Price Selection in the Microdata *

Peter Karadi† Raphael Schoenle‡ Jesse Wursten§

December 1, 2023

Abstract

The flexibility of the aggregate price level affects the ability of monetary policy to stabilize business cycles. Monetary policy can be ineffective even if only few prices adjust, as long as those price changes are disproportionately large and make the aggregate price level flexible. We show that the most misaligned prices change with the highest probability. Systematic selection of large price changes is absent, however, following aggregate shocks. Instead, aggregate shocks bring about a uniform shift between price increases and decreases. These results are consistent with a particular class of state-dependent pricing models and imply a sizable, real impact of monetary policy.

JEL codes: E31, E32, E52

Keywords: monetary non-neutrality, state-dependent pricing, identified credit and monetary policy shocks, price-gap proxy, scanner data, PPI microdata

*All opinions expressed are personal and do not necessarily represent the views of the European Central Bank or the Eurosystem, the Federal Reserve Bank of Cleveland or the Federal Reserve System. Part of this research was conducted with restricted access to the Bureau of Labor Statistics (BLS) data. The views expressed here are those of the authors and do not necessarily reflect the views of the BLS. We thank Andrey Alexandrov, David Berger, Florin Bilbiie, Luca Dedola, Yuriy Gorodnichenko, Cosmin Ilut, Oleksiy Kryvtsov, Anton Nakov, Luminita Stevens, Ludwig Straub, Ivan Werning and participants at the Ifo Conference 2020 and seminars at the Central Bank of the Netherlands (DNB), Duke University, the European Central Bank, and the University of Maryland for their comments and Lydia Cox for excellent research assistance.

†European Central Bank and CEPR

‡Federal Reserve Bank of Cleveland, Brandeis University and CEPR

§KU Leuven and IRLE, UC Berkeley. This author gratefully acknowledges financial support from the Research Foundation Flanders (FWO) under grant number 12B1722N.

1 Introduction

The ability of monetary policy to stabilize business cycles depends on the flexibility of the aggregate price level. If the overall price level is flexible, shocks are absorbed by "prices," but if the overall price level is sticky, real economic activity adjusts instead. However, the price level can be flexible even if only a few prices adjust, as long as the prices that react do so in a disproportionate manner. Such "selection" of large price changes tends to reduce the real effects of monetary policy shocks in price-setting frameworks where frictions to price adjustment are micro-founded by fixed (menu) costs (Golosov and Lucas, 2007). In this paper, we propose a precise definition of price selection and we use micro price data to measure its extent, imposing discipline on models of price setting.

In our analysis, we first show that a *broad* concept of selection is at work in the data: the probability of price adjustment increases with the extent of price misalignment, that is, with the gap between optimal and actual prices. This evidence shows that price setting is state-dependent and challenges time-dependent price-setting frameworks (Calvo, 1983). Second, we define a *narrow* concept of selection as being present if an *aggregate shock* triggers or cancels the adjustment of prices that are systematically far away from their optimal levels. That is, we characterize narrow selection by the *covariance* between the marginal change in the probability of adjustment in response to a shock and the size of the new adjustments. Our model-free finding is that narrow selection is absent in the data: prices do respond to aggregate shocks directly, but the probability of doing so is not a function of the extent of their misalignment. Instead, we detect adjustment through a uniform shift between price increases and price decreases, which we call the *gross* extensive margin. These results imply that price adjustment in macroeconomic models should generally feature state-dependent adjustment through the gross extensive margin, but no narrow selection. We show that narrow selection is a useful new measure to distinguish *among* state-dependent price-setting models, and characterize conditions under which particular state-dependent pricing models (such as Dotsey et al., 1999; Alvarez et al., 2021) are consistent with our findings.

