of a free-floating system. A week after the initial devaluation, the exchange rate had lost about 40 percent of its value vis-à-vis the U.S. dollar. Annualized inflation, which had been hovering around 7 percent in 1994, jumped to more than 40 percent in early 1995. The devaluation is a particularly interesting shock because it was large, broadly unexpected, and its timing is clear. Moreover, its occurrence late in December means that most prices had already been collected for that month.

The upper panel of figure 2 shows the distribution of monthly price changes observed in January 1995 (the “actual” distribution) and the distribution that prevailed 12 months earlier (the “counterfactual” distribution). The fraction of nonzero price changes doubled from 24.9 percent in January 1994 to 50.7 percent in January 1995, pointing to a major role of the extensive margin in the adjustment to the macro shock. The actual distribution lost some of its mass of negative price changes relative to the counterfactual, especially in the -10 percent to 0 percent range, while gaining significantly more mass on the positive side. The jump in the frequency of price changes was observed across all major groups of products and is perhaps clearest among product categories whose frequencies were initially low to moderate. Processed food and nonenergy industrial goods witnessed increases in the frequency of price changes of 30 percentage points or more. The distribution of unprocessed food, with its pre-shock frequency of price changes around 40 percent, already featured a large role for the intensive margin as a large number of price changes were predetermined. The arrival of the shock pushed the frequency up to about 55 percent.

We compute our two sets of lower bounds to gauge the importance of adjustment along the extensive margin. The first bound posits full immediate pass-through of the nominal exchange rate devaluation, Δe , to the price of imported final consumption items and imported intermediate inputs going into the production of final consumption items. If there are no real rigidities, the

change in the optimal reset prices is $\Delta\pi^* = \xi \cdot \Delta e$, where ξ measures the importance of imports in the consumption basket. Based on input-output tables, Burstein, Eichenbaum, and Rebelo (2005) estimate the total import content of Mexican consumption expenditures around that period to be 10.9 percent (and a direct import content only half as large), a figure reflecting Mexico’s low openness at the time. Setting $\xi = 0.109$ and $\Delta e = 0.5$ (roughly the peso depreciation in the month that followed its collapse), we get that individual reset prices should have jumped 5 percentage points under full pass through. The second bound is based on equation (12).

As seen in table 2, both approaches suggest that adjustment along the extensive margin was the dominant channel. Under the full-pass-through method, the extensive margin accounts for over two thirds of the initial price level response. The Calvo+ method provides a similar lower bound (63 percent). Table 2 also provides estimates by special groups of product using the same value of ξ . With the exception of unprocessed food, all special groups of items suggest contributions of the extensive margin in excess of 75 percent. The relatively small share for unprocessed food likely reflect their initially high frequency of price changes (38.4 percent) which, all else equal, tends to boost the role of the intensive margin. In addition, the import content of unprocessed food items may have been higher than 10.9 percent, which would tend to raise the lower bound on the extensive margin. Finally, we note that the shape of the actual distributions of price changes displayed in figure 2 appears inconsistent with the random selection of items whose price change is triggered by the shock. In particular, the positive skewness increased following the shock, which is suggestive of a role for the selection effect in the extensive margin, contrary to the assumption behind the Calvo+ method.

Our full-pass-through method assumed that the devaluation was the only shock to individual reset prices in late 1994. However, the rise in the price level in 1995 was much larger than implied by the devaluation alone, hinting that other factors also pushed up individual reset prices. Some of these factors arguably had limited impact initially. Notably, monetary policy was tight in early 1995 as the central bank pushed up overnight rates to contain the outflow of capital. In addition, labor costs did not accelerate until late 1995. Also consistent with a somewhat moderate initial response of individual reset prices is the change in inflation expectations. According to a survey of private forecasters conducted by the Bank of Mexico (reported by Capistrán and López-Moctezuma [2010]), respondents surveyed in January and February 1995 expected consumer prices to have risen 25 percent by the end of the year, roughly half the observed rise that year but significantly more than the 7 percent rise experienced in the twelve months to December 1994. As an alternative gauge of the amount of price pressure, we recomputed our bounds assuming that the initial shock to reset prices was 18 percentage points. This alternative scenario, also reported in table 2, implies that almost a third of the shock was passed-through to prices upon impact, with the extensive margin making a significant contribution in most groups of products.

On a related note, micro price studies have documented the presence of seasonal patterns in the timing of consumer price adjustments for many product categories. Several authors (e.g., Dhyne *et al.* [2005]) have related such patterns to time-dependent pricing models —and to fixed-

duration contracts in particular— in which the timing of price adjustments is exogenous to firms. A preference for implementing price changes at particular times of the year needs not imply that adjustment along the extensive margin is unimportant, however. The behavior of services prices in the wake of the peso crisis illustrates this point. From 1995 to 2000, a quarter of these prices were adjusted every January, a figure almost twice as large as in other months of the year (13 percent). As the inflationary consequences of the peso crisis drew to an end, the January frequency of price changes stepped down to below 15 percent in 2001 and stayed at that level thereafter. This step down offers further evidence that the number of firms adjusting their price is responsive to aggregate conditions, consistent with an economically important role of adjustment along the extensive margin.

4.2 Hikes in the Mexican Value-Added Tax

Changes in VAT are simple shocks to study because their timing and size are observable by economic agents. Mexican firms must include the VAT and other mandated duties in posted prices. Changes in VAT thus affect their effective mark-ups, creating an incentive for price adjustments. About a third of items in the Mexican CPI excluding housing rents are subject to either the VAT or special duties. Exempted items include food at home excluding beverages, most education services, patented drugs, and books and periodicals. Items subject to special duties include gasoline, phone lines, tobacco products, and alcoholic beverages. The most recent Mexican VAT hike occurred on January 1, 2010, when the rate rose by one percentage point throughout Mexico. The measure had been debated and adopted earlier in the fall, leaving firms a few price collection periods to prepare for its implementation. The VAT increase also coincided with a rise in special duties on a few products. A more substantial 5-percentage point hike took place on April 1, 1995, when the Mexican government, faced with an urgent need for revenues amid the peso crisis, increased the rate from 10 to 15 percent. Its large size aside, two features make the 1995 hike particularly interesting. First, retailers were given a short notice before its enactment: the measure was announced as part of an austerity package on March 9 and officially adopted in a decree on March 18. Second, firms located in the tourism regions of Baja California and Quintana Roo, and within a narrow corridor along the country's southern and northern international borders, were exempt from the hike, with their rate remaining at 10 percent. Prices in these regions can thus be used as a control group to analyze the impact of the tax.

Figures 3 and 4 show the observed distributions of price changes associated with the 2010 and 1995 VAT hikes, respectively. For the 2010 hike, the counterfactual distribution is the average density of price changes in the month of January for the years 2003 to 2007, a period over which the macroeconomic environment was rather stable. We use only January observations to control for seasonal repricing. For the 1995 hike, the counterfactual distribution is the one that prevailed in the regions not subject to the tax increase. This regional difference approach helps control for macroeconomic shocks other than the VAT hike, such as the peso devaluation and the contraction in aggregate demand, that likely influenced pricing decisions at the time.

For items subject to the VAT hike, we find a strong response of the frequency of price changes in the sample (see upper panels of figures 3 and 4). Price changes were more than twice as frequent in January 2010 (38.5 percent) than in January of the years 2003 to 2007 (average of 17.9 percent). The April 1995 hike is associated with an even larger difference in the frequency of price changes between the observed and counterfactual distributions; 75.6 percent of items subject to the tax hike changed their price that month compared to 48.0 percent for taxable items in regions exempt from the hike. As is apparent from figures 3 and 4, a large proportion of price changes clustered around 1 percent in 2010 and especially 5 percent in 1995. Many firms also seized the opportunity to raise prices by more than the hike, providing evidence that price pressure from other sources was being released in the process. For items exempted from the VAT hikes, such as food products, the observed and counterfactual distributions of price changes had similar shape during both VAT hike episodes.

Under full immediate pass-through of the tax and duty hikes, we estimate that the price of taxable items should have rise 0.90 percent in 2010 and 4.45 percent in 1995. As reported in table 3, the observed rate of inflation for taxable items in 2010 was 0.65 percentage point higher than the counterfactual rate, consistent with the immediate pass-through of roughly three quarters of the hike. Pass-through was especially high for taxable nonenergy industrial goods (83.6 percent) but only about a third for taxable services. The difference in inflation for taxable items in 1995 between regions subject to and exempt from the hike is consistent with a somewhat lower pass-through rate of 62.8 percent. However, this estimate likely understates the extent of actual pass-through in 1995. Inflation was running higher in exempt regions in early 1995, consistent with border regions experiencing more price pressure from the devaluation. Applying the inflation differential for non-taxable items in April 1995 to taxable items suggests a pass-through rate closer to 80 percent.

Table 3 also reports our lower bounds on the relative importance of adjustment along the extensive margin. For the 2010 hike, the full-pass-through method suggests that at least three quarters of the observed jump in $\Delta\pi$ was attributable to adjustment along the extensive margin, a finding that holds for both goods and services. The corresponding lower bounds for 1995 are noticeably smaller at about a quarter or less. This result should not come as a surprise in light of the high frequency of price changes at the time under the counterfactual scenario, which implies that a large number of price adjustments were determined ahead of the tax adjustment. These results offer an interesting contrast between the sources of adjustment to the two VAT hikes. While both episodes shared similar degrees of initial pass-through, price adjustment in the low-inflation environment of 2010 operated to a greater extent via the extensive margin—despite a smaller shock—than in the high-inflation environment in 1995. In both cases, the index of macroeconomic flexibility, which is captured by our initial pass-through measure, was very high.¹⁹

We also note that several implications of the Calvo+ lower bound method are counterintuitive. Absent a selection effect, π^* is revealed by the difference between the observed and counterfactual

¹⁹Our Mexican sample does not allow us to explore asymmetries in price adjustments to VAT increases and decreases. Karadi and Reiff (2010) report greater pass-through for VAT hikes than for VAT cuts in Hungarian data.

average nonzero price change. This difference was negative for some groups of products, notably services, following the VAT hikes, leading to the surprising implication that the shock was deflationary in nature. And for categories with positive estimates of π^* , we sometimes find that the initial price level response overshoots π^* . Such an overshooting is also at odds with the data because we find no evidence of inflation payback in the months that followed the hikes. The root of these counterintuitive implications is simple: The data strongly reject the absence of a selection effect which is built into the Calvo+ method. Allowing for the exogenous probability of price changes to fluctuate along with macroeconomic conditions is thus insufficient to capture the observed micro price adjustment in response to VAT hikes. We shall return to the importance of the selection effect in section 5.

4.3 The 2008 Collapse in Oil Prices

The bankruptcy of Lehman Brothers on September 15, 2008, sent shock waves throughout the global financial system, and helped transform an initially mild U.S. downturn into the longest and deepest recession of the post-war period. The U.S. price level dropped over two percentage points in the fourth quarter of 2008 alone, the largest quarterly decline on record. Most of that decline can be traced back to the plunge in the price of a barrel of oil from about \$140 to \$40 in just a few months. The upper panel of figure 5 shows that the drop in oil prices was accompanied by a concomitant 50-percent fall in the energy component of the U.S. CPI. The middle row of panels shows the distribution of retail energy price changes in the second half of 2008, along with the corresponding distributions in 2007 for comparison.²⁰ As oil prices fell, the distribution shifted toward price decreases. Some skewness appeared as pass-through to gasoline prices was noticeably larger than for other energy products.

To construct a counterfactual case under full immediate pass-through, we estimate the following equation

$$\pi_t^e = \alpha + \sum_{l=0}^6 \beta_l \Delta Earnings_{t-l} + \sum_{l=0}^6 \delta_l \Delta Wholesale_{t-l} + \varepsilon_t,$$

where π_t^e is the monthly change in the energy component of the CPI, $Earnings_t$ is a measure of labor earnings, and $Wholesale_t$ is the wholesale price of gasoline.²¹ The estimates indicate that a 1-percent increase in the wholesale price of gasoline is associated with a 0.41 percent increase in consumer energy prices after six months, with nearly all pass-through occurring upon impact of the shock or in the month immediately after. We use this six-month response as our measure of $\Delta\pi_t^*$ under the full-pass-through method. The counterfactual frequency of price changes is set to the

²⁰ *Gasoline and gas (piped) and electricity* each account for just under half of our micro data sample of energy prices, with *fuel oil and other household fuels* making up the residual. We first compute the distribution at the monthly frequency then report a simple average over pairs of adjacent months. We do not conduct a similar exercise with Mexican energy prices because they are heavily regulated.

²¹ We use the wholesale price of gasoline instead of WTI to control for movements in refiners' margins, which were volatile in 2008. For labor costs, we use hourly earnings of nonsupervisory employees in the manufacturing sector. The estimation period is January 1987 to December 2009. Aside from the constant, the only statistically significant coefficients are δ_0 and δ_1 .

average frequency in the corresponding month during the 2003-2007 period. Given that the slide in world energy prices occurred somewhat gradually rather than at a specific point in the price collection periods, we implement our method month by month and then average our estimates over two-month periods.

The table at the bottom of figure 5 shows that the cumulative change in consumer energy prices in September and October 2008 accounted for over 90 percent of the estimated shock to individual reset prices over that period, and an even higher fraction in the September-October period. It is clear from these results that the index of macroeconomic flexibility is close to 1 for the transmission of oil shocks to retail energy prices. And with three quarters of retail energy prices being adjusted every month, the full-pass-through method leaves only a minor role to the extensive margin.²² The findings under the Calvo+ method are similar. The frequency of price changes in 2008 was only 2 to 5 percentage points higher than its average for the corresponding months from 2003 to 2007. Shutting down the extensive margin would thus have only a limited impact on the transmission of oil prices.

Outside the energy sector, movements in inflation were rather limited around the Lehman collapse. Despite rapidly deteriorating activity, downward price pressures are hardly perceptible before November and December 2008, and even then price movements were muted.²³ Among unprocessed food categories, pass-through of earlier agricultural commodity price hikes was associated with unusually frequent price *increases* through October, before price drops became more common. Products other than energy and unprocessed food, which together provide a measure of core goods and services, displayed little if any downward drift at the time. We interpret the stability of core prices in the second half of 2008 as consistent with the Great Recession having initially small effects on individual reset prices.

5 Some Evidence of a “Selection Effect”

The contribution of adjustment along the extensive margin to macroeconomic price flexibility should be especially relevant if items with large deviations from their individual reset price are relatively likely to have the timing of their price adjustment altered by the occurrence of a shock. In this final section of results, we first investigate indirect manifestations of the selection effect through the observed distribution of price changes, and then attempt a more direct assessment by studying how a firm’s pricing decisions relate to those of its local competitors.

²²As reported by Eichenbaum *et al.* (2012), prices for natural gas and electricity in the CPI Research Database are based on unit value indexes, which can lead to an overestimation of the frequency of price changes. Imposing a lower frequency of price changes would lower the contribution of the intensive margin.

²³This apparent delay in the inflation response is not due to the timing of price collection. We verified that possibility by constructing distributions of price changes for both monthly and bimonthly observations in the price collection period immediately before and after Lehman’s bankruptcy filing.

5.1 Price Pressure and the Distribution of Price Changes

One consequence of the presence of a selection effect is that the distribution of deviations should remain relatively compact. This observation has testable implications for the shape of the distribution of nonzero price changes. Absent a selection effect, the size of individual price changes should equal, on average, the amount of aggregate reset price inflation having accrued since the last adjustment. Also, absent a selection effect, the dispersion of individual price changes should increase with the duration since the last price change due to the piling up of idiosyncratic shocks. As we now show, we find little empirical support for either phenomenon.

We first use the massive repricing of items associated with the April 1995 Mexican VAT hike to investigate the drift in the mean nonzero price changes conditional on the duration in a high-inflation environment. Although we do not observe aggregate reset price inflation directly, we know that it was substantial at the time because price adjustments were frequent and the CPI rose 40 percent in the 12 months that followed the hike. If aggregate reset price inflation introduces a drift in the average size of price changes, then it should be quite evident in this environment. To investigate this possibility, we consider only items that experienced a price change in April 1995, so that any price pressure from the VAT hike and other aggregate and idiosyncratic sources should have been released that month. We then look at the first price adjustment that followed the hike conditional on the duration. For example, we compute the distribution at a 3-month duration using only items that experienced a price change in April 1995 and had their next price change in July 1995. Such conditioning ensures that items have similar degrees of cumulated aggregate reset price inflation given the timing of their first price adjustment.

As figure 6 shows, the distributions of price changes are remarkably similar across durations. For all special groups of products considered, we find no apparent rise in the mean nonzero price change in the first six months following the hike. For instance, the average price change among nonenergy industrial goods was roughly 10 percent whether items had their first adjustment one month, three months, or six months after the VAT hike. The rise in the price index of nonenergy industrial goods differed much over these horizons, however, at 7 percent after one month and almost 20 percent after six months. Findings are similar for other special groups. A related observation is that the average price change taking place shortly after the VAT hike is larger than the amount of cumulated inflation, whereas the average price change taking place several periods after the VAT hike is smaller than the amount of cumulated inflation. This behavior is consistent with early price changers having initially more price pressure to release than others in the sample due to idiosyncratic factors, and late price changers initially benefiting from an offsetting of aggregate reset price inflation by idiosyncratic shocks, thus limiting the extent of the deviation from their individual reset price.

To gauge the degree of similarity across the distributions at various durations, we compute Kolmogorov-Smirnov statistics to test the hypothesis that the samples of nonzero price changes used to compute the distributions at various durations are drawn from the same population. In the vast majority of cases, we cannot reject at the 5-percent significance level that pairs of the

distributions shown in figure 6 are statistically different.²⁴ One may suspect that the absence of rejection is due to the low power of the Kolmogorov-Smirnov test, especially given that the number of usable observations declines rapidly with the duration considered. However, the test would have easily rejected the null of no difference in nearly all cases if the mean of the distributions had drifted by the amount of cumulated inflation.

Our finding that the mean of the distribution is rather insensitive to the duration since the last price change complements those of Eden (2001), Álvarez and Hernando (2004), and Klenow and Kryvtsov (2008). Our approach is different from that employed by these authors: By conditioning on the adjustment of prices at a particular point in time, we ensure that the amount of cumulated aggregate reset price inflation is similar across items. The extent of aggregate reset price inflation is also much higher than in the analysis of Álvarez and Hernando (2004) and Klenow and Kryvtsov (2008), who consider low-inflation environments, while we benefit from a substantially larger product coverage than Eden (2001).

Low-inflation environments arguably provide conditions in which it is easier to identify any drift in the dispersion of nonzero price changes as a function of duration because they feature a larger number of periods elapsing between price adjustments on average, thus allowing more time for idiosyncratic shocks to accumulate and create dispersion. Figure 7 presents the distributions in the U.S. CPI for four special groups of products. We find no obvious increase in the dispersion of price changes as a function of duration. If anything, some groups of products display a small decrease in the dispersion of price changes conditional on the duration. Arguably, our distributions do not control for heterogeneity in price-setting practices, which could imply, for example, that items with relatively elevated idiosyncratic noise choose to reprice more often than those facing less volatility. Nevertheless, the findings are consistent with that reported by Campbell and Eden (2010) and Klenow and Kryvtsov (2008), who control for such heterogeneity.

Another indirect sign of the selection effect being at play would be if the distribution of price changes had little mass near zero, consistent with firms waiting for the deviation from their individual reset price to be large enough to justify adjusting their posted price. Small price changes are quite common in CPI data overall but less so in some specific markets or for particular firms. For example, Cavallo and Rigobon (2011) find that the distribution of price changes is often bimodal in scrapped on-line supermarket data from 22 countries. We are reluctant to attribute this bimodality entirely to the selection effect, however, because it could be an artefact of the price-setting strategies employed by retailers. In particular, Boivin *et al.* (2012) show that some online book retailers commonly post large positive and negative price changes, whereas others mostly adjust by small amounts. A related finding reported by Klenow and Kryvtsov (2008) is that the average size of price changes in the U.S. CPI is noticeably larger for sales-related prices (25.1 percent) than for regular prices (11.3 percent).

²⁴The tests assume that observations used in the construction of the distributions are equally weighted, whereas our item weights are based on their spending share in the sample. Given that our objective is largely illustrative, we simply used the total (unweighted) number of observations in the computation of the test statistics.

5.2 Local Competition and the Adjustment Hazard

We next investigate how an item’s price adjustment probability relates to the deviation from its individual reset price. An immediate difficulty in pursuing this objective is that the deviation is typically not observed. There are two approaches in the literature for dealing with this hurdle. The first approach is to parametrize the shape of the adjustment hazard to target moments of some aggregate variables of interest (e.g., Caballero and Engel [1993a and 1993b]). It has the benefit of not requiring the identification of individual reset prices but comes at the cost of ignoring rich micro data information. The second approach, which we follow below, is to proxy the deviation using information from observables. This approach has notably been employed to study car purchase decisions by Eberly (1994) and Attanasio (2000), and hiring decisions by Caballero, Engel, and Haltiwanger (1997).

Applications of this latter approach to the study of pricing decisions have been constrained by the lack of micro data that can be used to proxy individual reset prices. We note two exceptions. Using weekly scanner data from a large U.S. retailer, Eichenbaum *et al.* (2011) show that the probability of observing a change in an item’s “reference” price (defined as the modal price in the quarter) is increasing in the deviation from the average mark-up over the vendor cost (measured at the UPC-store level). Their estimates imply a potent selection effect: a 5-percentage-point increase in the deviation from the average markup raises the probability of adjusting the reference price during a given week by over 10 percentage points. A related finding is that deviations from the average markup are small, with 90 percent of items having an absolute deviation of 10 percentage points or less.

The other notable exception is Campbell and Eden (2010), who use weekly scanner data from two small Midwestern cities. They report that the probability of observing a price change is increasing in the deviation between an item’s price and the average price of its local competitors selling the same UPC. A rationale for their approach is that individual reset prices, when measured at the UPC-market level, may comove strongly due to similarities in costs and the degree of local competition. Indeed, we will provide evidence of limited dispersion in the level of prices across stores. There are also limitations to this approach. Deviations from the average price of local competitors abstract from idiosyncratic factors influencing individual reset prices. And even if firms shared the same individual reset price, deviations from the average price may imperfectly proxy for the actual amount of price pressure because the average price only gradually reflect common shocks due to nominal rigidities.

5.2.1 Proxying for the Deviation

We use the IRI Marketing dataset to investigate pricing decisions and local competition. As appendix A details, this dataset contains UPC-level information on weekly sales and prices of U.S. grocery stores and drug stores. It allows us to control for the geographic location and to identify stores belonging to the same retail chain. The sample size is unusually large at about 300 billion individual price observations per year. The dataset covers staple food (e.g., carbonated beverages,

condiments, and cereals) and personal care products (e.g., toilet paper and laundry detergent).

To construct an estimate of the deviation, we first calculate the average price of store i 's local competitors in month t ,

$$\bar{p}_{-i,t} = \sum_{s \neq i} \omega_{-i,s,t} p_{s,t}.$$

The weight $\omega_{-i,s,t}$ is the share of total UPC sales by item i 's local competitors that is accounted for by store s . We then calculate retailer i 's deviation from $\bar{p}_{-i,t}$ prevailing at the beginning of the period, $\tilde{x}_{i,t} = p_{i,t-1} - \bar{p}_{-i,t}$. There could be factors leading to permanent differences in the level of prices across stores, even conditioning on the UPC and market. For example, the Chicago-area grocery chain Dominick's Finer Foods uses a four-tier system to set the level of prices across its stores in order to account for differences in the degree of local competition and consumer price sensitivity. For this reason, we use demeaned time series unless otherwise indicated.

As section 3.1 discussed, only price adjustments coinciding with a reoptimization should be comprised in the intensive and the extensive margins. Accordingly, in addition to using posted prices, we consider a "regular" price series excluding one-month temporary price drops and a "reference" price series computed as the modal posted price observed during a 13-week period centered around the 15th day of the month. Our objective is not to dismiss sales and nonreference prices as channels for the transmission of shocks. On the contrary, there is mounting evidence that these prices are responsive to aggregate conditions, although the extent of this response remains the subject of ongoing debate.²⁵ Instead, we use these alternative price series to provide statistics under a wider range of assumptions regarding the set of price adjustments coinciding with a reoptimization. Appendix A details the construction of these series.

5.2.2 Local Price Dispersion and the Nonparametric Adjustment Hazard

We first use $\tilde{x}_{i,t}$ to document the importance of dispersion in the level of prices across stores at the UPC-market level.²⁶ We sort $\tilde{x}_{i,t}$ from all UPCs and markets with sufficient observations into 1-percentage-point bins, pooling together observations from all months. Observations within a UPC-market combination are weighted uniformly to highlight price dispersion across stores. The distributions also weigh equally all UPC-market combinations. The upper panel of figure 8 shows the resulting distribution of deviations before demeaning $\tilde{x}_{i,t}$. The distributions are clearly skewed to the right: about 10 percent of observations for posted prices and regular prices have $\tilde{x}_{i,t} < 0$, a proportion that doubles to 20 percent for reference prices. The explanation for this skewness is that total UPC sales within local markets tend to be dominated by a small number of stores with relatively low prices. The lower panel of figure 8 reports the distributions after demeaning $\tilde{x}_{i,t}$. These distributions are relatively tight, which suggests that fluctuations in the level of prices are

²⁵See Klenow and Malin (2011) for a recent survey and several new facts, and Chevalier and Kashyap (2011) for evidence that multi-product retailers may adjust the frequency and intensity of sales in response to aggregate shocks.

²⁶Consistent with the CE decomposition, $\tilde{x}_{i,t}$ measures the deviation *before* an item's price adjustment decision, whereas previous studies of price dispersion (e.g., Reinsdorf 1994) focus of the deviation *after* price adjustments take place. The shape and positive skewness of the distribution are little changed when we adopt this alternative timing.

somewhat limited once one accounts for permanent differences across stores.

We note at the outset that some of the observed price dispersion is not due to nominal rigidities hampering price adjustment in response to shocks, but is rather traced back to firms' own pricing decisions. The upper panel of figure 9 illustrates this fact by showing how prices at the edges of the distribution tend to be relatively young. The average age of posted, regular, and reference prices peaks near 5 months, 6 months, and 10 months, respectively, for $\tilde{x}_{i,t}$ near zero. The average age then falls steadily as we consider increasingly large absolute deviations from the average price of local competitors. In the case of posted and regular prices, the average age of items with an absolute deviation around 20 percent is roughly 40 percent smaller than that of items with a deviation near zero. The relative age difference is even larger for reference prices. These findings are consistent with those of Campbell and Eden (2010) and caution against associating all forms of price dispersion with resource misallocation.

The middle panel of figure 9 plots the nonparametric adjustment hazard, $\Lambda^{np}(\tilde{x}_{i,t})$, computed for each 1-percentage point bin. Unsurprisingly, $\Lambda^{np}(\tilde{x}_{i,t})$ is highest for posted prices, which are relatively flexible, and lowest for reference prices, which are relatively sticky. Consistent with a selection effect, $\Lambda^{np}(\tilde{x}_{i,t})$ is increasing in the absolute deviation from the average price of local competitors. One may suspect that this phenomenon is driven by temporary sales and promotions, which can cause posted prices to be temporarily located on the edges of the distribution of deviations. However, the V shape is preserved when we exclude sales and nonreference price changes. For instance, the probability of observing a reference price change leaps from 7 percent when the item's price equals the average of its local competitors to between 20 percent and 30 percent when its (absolute) deviation reaches 20 percent.

The bottom panel of figure 9 shows that the median firm adjusts its reference price by the full size of its price gap vis-à-vis local competitors (i.e., $\Delta p_{i,t} \approx -\tilde{x}_{i,t}$). A similar conclusion holds for posted prices and regular prices (not shown). That said, there is some heterogeneity in the size of price changes conditional on $\tilde{x}_{i,t}$, as evidenced by the 10th and 90th percentile lines reported in the panel. This variation suggests that, while $\tilde{x}_{i,t}$ is a roughly unbiased estimate of the actual deviation, $x_{i,t}$, it misses other factors influencing individual reset prices. Eichenbaum *et al.* (2011) also find that deviations from their average markup over vendor costs are about erased on average

Figures 10 and 11 provide some evidence that the shape of $\Lambda^{np}(\tilde{x}_{i,t})$ is robust to alterations to our baseline methodology. First, it retains its V shape when we compute it for narrower product categories, as illustrated in figure 10 for salty snacks, carbonated beverages, cold cereal, and frozen dinners. Second, we allow for the possibility that pricing decisions are made at the chain level rather than at the store level by excluding stores belonging to the same chain as retailer i from the computation of $\bar{p}_{-i,t}$. The upper panel of figure 11 shows that the resulting adjustment hazard is only a bit flatter than our benchmark case for reference prices. The findings are also similar for posted price and regular prices when we control for chains (not shown). Third, the figure also presents the nonparametric hazard obtained when the deviation is constructed using only the average price of the local competitors that have reset it over the previous six month (inclusive). In

this case, $\Lambda^{np}(\tilde{x}_{i,t})$ is a bit steeper than in our benchmark case. Finally, the lower panel of figure 11 reports $\Lambda^{np}(\tilde{x}_{i,t})$ conditional on the age of price deviations. While the V shape is preserved at all durations considered, the level of the estimated adjustment hazard declines with the age of the price, consistent with negative duration dependence in the data.