Our analysis first formally defines three adjustment margins through a suitable generalization of the flexible accounting framework of Caballero and Engel (2007). According to this generalization, the price-level response to an aggregate shock can be decomposed into the intensive margin – the change in the average size of price changes; the gross extensive margin – the uniform shift between price increases and price decreases; and the narrow selection effect – whether goods with larger price misalignments respond with a marginally higher probability. This decomposition builds on the decomposition in Caballero and Engel (2007), where the gross extensive margin and the narrow selection effects jointly make up the total extensive-margin effect, our broad selection effect. A model analysis calibrates a set of price-setting models to show the relevance of our decomposition: broad selection is useful in distinguishing between time-dependent and state-dependent models (as shown by Caballero and Engel, 2007); our newly introduced narrow selection effect is useful in distinguishing among different state-dependent pricing models. In performing this model exercise, we also show how the magnitude of both the broad and the narrow selection effects relies on an unobserved object, the generalized hazard function, which expresses the probability of price changes as a function of the price misalignment.

To gauge the extent of selection empirically and provide discipline for the selection of price-setting models, we quantify the importance of these three adjustment margins. We do so in two ways. The main approach is based on the estimation of a generalized hazard function that involves two ingredients: micro price data and estimates of product-level price misalignments. In terms of the first ingredient, our baseline analysis builds on a detailed weekly panel of barcode-level prices in U.S. supermarkets between 2001 and 2012, compiled by the marketing company IRi.[1] Second, to measure price misalignments, our analysis uses the distance of prices from the average of close competitors, after controlling for permanent differences across stores coming from the variation in geography and amenities (Gagnon et al., 2012). This "competitor-price gap" is a relevant measure of product-level price pressures as long as sellers strive to keep their prices close to competitors' prices to maintain profitability and market share (Mongey, 2021).

When using these ingredients to estimate a generalized hazard function that relates the probability of price adjustment to the price gap, we find that our price misalignment measure is a strong predictor of the probability of future price adjustments, while the average size of future adjustments is proportional to the estimated misalignment. In particular, our results show that the price-adjustment probability increases *linearly* with the price misalignment after we control for unobserved heterogeneity. The increase, furthermore, is significantly different from zero, but fairly low: it reaches only around 40% as the price gap approaches 50%.

These estimates of the empirical hazard along with the relevant density function estimates from the data allow us to quantify the relative strength of the adjustment channels. In a first step, our analysis simply takes the empirical estimates as their theoretical counterparts and uses them to compute our three margins of adjustment. Under this assumption, the intensive margin turns out to contribute the most in our decomposition framework, namely, 73.4%. Broad selection captures the remainder. However, the gross extensive margin makes up nearly the entirety of this remainder, with 26.5% of the total, with only 0.2% stemming from narrow selection. According to this calculation, the linear generalized hazard function we estimate implies zero narrow selection, and its low slope implies a moderate gross-extensive-margin effect.

In a second step, we calibrate a flexible quantitative price-setting model to our empirical estimates and repeat the decomposition based on the model moments, confirming the results. In this exercise, we estimate the model parameters so as to match the empirical hazard function with its model counterpart. The estimated model also fits the price-change distribution and the price-gap density well, even though these moments are not directly targeted. Crucially, we then use the model to calculate the theoretical decomposition of the intensive, gross-extensive, and narrow selection margins. The model-based decomposition is close to the decomposition that uses the empirical moments as theoretical moments, though the quantitative magnitude of the estimates is slightly different: according to the model the gross extensive margin captures around 20% (rather than 26% using the empirical moments as their theoretical counterparts). Overall, the results of the model exercise again assign no role to narrow selection.

The flexible theoretical model also allows us to assess the strength of the various adjustment margins in counterfactual scenarios with time-dependent constant hazard functions (Calvo, 1983) or strongly state-dependent step hazard functions (Golosov and Lucas, 2007) – providing guidance for the choice of price-setting models. Our model setup nests these specific price-setting models and our analysis recalibrates the model in each case to match the frequency and the size of price changes. Because the model solution generates the endogenous price-gap distribution and hazard function, we can then use that solution to decompose the impact of the various adjustment margins. A clear finding emerges: the time-dependent model implies no

[1] We would like to thank IRi for making the data available. All estimates and analyses in this paper, based on data provided by IRi, are by the authors and not by IRi. In a complementary analysis, we also show that our results are robust to using the producer-price microdata that underlie the U.S. producer-price index (PPI).

adjustment on either the gross extensive or the narrow selection margin. In contrast, the total extensive margin is powerful in the strongly state-dependent Golosov and Lucas (2007) model and raises the impact effect of the extensive-margin to 77.3%. Furthermore, both the narrow selection and the gross-extensive margins contribute to this impact effect nearly equally. Crucially, this finding justifies using narrow selection to distinguish between different state-dependent price-setting models. As our main calibration exercise shows, second-generation state-dependent price-setting models such as those featuring random menu costs (Dotsey et al., 1999; Alvarez et al., 2021) can be consistent with our empirical evidence. Specifically, models with a *linear* price-adjustment hazard function are consistent with our evidence. They generate state-dependent adjustment on the gross extensive margin without selection (unlike the Golosov and Lucas, 2007, model), and sizable monetary non-neutrality.

A second empirical approach to gauging the extent of selection also finds results consistent with the results from this quantitative model-based analysis, confirming that the gross extensive margin is active in the data, but narrow selection is absent. This second approach in particular focuses on measuring narrow selection by estimating how the probability of price adjustment depends on the *interaction* of an aggregate shock and the product-level price-misalignment proxy. We estimate this relationship in a linear-probability panel-regression framework. Our dependent variables are the probability of price increases and price decreases, respectively, over the 24 months following a credit shock, which we identify using timing as in Gilchrist and Zakrajšek (2012).[2] Our analysis controls for the current aggregate shock and the lagged price-gap proxy separately, as well as the age of prices, and includes a rich set of product and time fixed effects. The regression specification clusters standard errors across time and product categories.

Results from estimating this specification show that both the product-level price-misalignment proxy and the aggregate shock strongly affect the probability of price adjustment. However, the impact of their interaction term is close to zero and consistently insignificant. These results indicate that (i) product-level price misalignment increases the probability of price adjustment, so broad selection is present; (ii) the aggregate shock shifts the share of price increases and price decreases, so there is an adjustment at the gross extensive margin; and (iii) the change in the probability of adjustment following a shock is independent of the extent of price misalignment, so narrow selection is absent, in line with the findings of the first approach. These results are robust to using different price-gap proxies (competitor-price gaps or competitor-reset-price gaps), two different aggregate shocks (credit shock versus monetary policy shock), different data sets (supermarket scanner versus producer-price microdata), and different specifications (linear versus non-linear probability models). The gross extensive margin is always active, but narrow selection is absent.

Overall, the evidence from our two approaches, therefore, indicates that broad selection is present in the data, which poses a challenge to time-dependent price-setting models (Calvo, 1983), in which this channel is inactive. At the same time, it implies that broad selection emerges from a uniform shift between price increases and price decreases, which implies the presence of the gross extensive margin but no narrow selection. This is inconsistent with standard menu cost models (Golosov and Lucas, 2007; Karadi and Reiff, 2019; Bonomo et al., 2023), where narrow selection is generally strong. As our model exercise shows, models with a *linear* price-adjustment (hazard) function are consistent with our evidence. They generate state-dependent adjustment on the gross extensive margin without selection, and sizable monetary non-neutrality.

[2] In the online appendix, we show robustness to a monetary policy shock identified using high-frequency surprises in interest rates around policy announcements (Gertler and Karadi, 2015; Jarociński and Karadi, 2020).