5.2.3 Parametric Estimation of the Adjustment Hazard

The nonparametric nature of our analysis so far has precluded us from controlling for various factors that could affect the shape of the adjustment hazard function, such as seasonal repricing and selection effects by which items with large deviations are drawn from UPC-market combinations with relatively flexible prices. To assess whether these elements impacted our findings, we estimate linear probability models with item fixed effects in the spirit of Campbell and Eden (2010). We focus on reference prices to limit the risk that sales, promotions, and other transitory price movements orient us toward finding a V-shaped adjustment hazard. The presence of item fixed effects also seeks to limit that possibility. Our specification is

$$I(\Delta p_{i,t} \neq 0) = \beta_{i,0} + \beta_1 \tilde{x}_{i,t}^{\text{neg}} + \beta_2 \tilde{x}_{i,t}^{\text{pos}} + Z'_{i,t} \Gamma + \varepsilon_{i,t},$$

where $I(\Delta p_{i,t} \neq 0)$ is an indicator variable that item i 's price has changed, $\tilde{x}_{i,t}^{\text{neg}}$ is the $\tilde{x}_{i,t}$ if $\tilde{x}_{i,t} < 0$ and zero otherwise, $\tilde{x}_{i,t}^{\text{pos}}$ is similarly defined for positive deviations, and $Z_{i,t}$ is a vector of control variables that includes period dummies and a measure of store size (total yearly revenues). $\tilde{x}_{i,t}^{\text{neg}}$ and $\tilde{x}_{i,t}^{\text{pos}}$ allow the adjustment hazard to have different slopes for negative and positive deviations. We also considered higher-order terms but their economic significance was minimal. We cluster the standard errors at the store level and, given the expansive size of the IRI Marketing dataset, run the regressions at the product category-market level. We only include in our sample items from UPC-market combinations that have at least 10 observations.

The first three columns of table 4 display the estimated coefficients for three product category-market combinations: salty snacks in New York City, carbonated beverages in Los Angeles, and frozen dinners in Boston. All slope coefficients are statistically significant and increasing in the (absolute) size of positive and negative deviations, consistent with our benchmark findings. In addition, the magnitudes of the slopes are comparable to the ones obtained in our nonparametric exercise. As was the case earlier, the results point to a steeper adjustment hazard for negative deviations than positive ones, suggesting that relatively low prices are especially unlikely to persist for long.

For salty snacks in New York City, our largest UPC-market combination, we also ran separate regressions for price increases and price decreases (i.e., used $I(\Delta p_{i,t} > 0)$ and $I(\Delta p_{i,t} \neq 0)$ as the dependent variable), which are reported in columns (4) and (5) of table 5. Comparing these coefficients, we find that the magnitude of the negative deviation ($\tilde{x}_{i,t} < 0$) has a larger impact on the probability of a price increase than a price decrease, and vice versa for positive deviations. And to assess the effect of the age of the price on the estimates, we ran our regression again retaining

only observations at specific durations. The coefficients, shown in columns (6) to (8) of table 5 for durations of 2 months, 6months, and twelve months, respectively, suggest that the shape of the adjustment hazard is not driven by these young prices.

5.2.4 Local Competition and the Index of Macroeconomic Flexibility

Our benchmark nonparametric hazard has proven robust to controlling for several elements unrelated to the selection effect that could influence its slope. Consistent with the limited amount of price dispersion across stores shown in figure 8, we see our estimation results as evidence that the timing of pricing decisions is sensitive to deviations from individual reset prices, and that adjustment along the extensive margin may play a role at the macroeconomic level. If we assume that the deviation from the price of local competitors ($\tilde{x}_{i,t}$) and nonparametric adjustment hazard ($\Lambda^{np}(\tilde{x}_{i,t})$) are good proxies for the deviation from the actual individual reset price ($x_{i,t}$) and adjustment hazard function ($\Lambda(x)$), respectively, then we have in hand all the elements needed to estimate the intensive and the extensive margins in equation 4 directly on the micro data. Specifically, the intensive margin can be obtained by integrating $\Lambda^{np}(\tilde{x}_{i,t})$ over the distribution of $\tilde{x}_{i,t}$, and the extensive margin can be recovered by integrating the product of $\tilde{x}_{i,t}$ and the derivative of $\Lambda^{np}(\tilde{x}_{i,t})$ over the distribution of $\tilde{x}_{i,t}$.

Table 6 shows the estimated intensive and extensive margins when we perform such exercise. For the slope of the nonparametric hazard function, we recover the slope of $\Lambda^{np}(\tilde{x}_{i,t})$ separately for positive and negative deviations using linear regressions. We find that the extensive margin is an economically important contributor to the index of macroeconomic flexibility irrespective of whether we use posted, regular, or reference prices. Depending on the particular price measure used, the extensive margin is 35 percent to 45 percent as large as the intensive margin. Absent real rigidities, the index of macroeconomic flexibility ranges from 0.13 for reference prices to 0.37 for posted prices. If we instead assume moderate real rigidities similar to those estimated on retail data by Burstein and Hellwig (2007), then the point estimates of the index of macroeconomic flexibility roughly halve. Although these results imply greater macro price stickiness, the extensive margin is nevertheless an important contributor to F^{macro} .

There is admittedly much uncertainty surrounding these estimates. Importantly, they rely on an imperfect measure of the true deviation from individual reset prices. Whether this measure leads us to underestimate or overestimate the slope of the adjustment hazard is unclear. On the one hand, by omitting idiosyncratic factors, we add noise to the deviation that could bias downward the slope of the adjustment hazard. Indeed, the probability of price changes reported by Eichenbaum *et al.* (2011) appears more responsive to their measure of the deviation than ours. The importance of the extensive margin in table 6 is also near or below the lower bounds derived in section 4 in the case of adjustment to large shocks. On the other hand, sluggish adjustment of the price of local competitors to shocks could bias upward the slope of the nonparametric adjustment hazard function. For these reasons, we see these estimates as suggestive that the extensive margin is an important channel for macro price adjustment, but more work remains to be done to refine these

estimates and broaden the set of products to which they apply.

6 Conclusion

In this paper, we have sought to translate micro price stickiness into macro price stickiness by distinguishing between price changes that are determined ahead of shocks and price changes that are either triggered or cancelled in response to shocks. We have discussed the assumptions under which observed price behavior can be used to infer the importance of each margin of price adjustments. We have found compelling evidence that some shocks alter the timing of price changes, contributing significantly to the flexibility of the aggregate price index. In the context of the CE decomposition used throughout the paper, such state-dependent price adjustments are described as operating along the extensive margin.

Our findings have several implications for macroeconomic modelling. First, researchers primarily interested in fitting macro facts should be careful when claiming that their models are consistent with micro facts. The current generation of macro models is still a long way from capturing the richness of pricing strategies employed by firms; some simplification is unavoidable. As emphasized by CE, models featuring no extensive margin abstract from an important channel of adjustment and their calibration should take that aspect into account. Our findings also suggest that micro-level price stickiness is at most a moderate contributor to the transmission mechanism. Moreover, rapid initial pass-through of the shocks covered in our study seems mostly incompatible with the presence of strong real rigidities, a conclusion echoing that of Bils and Klenow (2004) and BKM based on the dynamics of sectoral U.S. CPI inflation series.

Obviously, our work leaves several questions unanswered. We have considered only a subset of sectoral and aggregate shocks with the peculiarity of being relatively easy to identify. Micro price adjustment to other types of aggregate shocks —notably monetary shocks— could be more sluggish. Also, there continues to be a lack of microeconomic data on both costs and local competition, which could help refine estimates of individual reset prices between price adjustment periods. Moreover, micro-level price evidence on the importance of real rigidities, which is important for refining estimates of adjustment along the extensive margin, remains constrained by limited data availability.

References

- [1] Alvarez, Fernando, Francesco Lippi, and Luigi Paciello. 2011. “Optimal Price Setting with Observation and Menu Costs.” *Quarterly Journal of Economics* 126, no.4:1909-1960.
- [2] Álvarez, Luis J. and Ignacio Hernando. 2004. “Price Setting Behaviour in Spain: Stylised Facts using Consumer Price Micro Data.” Working Paper Series 416, European Central Bank.
- [3] Attanasio, Orazio P. 2000. “Consumer Durables and Inertial Behaviour: Estimation and Aggregation of (S,s) Rules for Automobile Purchases.” *Review of Economic Studies* 67, no.4:667-96.

- [4] Baharad, Eyal and Benjamin Eden. 2004. “Price Rigidity and Price Dispersion: Evidence from Micro Data.” *Review of Economic Dynamics* 7, no.3:613-641.
- [5] Ball, Laurence M., and David Romer. 1990. “Real Rigidities and the Non-neutrality of Money.” *Review of Economic Studies* 57, no.2:183–203.
- [6] Barsky, Robert B., Christopher L. House, and Miles S. Kimball. 2007. “Sticky-Price Models and Durable Goods.” *American Economic Review* 97, no. 3:984-998.
- [7] Barro, Robert J. 1972. “A Theory of Monopolistic Price Adjustment.” *Review of Economic Studies* 39, no.1:17-26.
- [8] Berger, David and Joseph Vavra. 2011. “Dynamics of the U.S. Price Distribution.” Yale University.
- [9] Bills, Mark and Peter J. Klenow. 2004. “Some Evidence on the Importance of Sticky Prices.” *Journal of Political Economy* 112, no.5:947-985.
- [10] Bills, Mark, Peter J. Klenow, and Benjamin A. Malin. 2009. “Reset Price Inflation and the Impact of Monetary Policy Shocks.” NBER Working Paper 14787.
- [11] Bills, Mark, Peter J. Klenow, and Benjamin A. Malin. (forthcoming). “Reset Price Inflation and the Impact of Monetary Policy Shocks,” *American Economic Review*.
- [12] Blanchard, Oliver J. and S. Fisher. 1989. *Lectures on Macroeconomics*. Cambridge: MIT Press.
- [13] Blinder, Alan S. 1994. “On Sticky Prices: Academic Theories Meet the Real World.” In *Monetary Policy*, ed. N. Gregory Mankiw, 117-154. University of Chicago and NBER.
- [14] Boivin, Jean, Robert Clark, and Nicolas Vincent. 2012. “Virtual Borders.” *Journal of International Economics*, 86, No.:2012: 327-335.
- [15] Bonomo, Marco and Carlos Carvalho. 2004. “Endogenous Time-Dependent Rules and Inflation Inertia.” *Journal of Money, Credit and Banking* 36, no.6:1015-41.
- [16] Bonomo, Marco, Carlos Carvalho, and René Garcia. 2010. “State-dependent Pricing Under Infrequent Information: A Unified Framework.” Staff Reports 455, Federal Reserve Bank of New York.
- [17] Bonomo, Marco and René Garcia. 2001. “The Macroeconomic Effects of Infrequent Information with Adjustment Costs.” *Canadian Journal of Economics* 34, no.1:18-35.
- [18] Bronnenberg, Bart J., Michael W. Kruger, and Carl F. Mela. 2008. “Database paper: The IRI Marketing Data Set.” *Marketing Science* 27, no.4:745–748.
- [19] Bureau of Labor Statistics (2011), *Handbook of Methods*, Bureau of Labor Statistics, Washington, DC.

- [20] Burstein, Ariel T. 2006. "Inflation and Output Dynamics with State-Dependent Pricing Decisions." *Journal of Monetary Economics* 53, no.7:1235-1257.
- [21] Burstein, Ariel, Martin Eichenbaum, and Sergio Rebelo. 2005. "Large Devaluations and the Real Exchange Rate." *Journal of Political Economy* 113, no.4:742-784.
- [22] Burstein, Ariel and Christian Hellwig. 2007. "Prices and Market Shares in a Menu Cost Model." NBER Working Papers 13455, National Bureau of Economic Research.
- [23] Caballero, Ricardo J. and Eduardo M.R.A. Engel. 1993a. "Microeconomic Adjustment Hazards and Aggregate Dynamics." *Quarterly Journal of Economics* 108, no.2:359-383.
- [24] Caballero, Ricardo J. and Eduardo M. R. A. Engel. 1993b. "Microeconomic Rigidities and Aggregate Price Dynamics." *European Economic Review* 37, no.4:697-711.
- [25] Caballero, Ricardo J. and Eduardo M. R. A. Engel. 1999. "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach." *Econometrica* 67, no.4:783-826.
- [26] Caballero, Ricardo J. and Eduardo M.R.A. Engel. 2007. "Price Stickiness in Ss Models: New Interpretations of Old Results." *Journal of Monetary Economics* 54, (Supplement):100-121.
- [27] Caballero, Ricardo J., Eduardo M.R.A. Engel, and John Haltiwanger. 1997. "Aggregate Employment Dynamics: Building from Microeconomic Evidence." *American Economic Review* 87, no.1:115-137.
- [28] Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12, no.3:383-398.
- [29] Campbell, Jeffrey R. and Benjamin Eden. 2010. "Rigid Prices: Evidence from U.S. Scanner Data." Working Paper Series WP-05-08, Federal Reserve Bank of Chicago.
- [30] Caplin, Andrew S. and Daniel F. Spulber. 1987. "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics* 102, no.4:703-25.
- [31] Capistrán, Carlos and Gabriel López-Moctezuma. 2010. "Forecast Revisions of Mexican Inflation and GDP Growth." Working Paper 2010-11, Banco de México.
- [32] Carlton, Dennis W. 1986. "The Rigidity of Prices." *American Economic Review* 76, no.4:637-58.
- [33] Carlton, Dennis W. 1989. "The Theory and the Facts of How Markets Clear: Is Industrial Organization Valuable for Understanding Macroeconomics?" *Handbook of Industrial Organization*, ed. by R. Schmalensee & R. Willig, 1 ch.15:909-946.
- [34] Cavallo, Alberto and Roberto Rigobon. 2011. "The Distribution of the Size of Price Changes." NBER Working Paper 16760.

- [35] Cecchetti, Stephen G. 1986. "The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines." *Journal of Econometrics* 31, no. 3:255-274.
- [36] Chahrour, Ryan A. 2011. "Sales and Price Spikes in Retail Scanner Data." *Economics Letters* 110, no.2:143-146.
- [37] Chevalier, Judith A. and Anil K. Kashyap. 2011. "Best prices." NBER Working Paper 16680.
- [38] Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113, no.1:1-45.
- [39] Costain, James and Anton Nakov. 2011. "Distributional Dynamics under Smoothly State-Dependent Pricing." *Journal of Monetary Economics* 58, no.6-8:646-665.
- [40] Dhyne, Emmanuel, Luis J. Álvarez, Hervé Le Bihan, Giovanni Veronese, Daniel Dias, Johannes Hoffmann, Nicole Jonker, Patrick Lünemann, Fabio Rünstler, and Jouko Vilmunen, 2005. "Price Setting in the Euro Area: Some Stylized Facts from Individual Consumer Price Data," Working Paper 524, European Central Bank.
- [41] Dotsey, Michael and Robert G. King. 2005. "Implications of State-dependent Pricing for Dynamic Macroeconomic Models." *Journal of Monetary Economics* 52, no.1:213-242.
- [42] Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo. 2011. "Reference Prices and Nominal Rigidities." *American Economic Review* 101, no.1:234-262.
- [43] Eichenbaum, Martin, Nir Jaimovich, Sergio Rebelo, and Josephine Smith. 2012. "How Frequent Are Small Price Changes?" mimeo, Northwestern University.
- [44] Eberly, Janice C. 1994. "Adjustment of Consumers' Durables Stocks: Evidence from Automobile Purchases." *Journal of Political Economy* 102, no.1:403-36.
- [45] Eden, Benjamin. 1994. "The Adjustment of Prices to Monetary Shocks When Trade Is Uncertain and Sequential." *Journal of Political Economy* 102, no.3:493-509.
- [46] Eden, Benjamin. 2001. "Inflation and Price Adjustment: An Analysis of Microdata." *Review of Economic Dynamics* 4, no.3:607-636.
- [47] Fabiani, Silvia, Martine Druant, Ignacio Hernando, Claudia Kwapil, Bettina Landau, Claire Loupias, Fernando Martins, Thomas Mathä, Roberto Sabbatini, Harald Stahl, and Ad Stokman. 2006. "What Firms' Surveys Tell Us about Price-Setting Behavior in the Euro Area." *International Journal of Central Banking*, September:3-48.
- [48] Gagnon, Etienne. 2009. "Price Setting during Low and High Inflation: Evidence from Mexico." *Quarterly Journal of Economics* 124, no.4:1221-1263.

- [49] Gertler, Mark and John Leahy. 2008. "A Phillips Curve with an Ss Foundation." *Journal of Political Economy* 116, no.3:533-572.
- [50] Golosov, Mikhail and Robert E. Lucas Jr. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy*, 115:171-199.
- [51] Gopinath, Gita and Oleg Itskhoki. 2010. "Frequency of Price Adjustment and Pass-Through." *Quarterly Journal of Economics* 125, no.2:675-727.
- [52] Gopinath, Gita and Oleg Itskhoki. 2011. "In Search of Real Rigidities." NBER Chapters, *NBER Macroeconomics Annual 2010* 25:261-309.
- [53] Gorodnichenko, Yuriy. 2010. "Endogenous Information, Menu Costs and Inflation Persistence." mimeo, University of California - Berkeley.
- [54] Head, Allen, Lucy Qian Liu, Guido Menzio, and Randall Wright. (forthcoming). "Sticky Prices: a New Monetarist Approach." *Journal of the European Economic Association*.
- [55] Kara, Engin. 2011. "Understand and Modelling Reset Price Inflation." Discussion Paper 11.623. University of Bristol.
- [56] Karadi, Peter and Adam Reiff. 2010. "Inflation Asymmetry, Menu Costs and Aggregation Bias – A Further Case for State Dependent Pricing." Working Papers 20103, Magyar Nemzeti Bank.
- [57] Kashyap, Anil K. 1995. "Sticky Prices: New Evidence from Retail Catalogs." *Quarterly Journal of Economics* 110, no.1:245-74.
- [58] Klenow, Peter J. and Oleksiy Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 123, no.3:863-904.
- [59] Klenow, Peter J. and Benjamin A. Malin. 2010. "Microeconomic Evidence on Price-Setting." *Handbook of Monetary Economics*, Volume 3, Chapter 6, in B. M. Friedman and M. Woodford (editors), pages 231-284.
- [60] Klenow, Peter J. and Jonathan L. Willis. 2006. "Real Rigidities and Nominal Price Changes." Research Working Paper RWP 06-03, Federal Reserve Bank of Kansas City.
- [61] Lach, Saul and Daniel Tsiddon. 1992. "The Behavior of Prices and Inflation: An Empirical Analysis of Disaggregated Price Data." *Journal of Political Economy* 100, no.2:349-89.
- [62] Levy, Daniel, Dongwon Lee, Haipeng Chen, Robert J. Kauffman, and Mark Bergen. 2011. "Price Points and Price Rigidity." *Review of Economics and Statistics* 93, no.4:1417-1431.
- [63] Mackowiak, Bartosz and Frank Smets. 2008. "On Implications of Micro Price Data for Macro Models." Working Paper 960, European Central Bank.

- [64] Mackowiak, Bartosz and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review* 99, no.3:769-803.
- [65] Mankiw, N. Gregory and Ricardo Reis. 2002. "Sticky Information Versus Sticky Prices: A Proposal To Replace The New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117, no.4:1295-1328.
- [66] Matejka, Filip and Christopher A. Sims. 2011. "Discrete Actions in Information-Constrained Tracking Problems." CERGE-EI Working Papers wp441, The Center for Economic Research and Graduate Education - Economic Institute, Prague.
- [67] Midrigan, Virgiliu. 2011. "Menu-Costs, Multi-Product Firms and Aggregate Fluctuations." *Econometrica* 79, no.4:1139-1180.
- [68] Nakamura, Emi and Jón Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *The Quarterly Journal of Economics* 123, no.4:1415-1464.
- [69] Reinsdorf, Marshall. 1994. "New Evidence on the Relation between Inflation and Price Dispersion." *American Economic Review* 84, no.3:720-731.
- [70] Rotemberg, Julio J. 2004. "Fair Pricing." NBER Working Papers 10915, National Bureau of Economic Research.
- [71] Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50, no.3:665-690.
- [72] Sheshinski, Eytan and Yoram Weiss. 1977. "Inflation and Costs of Price Adjustment." *Review of Economic Studies* 44, no.2:287-303.
- [73] Smets, Frank and Rafael Wouters, 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review* 97, no.3:586-606.
- [74] Taylor, John B. 1980. "Aggregate Dynamics and Staggered Contracts." *Journal of Political Economy* 88, no.1:1-23.
- [75] Woodford, Michael. 2003a. *Interest and Prices*. Princeton, NJ, Princeton University Press.
- [76] Woodford, Michael. 2003b. "Imperfect Common Knowledge and the Effects of Monetary Policy." In P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, eds. *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*. Princeton, NJ, Princeton University Press.
- [77] Woodford, Michael. 2008. "Inattention as a Source of Randomized Discrete Adjustment." mimeo, Columbia University.
- [78] Woodford, Michael. 2009. "Information-Constrained State-Dependent Pricing." *Journal of Monetary Economics* 56, (Supplement):S100-S124.

- [79] Wulfsberg, Fredrik. 2009. "Price Adjustments and Inflation - Evidence from Norwegian Consumer Price Data 1975-2004." Working Paper 2009/11, Norges Bank.

Appendix A. Data Description

We use data from three sources to study a variety of macroeconomic shocks and price-setting behavior. All three datasets have been described in earlier work, so we shall concentrate on presenting the features that are most relevant for our paper.

U.S. CPI Database

The Bureau of Labor Statistics (BLS) makes available most of the micro data behind the official U.S. CPI on a restricted-access basis through its CPI-Research Database. This database excludes housing rents, whose index is computed separately due to differences in sampling and treatment. With the exceptions of food, energy, and a few additional product categories, item prices are collected every other month in all metropolitan areas but the largest three (New York City, Los Angeles, and Chicago), where price collection takes place monthly for all items. The database starts in January 1988 and we extracted data through December 2009. We restrict the sample to the largest three metropolitan areas, leaving us with between 12,000 and 19,000 observations per month. The individual observations are weighted using the same procedure and product category weights as BKM.²⁷ We classified the product categories according to the United Nations *Classification of Individual Consumption According to Purpose* (COICOP) in order to compute statistics for special groups of products. For additional details on this dataset, see the BLS Handbook of Methods, Bils and Klenow (2004), Klenow and Kryvtsov (2008), and Nakamura and Steinsson (2008).

Mexican CPI Database

The Mexican CPI dataset covering the period January 1994 to December 2004 was constructed by Gagnon (2009) based on monthly price lists and item substitution lists published by the Bank of Mexico. We extended the sample period back to July 1993 and forward to December 2010 using price lists released in the official gazette of the Mexican government. The 1993 price lists were not available on electronic support; the data were typed in by a data entry firm using scans of the original documents obtained from Mexico's Archivo General de la Nación. The dataset excludes items whose prices are regulated, such as taxi fares, phone services, gasoline, and tuition fees. The sample covers 54 percent of Mexican consumer price expenditures prior to the basket update of July 2002, and about 61 thereafter. The number of usable monthly observations ranges from 21,000 in 1994 to 57,000 in 2010. Prices are inclusive of sales as long as these sales are conditional on the purchase of a single item. There is no flag in the dataset indicating that an item is on sale but the application of filters to remove V-shaped price patterns suggests that sales are less prominent than in the U.S. CPI. The reader is referred to Gagnon (2009) for further details on the database construction.

SymphonyIRI Group Scanner Database

The IRI Marketing dataset includes scanner data from U.S. grocery stores and drugstores. Its content is detailed in Bronnenberg, Kruger, and Mela (2008). An observation corresponds to a specific item sold by a given store in a particular week. Available information includes total revenue

²⁷We thank them for sharing these weights with us.

from the sale of the item and the number of units sold. As is customary with scanner data, the unit price is obtained by dividing total revenue by the quantity sold.

The dataset has several appealing features. First, observations are available at a high collection frequency. Second, the dataset covers a relatively long time span, January 2001 to December 2007, that allows us to look at price dynamics. Third, information is available for 31 food and personal care product categories, such as beer, household cleaners, and milk. Fourth, the geographical coverage—50 U.S. markets—is more extensive than that of other scanner datasets used in the literature. Also, it includes several retailers within each local market. The number of stores ranges from 7 in Eau Claire, WI, and Pittsfield, MA, to 97 in New York City, NY. To study pricing decisions in competitive environments, we restrict the sample to Boston, Chicago, Houston, Los Angeles, New York, Philadelphia, San Diego, and Washington D.C.. These markets have at least 30 stores each, which limits the risk that most observations belong to the same chains. The high frequency of price collection, long time span, and wide product, market, and store coverage mean that the number of usable observations is enormous at about 300 billion per year. One limitation of the dataset is that it does not contain information about costs.

We applied a number of filters to the data to make them suitable for our analysis. First, we censored fractional price observations. For various reasons (use of membership cards, price changes during the week, etc.), dividing total revenue by the number of units sold occasionally yields a price with fractional cents (e.g., \$3.7485). For New York City, fractional prices represent less than 0.2 percent of posted prices, and a negligible proportion of reference prices; censoring fractional prices is thus mostly inconsequential for our results. Second, we occasionally filtered out temporary sales, which we define as a price drop that is offset within two weeks by a price increase of equal magnitude. Third, we consider a slight variant of the reference price filter of Eichenbaum *et al.* (2011). Specifically, the reference price of a given item in week t corresponds to the mode of the prices observed between week $t - k$ and $t + k$. For robustness, we considered a 13-week centered window by setting $k = 6$. Finally, we converted weekly statistics to a monthly frequency by selecting weekly observations encompassing the 15th day of every month.

Appendix B. Adjustment along the Intensive Margin

We briefly review the empirical evidence on the average frequency of price changes in the U.S. CPI, and offer matching statistics for the IRI Marketing scanner database. The average frequency plays a central role in our analysis. This frequency is equal to the intensive margin and, provided real rigidities are not too strong, offering a lower bound on the index of macroeconomic flexibility. For the revelation principle to apply, only price adjustments coinciding with a reoptimization should be comprised in the intensive margin. For this reason, we report alternative measures of the average frequency that exclude sales and nonreference prices changes. These alternative measures provide a range of estimates for the intensive margin under widely used assumptions regarding the set of price adjustments coinciding with a reoptimization.

The top of table 1 reports the average frequency of price changes in the CPI excluding shelter in the three largest U.S. metropolitan areas. According to Klenow and Kryvtsov (2008), items representing 36.2 percent of consumer expenditures experience a price adjustment every month. This statistic suggests that, absent strong real rigidities, nominal shocks are initially passed through rapidly to individual reset prices even if no adjustment along the extensive margin were to take place. The average frequency of price changes slides to a still-elevated 29.9 percent when one excludes price changes flagged by the BLS as related to sales and promotions. Filtering out nonreference price changes is more consequential. Klenow and Malin (2010) find an average frequency of 14.6 percent when they define the reference price as the modal price over a 13-month window centered on the current month, implying that reference prices are considerably stickier than regular prices. Whether one should ignore all nonreference price changes is debatable, however. Klenow and Malin show that price changes exhibit considerable novelty: Every month, the price of a quarter of items in the U.S. CPI is adjusted to a level not seen over the previous 12 months. Moreover, changes in nonreference prices are an important contributor to short-run inflation dynamics. Also, some frequently occurring prices could be selected by firms because consumers find them more attractive than others (e.g., \$9.99 versus \$10.47), in which case some of the stickiness in reference prices could be due to features of consumer demand rather than rigidity in pricing strategy.²⁸ Finally, Eichenbaum *et al.* (2011) show that replacement costs also exhibit reference behavior. It is thus plausible that reference price behavior reflects to some degree a similar feature of individual reset prices.