Related literature Our work contributes to the strand of literature that imposes minimal structure to estimate the strength of price selection in micro price data. In particular, a subset of these papers relies on Caballero and Engel (2007), who derive an indirect measure of the extensive-margin effect from the unobservable adjustment hazard and density functions of price gaps. Berger and Vavra (2018) impose flexible functional forms and match unconditional moments of the price-change distribution to estimate hazard functions, finding a sizable extensive-margin effect (see also Petrella et al., 2019). Luo and Villar (2021) match moments conditional on trend inflation and, by contrast, find hazard functions that imply weak extensive-margin effects closer to our results. Relative to these papers, we generate proxies for the price gaps and report non-parametric estimates of the hazard functions as in Gagnon et al. (2012). Akin to Gagnon et al. (2012) and Campbell and Eden (2014), we find that the absolute value of the price gap increases the probability of price change. Additionally, we document that the relationship is linear after one takes into account unobserved heterogeneity. Linearity has also been detected, though not emphasized, in Eichenbaum et al. (2011) using price and cost data from a single major U.S. retailer and in Carlsson (2017) using Swedish producer-price microdata. A central relative contribution of our analysis lies in establishing that this linear hazard exerts its influence through the gross extensive margin, but does not imply active narrow selection.

Our analysis is complementary to the work of Carvalho and Kryvtsov (2021), who find no indication for price selection in U.S. supermarket data, in line with our results. They compare the average price of close substitutes to the preset prices of products that eventually adjust. They show that this preset-price-relative does not interact with aggregate inflation in a time-series setting. Our analysis instead constructs price-gap measures for both adjusted and unadjusted prices and shows in a panel-data setting that the price gaps do not interact with identified aggregate shocks. Our work is also complementary to Dedola et al. (2021), who assess whether *unobserved* product-level shocks generate a selection bias in a Heckman-type model using Danish producer-price data. Relative to this paper, our regression analysis contributes by measuring selection caused by the interaction of aggregate shocks with *observable* proxies of price gaps.

Our work is also related to the long strand of the literature that uses observations about micro-level price behavior to infer the level of monetary non-neutrality in fully specified state-dependent price-setting models (Dotsey et al., 1999; Golosov and Lucas, 2007; Gertler and Leahy, 2008; Midrigan, 2011), or derive sufficient statistics in a class of models (Alvarez et al., 2016; Baley and Blanco, 2021; Alvarez et al., 2021). These papers predominantly use unconditional moments of price changes to calibrate their parameters and deliver predictions that are model-dependent. We measure price-level adjustments conditional on product-level price misalignments and on identified aggregate shocks and our predictions can be consistent with multiple theoretical models.[3] The strength of monetary non-neutrality in these models depends on the strength of price selection. As our model exercise shows, our results favor those theoretical models that find a linear hazard and a minimal role for narrow selection. We also find that the realistic hazard is flat, with values well below 50% even for very large gaps. Our baseline model matching these features predicts sizable monetary non-neutrality.

The rest of this paper is structured as follows. We present a new framework to decompose inflation's

[3] Hong et al. (forthcoming) study the informativeness of pricing moments jointly with aggregate price responses to monetary policy shocks, indirectly testing the importance of selection through moments such as kurtosis. Balleer and Zorn (2019) document a small frequency response – similar to us – and a *decline* in the average absolute size of price changes after an identified monetary policy easing using German PPI microdata, arguing that these conditional moments are inconsistent with a large selection effect.

response to aggregate shocks into an intensive margin, a gross extensive margin, and a narrow selection effect in Section 2. We describe the data in Section 3. We generate an empirical hazard function and use it to measure selection using a naive and a structural method in Section 4. In Section 5, we measure selection in a panel-data framework conditional on an aggregate shock and show its robustness to the use of a non-linear specification, and producer-price microdata.[4] Section 6 concludes.

2 Selection: A conceptual framework

The real effects of a nominal shock can depend as much on *which* prices change as on *how many* prices adjust (Golosov and Lucas, 2007). In fact, if prices that are far from their optimal levels adjust, then nominal shocks and consequently monetary policy can be completely neutral even if only a small subset of prices adjust (Caplin and Spulber, 1987). Selection measures how far new adjusters are from their optimal prices when an aggregate shock hits. We preface our empirical analysis with a detailed discussion of selection, the presentation of our definitions, and how selection is tied to the hazard of price adjustment. A reader less interested in the technical details of selection may skip to our empirical results and the modeling exercise in Sections 3 and 4.4.