In the bottom of table 1, we present related statistics for the IRI Marketing database. They are computed using all available product categories over the 2001-2007 period in eight markets: New York City, Los Angeles, Washington D.C., Boston, Philadelphia, Chicago, Houston, and San Diego. These markets were selected to provide a wide coverage in terms of size and geography while keeping the number of observations to a manageable level. Price adjustments are very common: The average *weekly* frequency is 22.6 percent for posted prices, and 18.1 percent when short-lived V-shaped price drops are filtered out. If we only use the last price observation of each month, we

²⁸See Levy *et al.* (2011) for evidence that some digits are overrepresented in posted prices.

obtain a monthly average frequency of 37.2 percent for posted prices, and 33.6 percent after filtering out V-shaped price drops lasting 1 month or less. As was the case with CPI data, reference prices are much more rigid than posted or regular prices. The weekly frequency of reference price changes is only 4.1 percent when we define the reference price as the modal price over a centered 13-week window.²⁹ This statistic is consistent with an average duration just shy of 1 year. However, as was the case with CPI prices, scanner prices exhibit considerable novelty.

Similar to findings from the CPI database, there is some variation in the average frequency of price changes across time, product categories, and markets in the IRI Marketing dataset. For example, New York and Houston have a lower average weekly frequency (19.2 percent and 19.8 percent, respectively) than San Diego and Chicago (27.0 percent and 27.3 percent, respectively). In the beverage category, the average weekly frequency ranges from 13.2 percent for milk to 31.9 percent for carbonated beverages. There is even substantial variation across retailers for narrowly-defined products.

²⁹We depart slightly from the methodology in Eichenbaum *et al.* (2011), who compute the reference price as the modal price within a quarter. Their approach at times imprecisely identifies the number and timing of reference price adjustments. See Chahour (2011) for a discussion.

Table 1: Standard measures of the intensive margin in the U.S. CPI and the IRI Marketing database

U.S. CPI Database	Weekly frequency (percent)	Monthly frequency (percent)	Sample period
<i>Posted prices</i> ^a	n.a.	36.2	1/1988 – 1/2005
<i>Regular prices</i> ^a	n.a.	29.9	1/1988 – 1/2005
<i>Reference prices</i> ^b <i>13-month window</i>	n.a.	14.6	1/1988 – 10/2009
<i>Novel prices</i> ^b	n.a.	25.0	1/1988 – 10/2009
IRI Marketing Database	Weekly frequency (percent)	Monthly frequency (percent)	Sample period
<i>Posted prices</i>	22.6	37.2	1/2001 – 12/2007
<i>Regular prices</i> ^c	18.1	25.6	1/2001 – 12/2007
<i>Reference prices</i> <i>13-week window</i>	4.0	13.5	1/2001 – 12/2007
<i>Novel prices</i>	n.a.	23.8	1/2001 – 12/2007

Notes: (a) Source: Klenow and Kryvtsov (2008). Regular prices correspond to posted prices excluding sales and promotions as flagged by the BLS. (b) Source: Klenow and Malin (2010). Novel prices are defined as prices that have not been observed for at least 12 months. An item's reference price is defined as the modal price over a centered 13-month window. (c) Regular prices exclude price drops returning to their original level within 2 weeks. The reference price is defined as the modal price over a centered window.

Table 2: Lower bounds on the share of the initial price level response to the late-1994 Mexican peso devaluation accounted for by adjustment along the extensive margin

Bound method	Initial pass-through ($\Delta p/\Delta m$, in %)	Share attributed to extensive margin (in %)
Full pass-through ($\Delta e = 50\%$)		
posted prices	84.7	70.5
unprocessed food	72.9	47.4
processed food	85.4	75.2
nonenergy industrial goods	95.9	84.8
services	57.8	84.5
Full pass-through (forecast revision)		
posted prices	30.8	18.9
unprocessed food	26.5	-44.9
processed food	31.0	31.7
nonenergy industrial goods	34.9	58.1
services	21.0	57.4
Calvo +		
posted prices	67.5	63.0
unprocessed food	60.1	36.2
processed food	83.9	74.7
nonenergy industrial goods	64.2	77.3
services	46.2	80.6

Notes: The “full-pass-through” method assumes no real rigidities and full immediate pass-through of the shock to individual reset prices. The long-run change in consumer prices implied by a 50-percent devaluation assumes a total import content of consumption expenditures equal to 10.9 percent, as estimated by Burstein *et al.* (2005). The inflation forecast revision corresponds to the mean inflation forecast for the calendar year 1995 observed between November 1994 and February 1995 in the Bank of Mexico’s *Encuesta sobre las expectativas de los especialistas en economía del sector privado*. The “Calvo+” method assumes no real rigidities and the random selection of items whose price adjustment is triggered by the shock.

Table 3: Lower bounds on the contribution of adjustment along the extensive margin to the initial price level response to the January 2001 and April 1995 VAT hikes

	Initial pass-through ($\Delta p/\Delta m$, in %)	Share attributed to extensive margin (in %)
January 2010 hike		
<i>Full pass – through</i>		
posted prices	72.3	75.2
nonenergy industrial goods	83.6	75.3
services	34.8	66.2
<i>Calvo +</i>		
posted prices	266.9	93.3
nonenergy industrial goods	104.0	80.1
services	–38.8	130.3
April 1995 hike		
<i>Full pass – through</i>		
posted prices	62.8	23.5
nonenergy industrial goods	67.8	15.0
services	41.3	24.7
<i>Calvo+</i>		
posted prices	–249.1	119.3
nonenergy industrial goods	134.4	57.1
services	–356.1	108.7

Notes: The “full-pass-through” method assumes no real rigidities and full immediate pass-through of the VAT hike to individual reset prices of taxable items. The “Calvo+” method assumes no real rigidities and the random selection of items whose price adjustment is triggered by the shock. Only taxable items were used in the computation of the statistics.

Table 4: Probability of Price Adjustment

		New York City	Los Angeles	Boston
		Salty snacks	Carbonated beverages	Frozen dinner entrees
		(1)	(2)	(3)
Regressor	Coef.			
$\tilde{x}_{i,t}^{\text{neg}}$	β_1	-1.220 (0.029)	-0.979 (0.018)	-0.891 (0.053)
$\tilde{x}_{i,t}^{\text{pos}}$	β_2	0.665 (0.035)	0.574 (0.015)	0.495 (0.042)
R^2		0.11	0.06	0.07
N		1,056,906	938,349	396,471

Notes: Fixed effect (store/UPC) linear probability regression model for the monthly change in reference prices across different categories and markets:

$$I(\Delta p_{i,t} \neq 0) = \beta_{i,0} + \beta_1 \tilde{x}_{i,t}^{\text{neg}} + \beta_2 \tilde{x}_{i,t}^{\text{pos}} + Z'_{i,t} \Gamma + \varepsilon_{i,t},$$

where $I(\Delta p_{i,t} \neq 0)$ is an indicator variable that a reference price change has occurred, $\tilde{x}_{i,t}^{\text{neg}}$ equals $\tilde{x}_{i,t}$ if $\tilde{x}_{i,t} < 0$ and zero otherwise (with $\tilde{x}_{i,t}^{\text{pos}}$ similarly defined for positive deviations), and $Z_{i,t}$ is a vector of control variables that includes price spell duration, monthly dummies, store dummies, and a measure of store size (total yearly revenues). Standard errors (in parentheses) are corrected for heteroskedasticity and clustered at store level. N is the number of monthly observations from January 2001 to December 2007.

Table 5: Probability of Price Adjustment by Spell Duration Bins

		Salty Snacks in New York City				
		Increases vs. Decreases		Price Duration Spells (D , in months)		
		$\Delta p > 0$	$\Delta p < 0$	$D = 2$	$D = 6$	$D = 12$
		(4)	(5)	(6)	(7)	(8)
Regressor	Coef.					
$\tilde{x}_{i,t}^{\text{neg}}$	β_1	-1.304 (0.024)	0.084 (0.009)	-1.311 (0.036)	-1.113 (0.071)	-1.521 (0.013)
$\tilde{x}_{i,t}^{\text{pos}}$	β_2	0.162 (0.009)	0.504 (0.034)	0.652 (0.034)	0.622 (0.059)	0.498 (0.090)
R^2		0.17	0.05	0.15	0.06	0.06
N		1,056,906	1,056,906	105,058	52,669	29,333

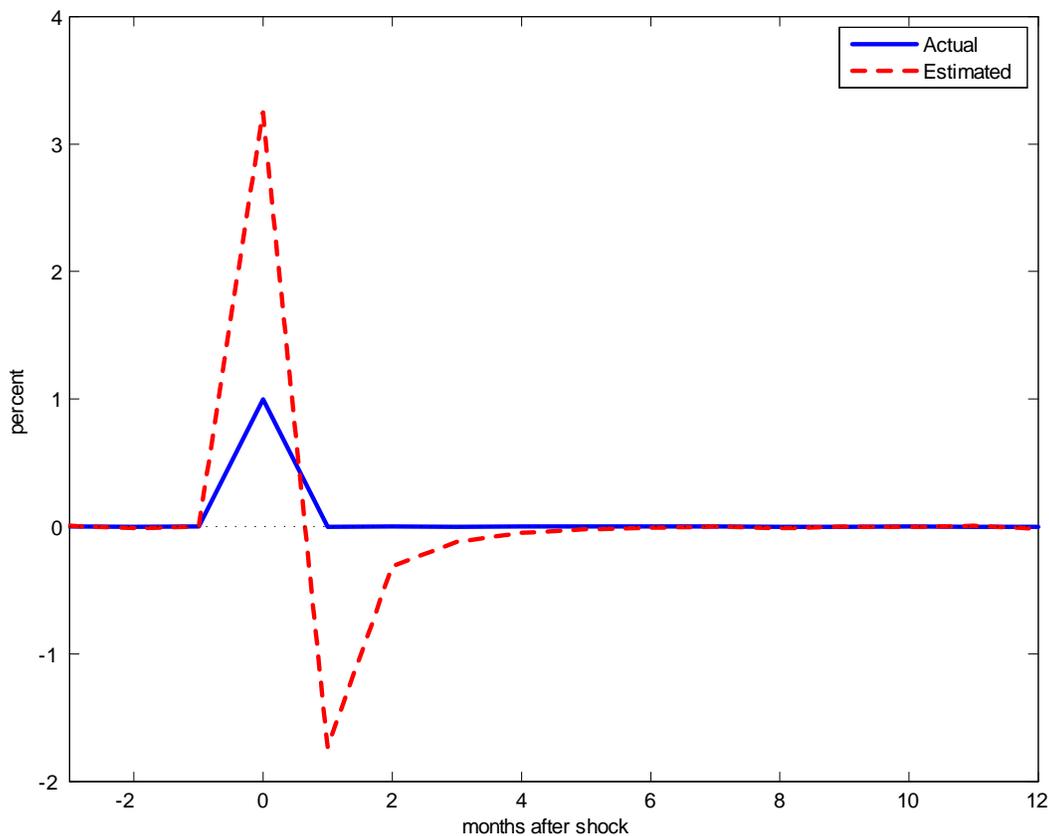
Notes: All the estimates correspond to salty snacks in New York City. Columns (4) and (5) are the linear probability models estimated for increases and decreases in prices, respectively. Columns (6) to (8) are the estimates for different age groups of prices (spells in months). See also the previous table for details. N is the number of monthly observations from January 2001 to December 2007.

Table 6: Index of Macroeconomic Flexibility

	\mathcal{A}_t	\mathcal{E}_t	F_t^{macro}	
			$a = 0$	$a = 0.6$
<i>Posted Prices</i>	0.27	0.10	0.37	0.19
<i>Regular Prices</i>	0.18	0.08	0.25	0.12
<i>Reference Prices</i>	0.09	0.04	0.13	0.06

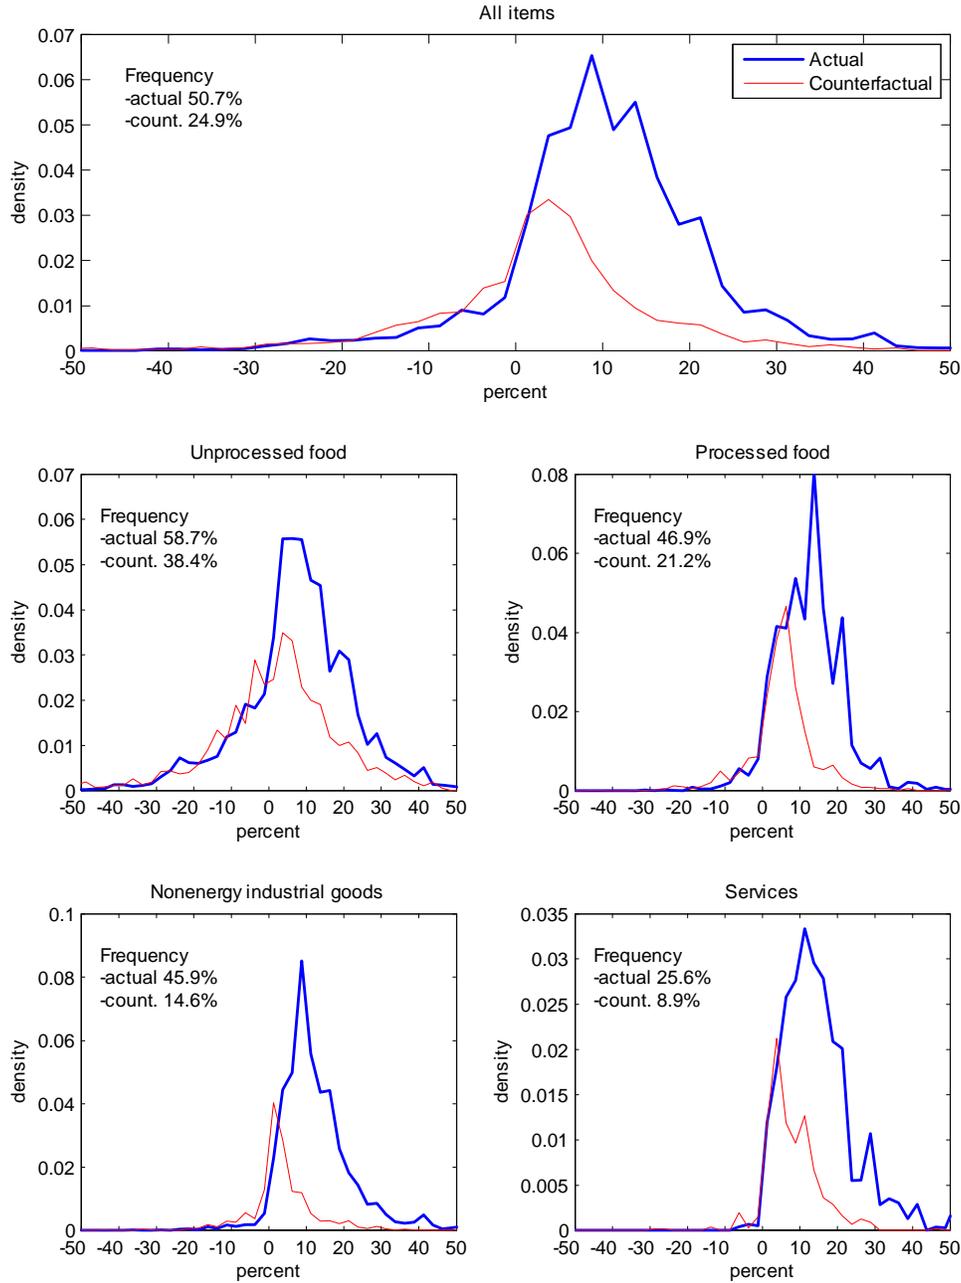
Notes: This table reports estimates of the intensive and the extensive margins based on equation 7 for the IRI Marketing dataset. We use deviations from the average price of local competitors to proxy for deviations from individual reset prices and to estimate a nonparametric adjustment hazard function. The index of macroeconomic flexibility is computed given no real rigidities ($a = 0$) and given moderate real rigidities ($a = 0.6$, as estimated by Burstein and Hellwig [2007])