2.1 An inflation decomposition: The hazard function and the price-gap density

Before formally defining selection, we introduce in the next two sections a few relevant objects that we use to propose a decomposition of inflation that directly connects to selection. This decomposition builds directly on the general price-setting framework of Caballero and Engel (2007). Both the time-dependent Calvo (1983) model and the strongly state-dependent models with an (S,s) adjustment rule are nested within this framework, as well as a continuum of intermediate cases in line with random menu costs (Dotsey et al., 1999) and rational inattention (Woodford, 2009).

In this framework, time is discrete and there are a continuum of firms, each producing a single product i. Firms set the (log nominal) prices of their product (p_{it}) subject to a price-adjustment friction. If these frictions were temporarily absent, the optimal price in period t would be p^*_{it}. The optimal price is driven by both aggregate and idiosyncratic factors $p^*_{it} = m_t + \nu_{it}$. For simplicity, we assume that shocks to both m_t and ν_{it} are permanent. The aggregate shock m_t shifts the optimal nominal price of all firms, whereas the idiosyncratic shock ν_{it} affects only firm i. The gap between the price and its optimal value $x_{it} = p_{it} - p^*_{it}$ is the relevant state variable and is sufficient to characterize each firm's price-setting choice.

The firm's price-adjustment decision can be described by a generalized hazard function $\Lambda(x)$. The function takes values between 0 and 1, and its value expresses the probability of price adjustment for a firm with a price gap x. As is well known, the hazard function is constant in the time-dependent Calvo (1983) model: there, the probability of adjustment is independent of the price gap $\Lambda^{\text{calvo}}(x) = \bar{\Lambda}$. At the other extreme, in the fixed menu cost model (Caplin and Spulber, 1987; Golosov and Lucas, 2007), the hazard function $\Lambda^{(\text{S,s})}$ is a step function equal to 0 when the gap is within the lower and upper inaction thresholds ($x \in [x^l, x^u]$), and 1 otherwise. Caballero and Engel (2007) show that a continuum of intermediate hazard functions can

[4] In the appendix, we present additional robustness tests, including using identified monetary policy shocks instead of our baseline credit shocks; two alternative price-gap measures (the competitors-reset-price gap and the reset-price gap (Bils et al., 2012)), as well as robustness across heterogeneous product categories.

arise when the menu cost is an i.i.d. random variable (from a cumulative density function $G(\omega)$, as in Dotsey et al., 1999), or when the firm is subject to a rational inattention friction as in Woodford (2009) (see also Alvarez et al., 2021). We pay particular attention to a special intermediate case where the generalized hazard function is piece-wise linear as follows:

$$\Lambda^{\text{linear}}(x) = \begin{cases} a + bx & \text{if } x \geq 0 \text{ and } a + bx \leq 1 \\ a - cx & \text{if } x < 0 \text{ and } a - cx \leq 1 \\ 1 & \text{otherwise} \end{cases} \tag{1}$$

where a, b, c are parameters. As we show below, this case is broadly in line with our evidence.

In this economy, inflation is simply

$$\pi = \int -x\Lambda(x)f(x)dx \tag{2}$$

where $f(x)$ is the density of price gaps across firms, and we suppressed subscripts for notational convenience. The expression is intuitive: the opposite price gap is the size of price adjustment and the hazard is its probability. Their product summed across the gap distribution and weighted by the density is the inflation rate.