Figure 1: Estimated aggregate reset price inflation using BKM's methodology in response to a 1-percent shock to actual aggregate reset price inflation (menu-cost model)



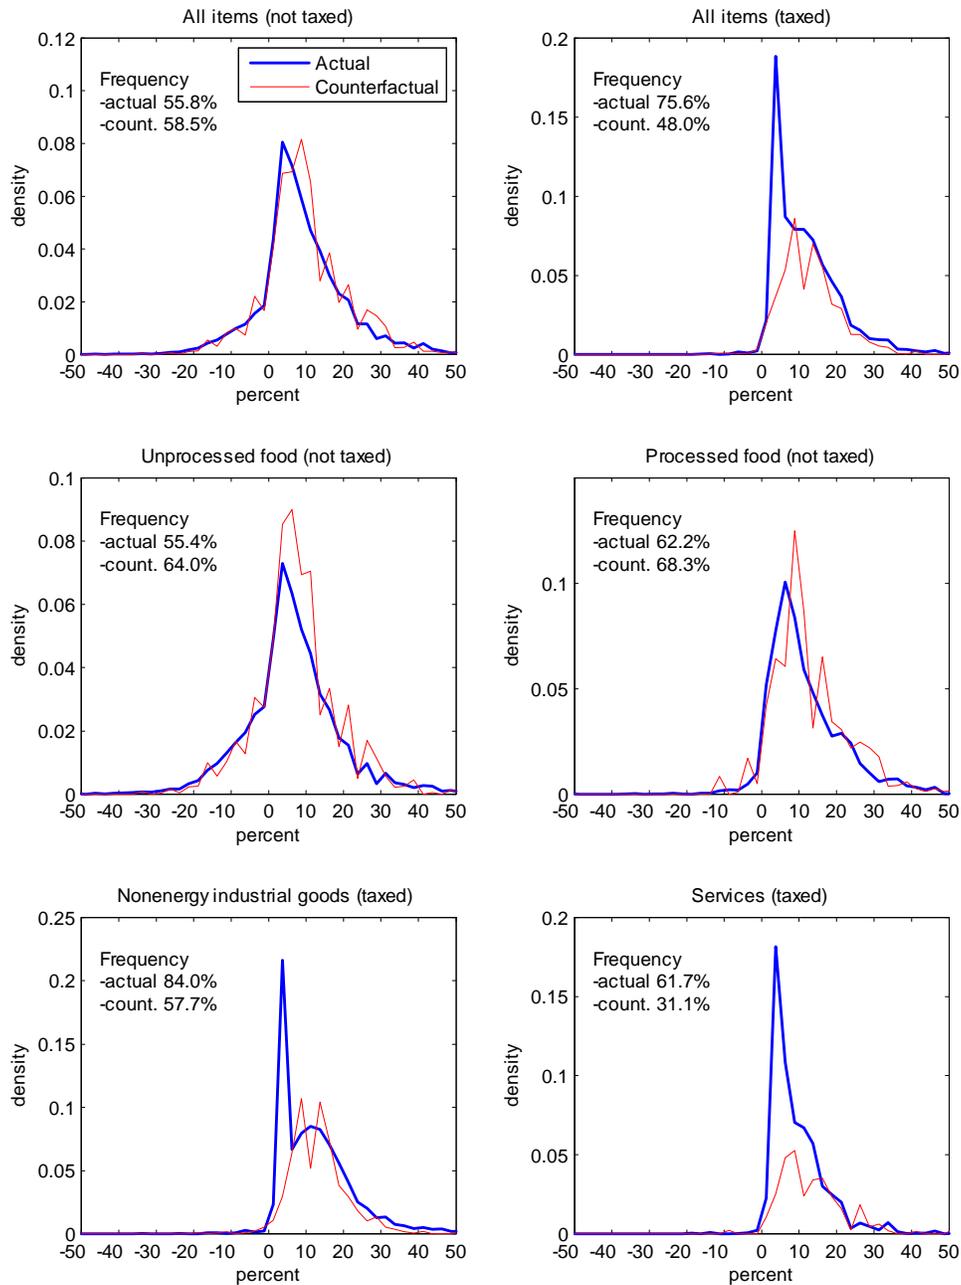
Notes: This figure applies BKM's methodology for identifying aggregate reset price inflation to the study of a 1-percent jump in the level of individual reset prices in a baseline menu-cost model. The model assumes a symmetric Ss band, a process for individual reset prices as in equation (9), and normally distributed idiosyncratic shocks. The model is calibrated to match an average frequency of price changes of 25 percent and an average size of price changes of 10 percent.

Figure 2: Initial impact of the Mexican peso devaluation on the distribution of price changes



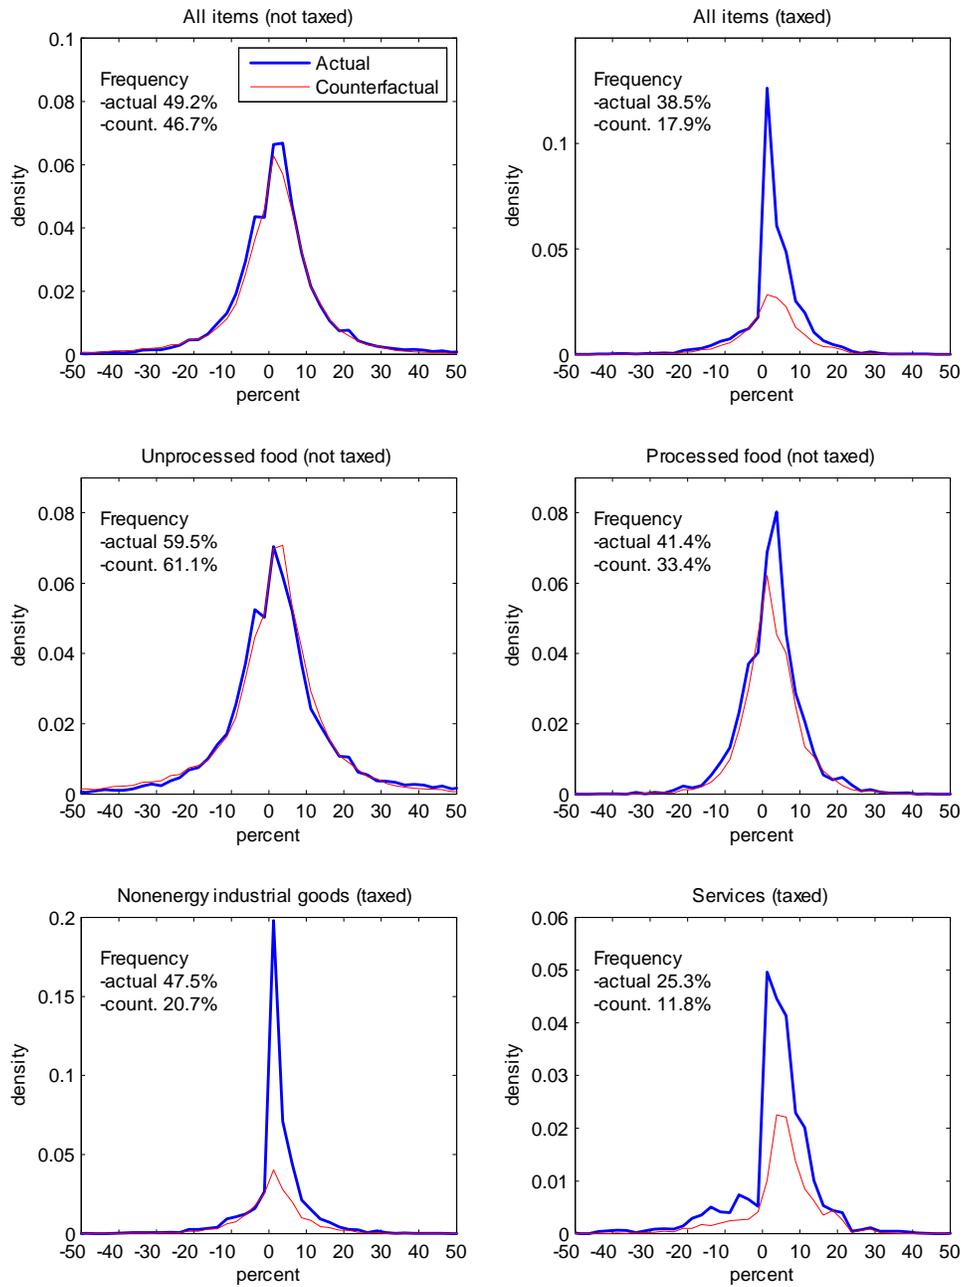
Notes: This figure shows the distribution of nonzero price changes in the Mexican CPI observed in January 1995 (the “actual” distribution) and in January 1994 (the “counterfactual” distribution) for special groups of products. Observations are grouped into bins of 2.5 percentage points and weighted by their relative importance in the CPI. The inserts in the panels report the monthly frequency of price changes.

Figure 3: Initial impact of the April 1995 5-percent hike in the Mexican VAT on the distribution of price changes



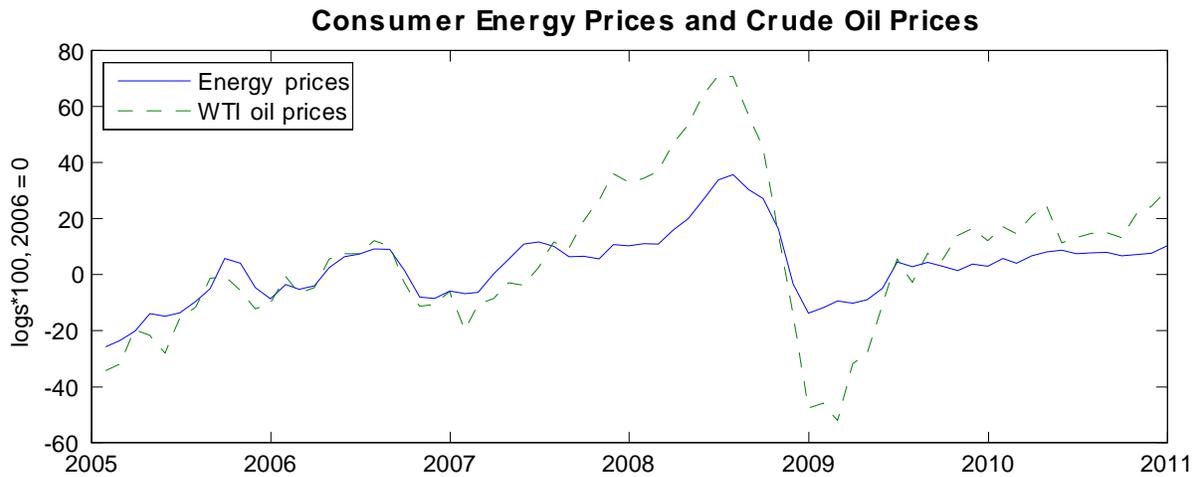
Notes: This figure shows the distribution of nonzero price changes in the Mexican CPI observed in April 1995 in regions subject to the VAT hike (the “actual” distribution) and in regions not subject to the VAT hike (the “counterfactual” distribution) for special groups of products. Observations are grouped into bins of 2.5 percentage points and weighted by their relative importance in the CPI. Inserts in the panels report the monthly frequency of price changes.

Figure 4: Initial impact of the January 2010 1-percent hike in the Mexican VAT on the distribution of price changes

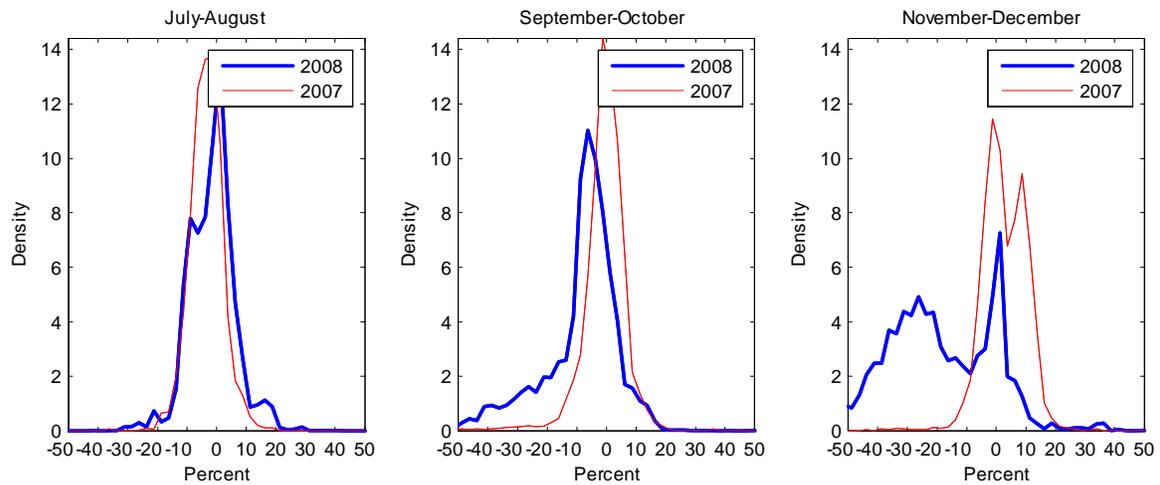


Notes: This figure shows the distribution of nonzero price changes in the Mexican CPI observed in January 2010 (the “actual” distribution) and an average of distributions observed in January from 2003 to 2007 (the “counterfactual” distribution) for special groups of products. Observations are grouped into bins of 2.5 percentage points and weighted by their relative importance in the CPI. Inserts in the panels report the monthly frequency of price changes.

Figure 5: Response of U.S. consumer energy prices to late 2008 crude oil price movements



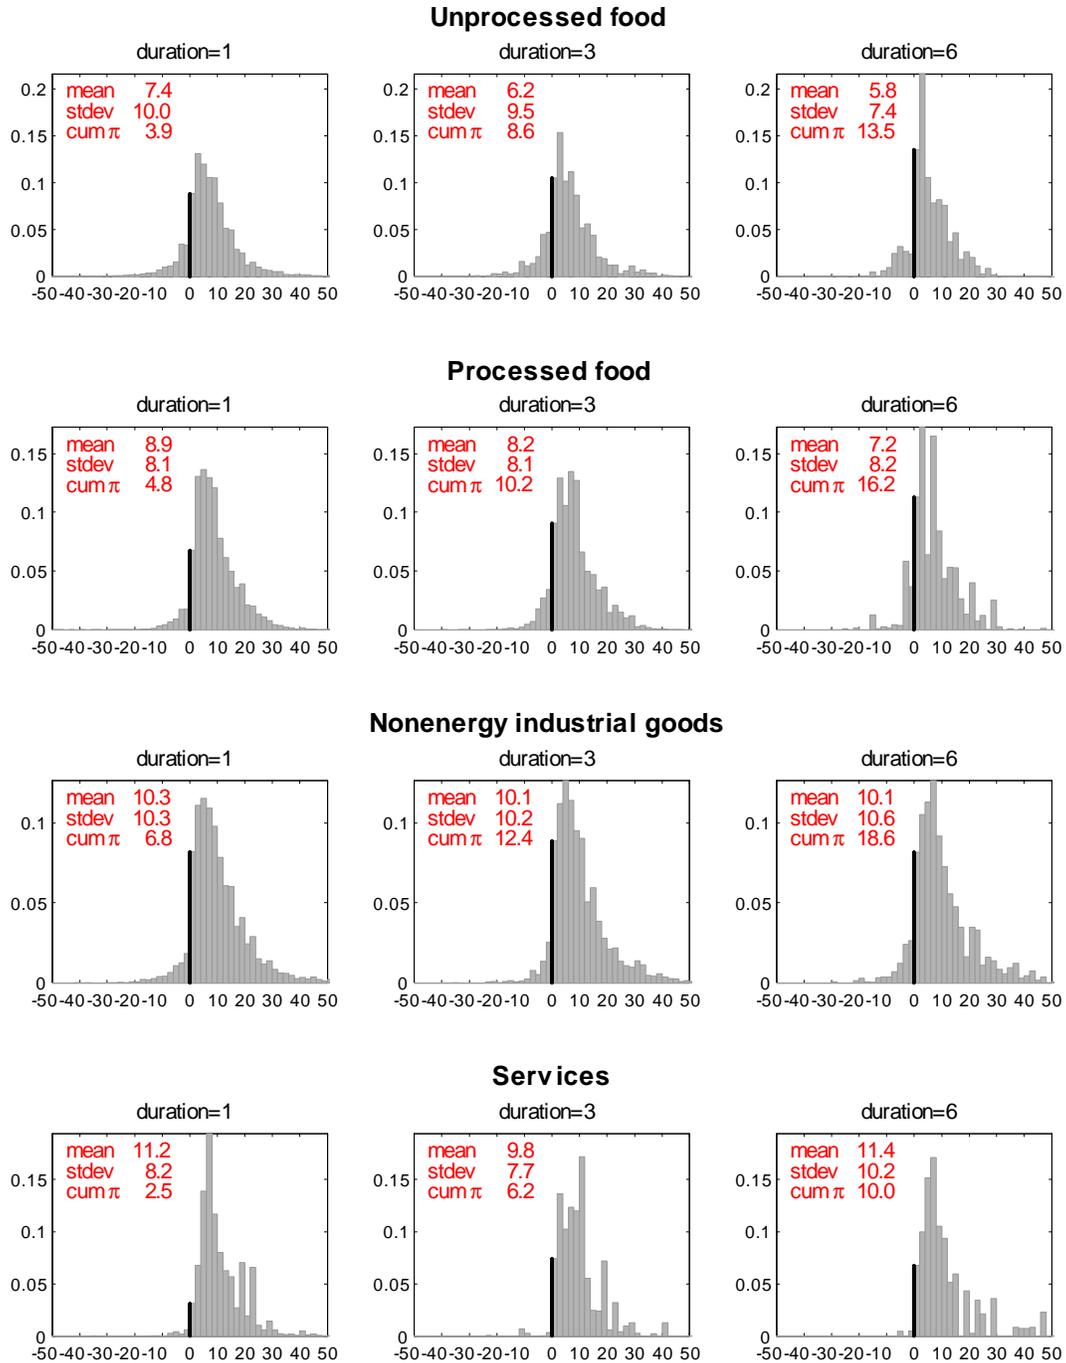
Distribution of Consumer Energy Price Changes



Estimated initial response of U.S. consumer energy prices to oil price movements in 2008

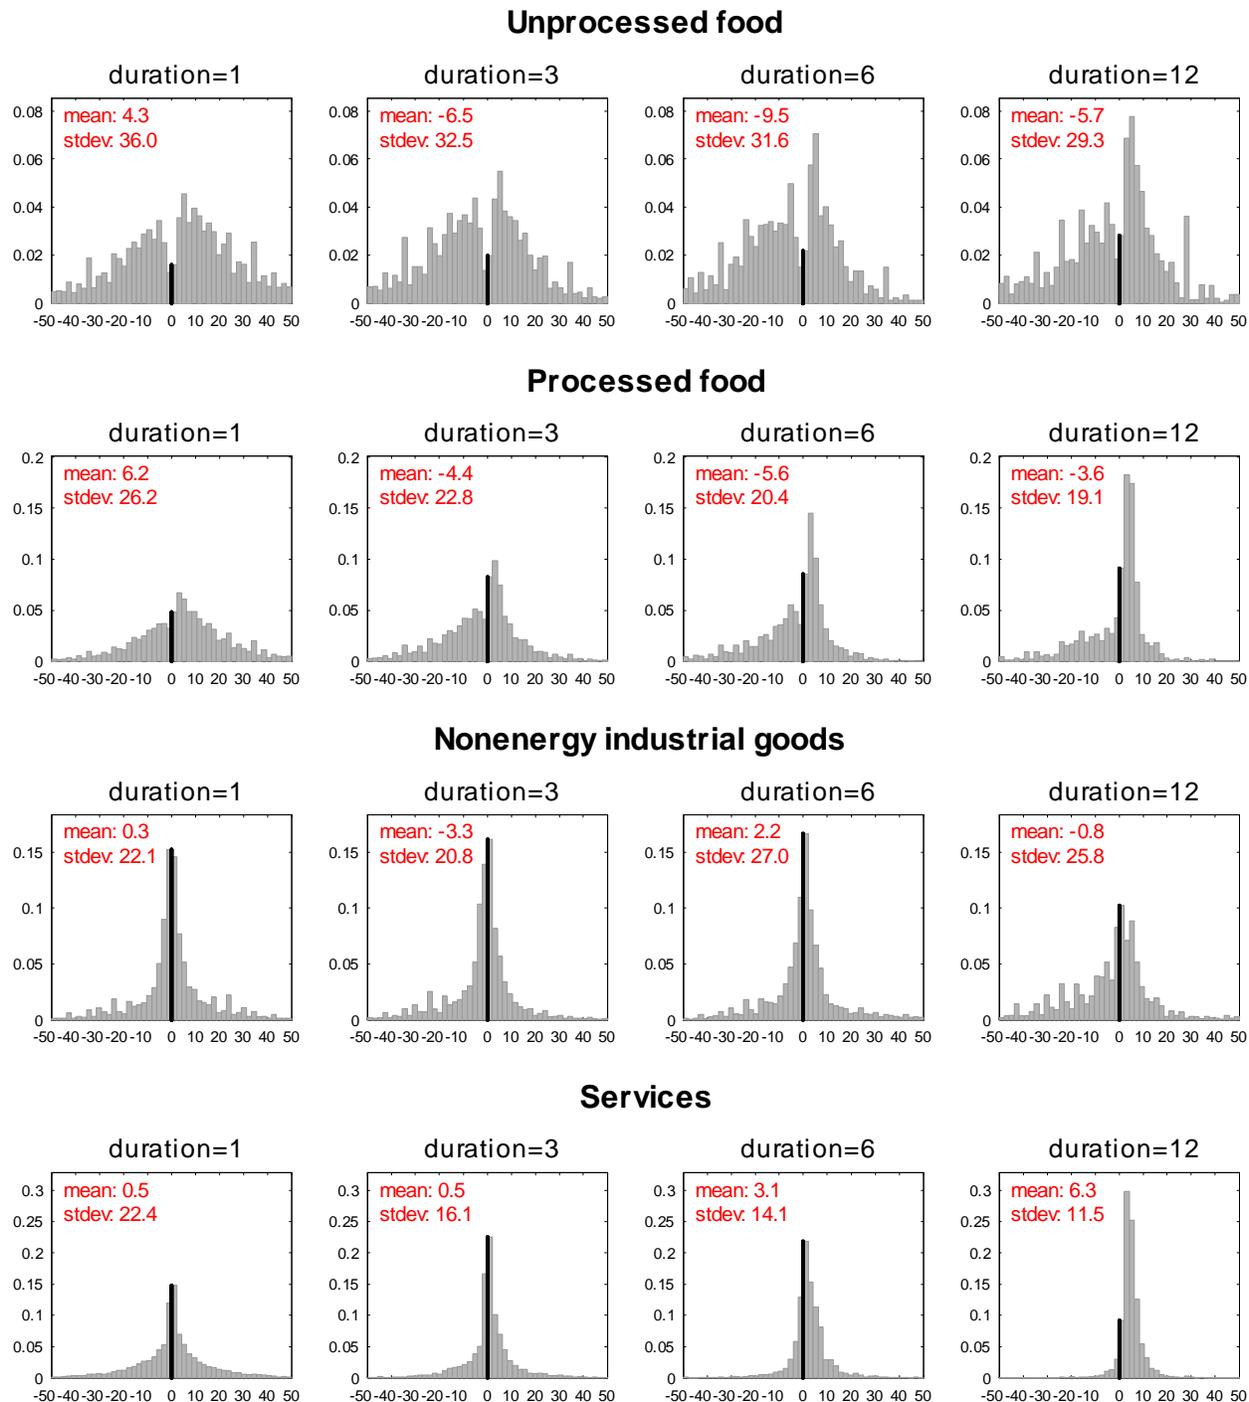
Bound method	Initial pass-through ($\Delta p / \Delta p^*$, in percent)	Lower bound (percent)
Full pass-through		
September-October	88.4	16.0
November-December	97.3	24.9
Calvo+		
September-October	78.5	5.2
November-December	79.8	8.1

Figure 6: Distribution of nonzero price changes in the Mexican CPI conditional on a price change in April 1995



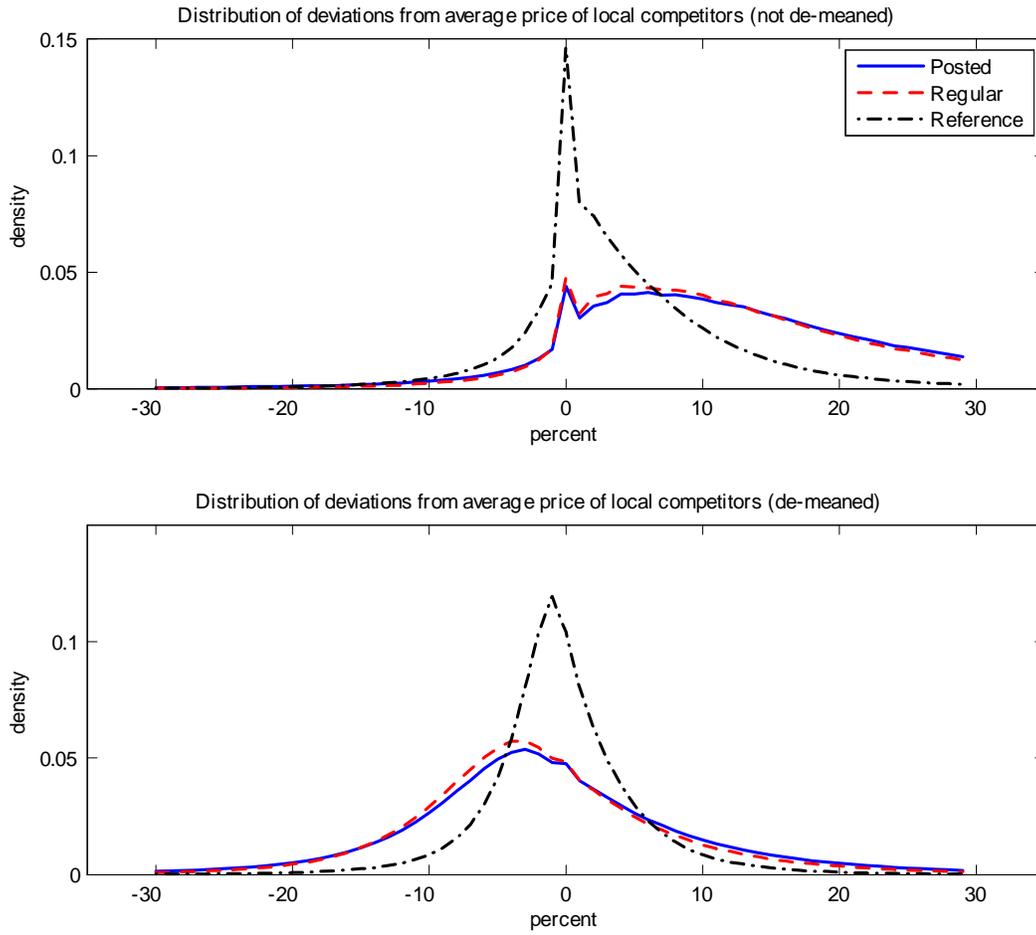
Notes: This figure shows the size distribution of first nonzero price changes in the Mexican CPI following the VAT hike in April 1995. The distributions are conditional observing a nonzero price adjustment in April 1995. The upper-left corner of each panel reports the mean of the distribution (in percent), its standard deviation (in percentage points), and the amount of inflation that has accumulated in the special group since April 1995 (in percent).

Figure 7: Distribution of nonzero posted price changes in the U.S. CPI conditional on the duration since the last price change



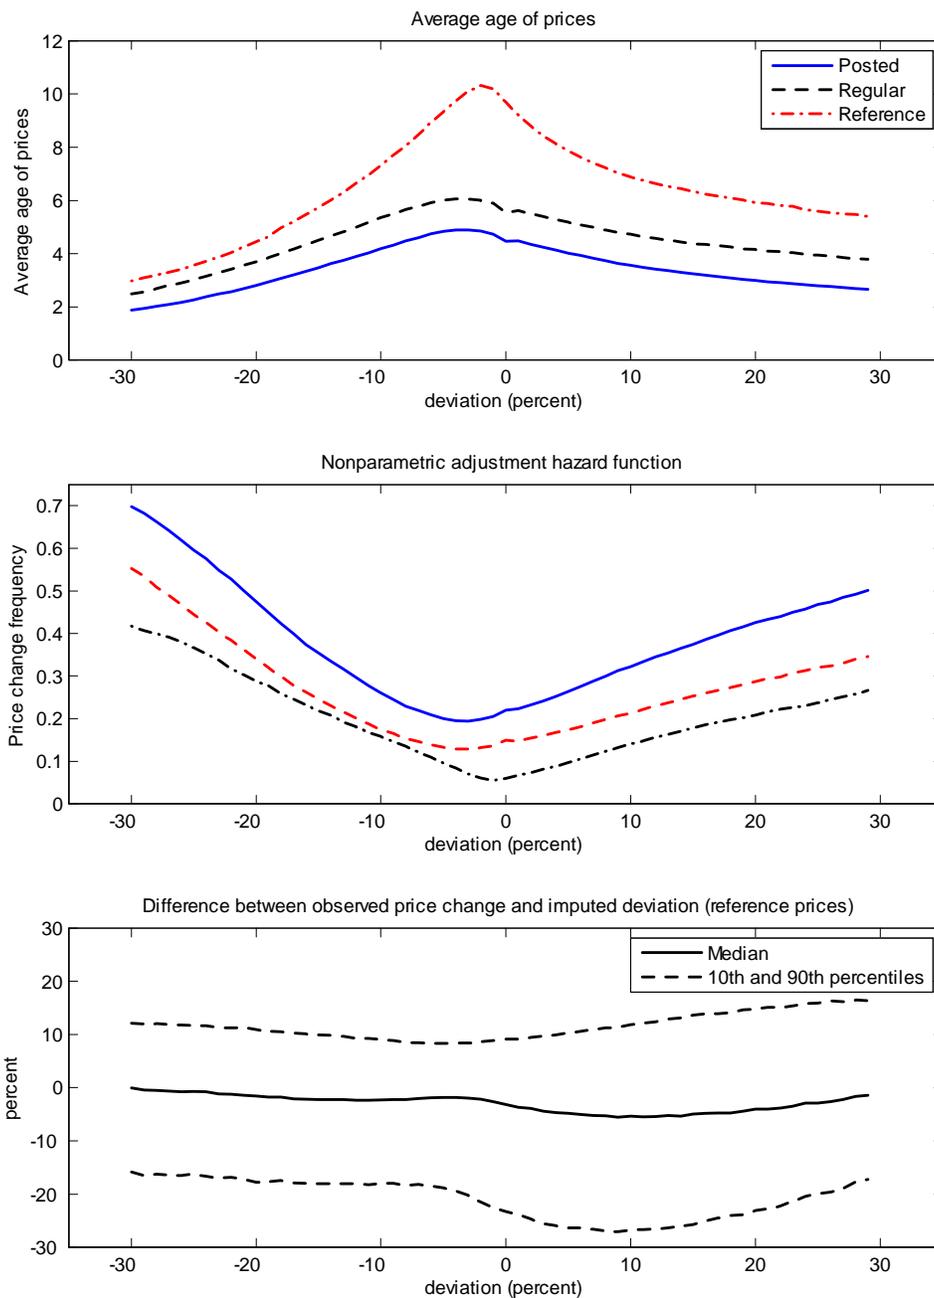
Notes: This figure shows the size distribution of nonzero posted price changes in the U.S. CPI conditional on the duration (in months) since the last price change. The sample period is January 1988 to August 2010.

Figure 8: Deviation from the average posted price of competitors and individual price adjustments



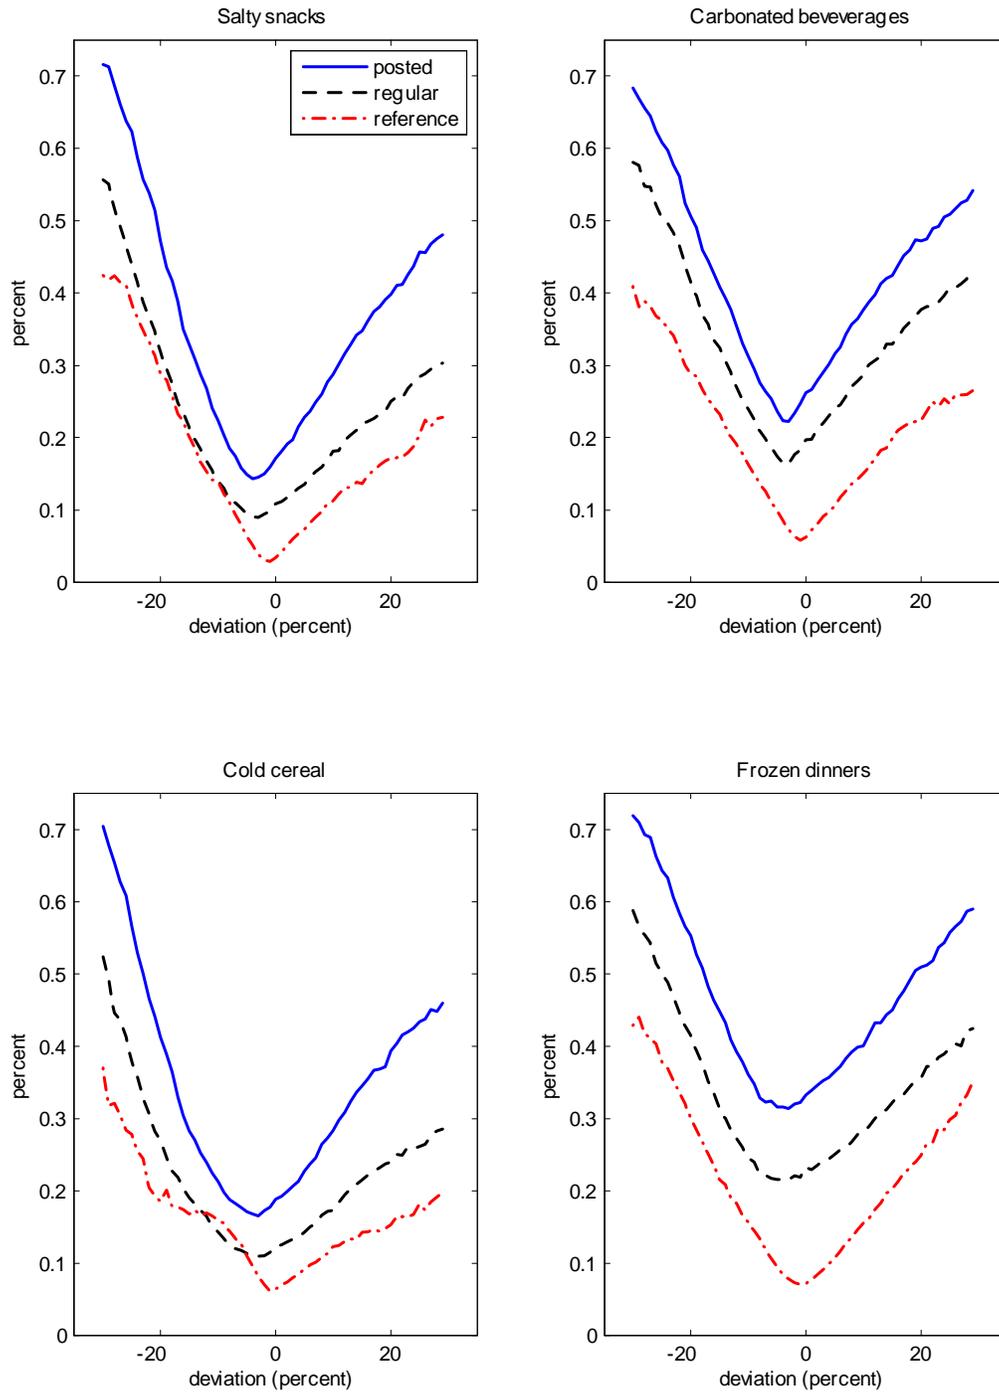
Notes: The upper and lower panels show the distribution of deviations from the average price of local competitors, computed at the UPC-market level in the IRI Marketing dataset, before and after removing item fixed effects to account for permanent differences in the level of prices across stores.

Figure 9: Deviation from the average reference price of competitors and individual price adjustments



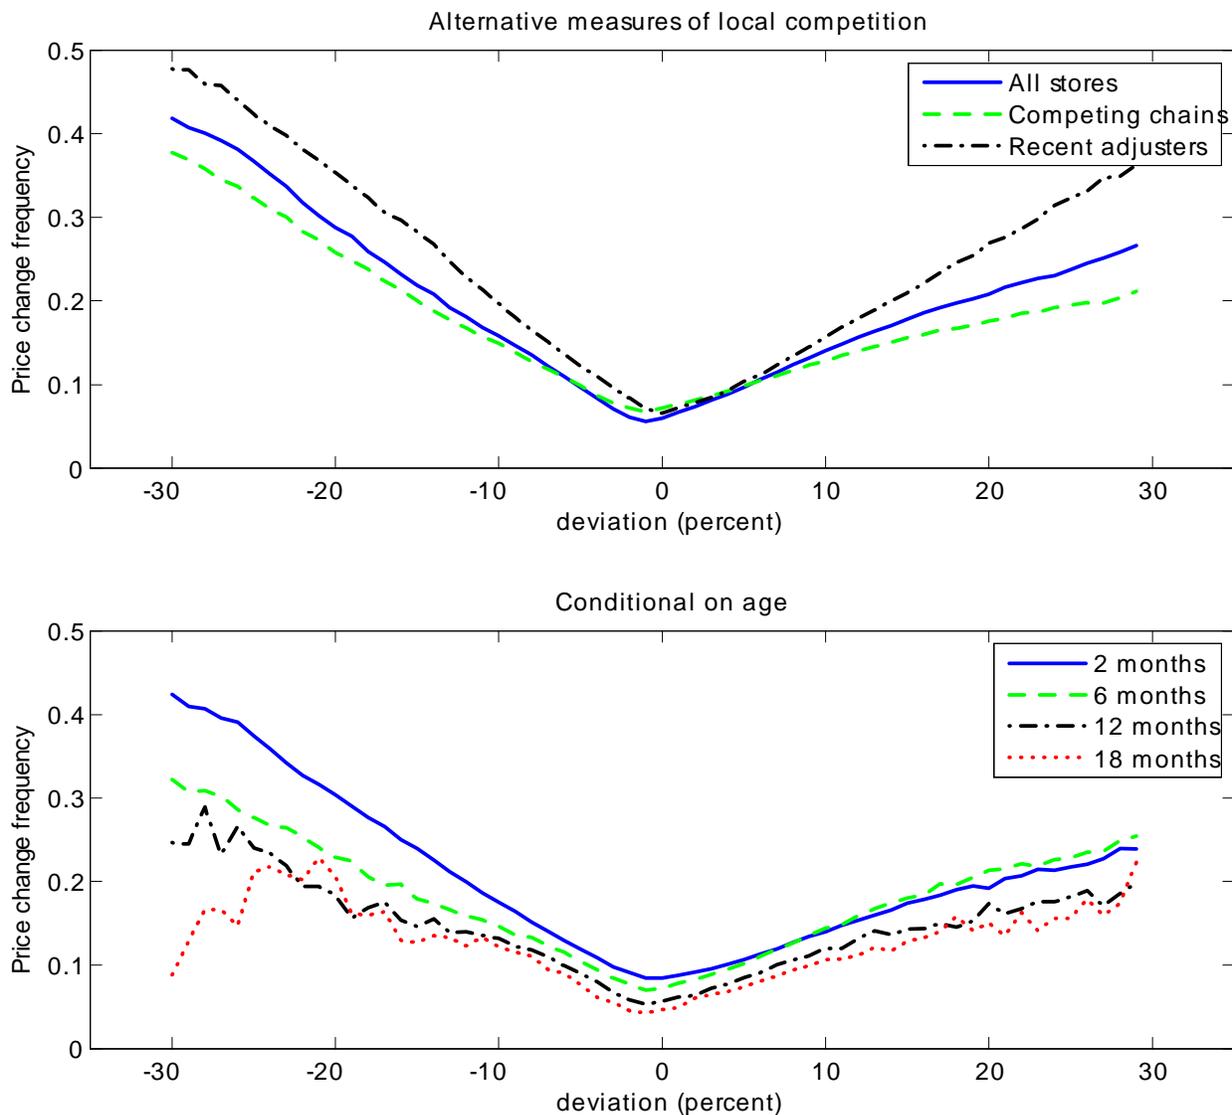
Notes: The upper panel shows the average age of monthly posted, regular, and reference prices in the IRI Marketing dataset conditional on the (demeaned) deviation from the corresponding average price of local competitors. The middle panel shows the fraction of items experiencing a price change during the month conditional on the deviation. The lower panel shows the median difference between the observed price change ($\Delta p_{i,t}$) and the imputed price pressure ($-\tilde{x}_{i,t}$), along with the 10th and 90th percentiles differences.

Figure 10: Nonparametric adjustment hazard function for four product categories



Notes: The figure shows the monthly probability of observing a change in the reference price conditional on the (demeaned) deviation from the average reference price of local competitors for four product categories in the IRI Marketing dataset.

Figure 11: Nonparametric adjustment hazard conditional conditional on age and alternative measures of local competition



Notes: The upper panel shows the probability of a reference price change conditional on alternative measures of the (demeaned) deviation from the prices of local competitors. The measures use either all other stores (“All stores”), stores belonging to competing chains (“Competing chains”), or retailers that have adjusted their price over the current or previous six months (“Recent adjusters”). The lower panel shows the probability of a reference price change as a function of the deviation from the average price of competitors, conditional on the duration since the last price change. The data source is the IRI Marketing dataset .