We further decompose inflation (π) as the sum of price changes driven by negative gaps, denoted by π^+, and price changes driven by positive gaps, denoted by π^-:

$$\begin{aligned} \pi &= \int -x\Lambda(x)f(x)dx \\ &= F(0)\int_{x<0} -x\Lambda(x)g(x)dx \;+\; \big[1 - F(0)\big]\int_{x\geq 0} -x\Lambda(x)h(x)dx \\ &= F(0)\,\pi^+ + \big[1 - F(0)\big]\,\pi^- \end{aligned} \tag{3}$$

where $g(x) = f(x)/\int_{x<0} f(x)dx$, and $h(x) = f(x)/\int_{x\geq 0} f(x)dx$ are new functions that are proper densities over negative gaps and positive gaps, respectively. $F(0) = \int_{x<0} f(x)dx$ is the share of products with negative gaps, and, naturally, its complement $\big[1 - F(0)\big]$ measures the share of positive gaps. This decomposition is novel, and we propose it to highlight the relevance of the gross extensive margin, which emerges from the shift in the relative share of price increases and price decreases as a response to an aggregate shock as we explain below. In the expressions below, we concentrate on the positive price gaps. The behavior of inflation caused by products with negative price gaps is analogous.

To illustrate how the inflation decomposition connects to different pricing models, we contrast how it works for three different pricing frameworks: a state-dependent model with a fixed menu cost and an (S,s) type adjustment (Golosov and Lucas, 2007), the time-dependent model of Calvo (1983), and an intermediate case with a linear price-adjustment hazard. As discussed above, the generalized hazard function ($\Lambda(x)$) is a step function in the first case and a constant in the second case and has a censored V shape in the third case. Figure 1 provides a graphical illustration of the key features of price setting in these models, focusing on the main objects of interest: the density of price gaps $f(x)$ is depicted by the grey shaded area[5] and the

[5]For ease of presentation, the figures do not reflect the fact that the gap density is an endogenous object and depends on the nature of price setting. Our model-based analysis, as well as the empirical analysis, however, explicitly takes this into account.

Figure 1: Price-setting frameworks

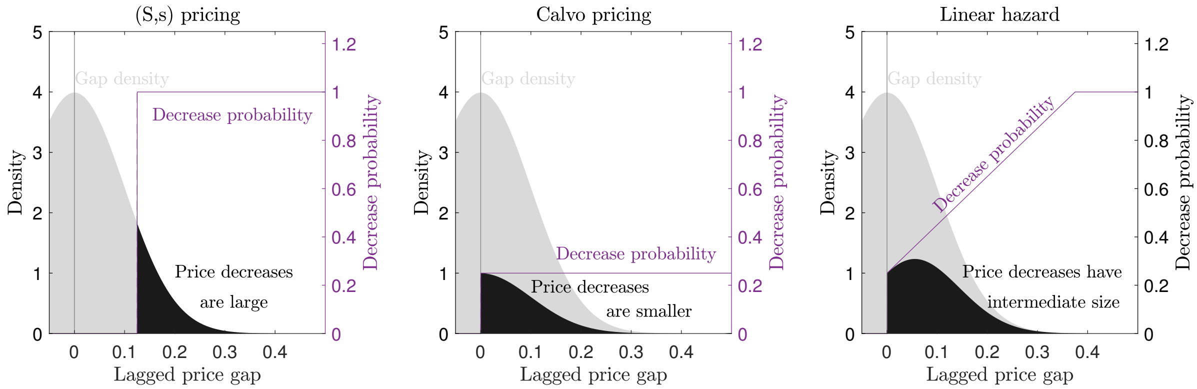

Note: The figure shows the density of the price gap, the hazard over positive price gaps, and the price-decrease density as a function of gaps in an (S,s) pricing model (left panel), a time-dependent model (middle panel), and a linear hazard model. The differences between the hazard functions imply large differences in the predicted relationship between the gap and the probability of price decreases.

density of price decreases is shown by the black shaded area. The latter is the product of the adjustment hazard and the density of gaps.

In particular, the panels in Figure 1 illustrate the contrast between the size distribution of price decreases in the pricing frameworks: while price changes are on average large in the (S,s) model (left panel), they are on average small in the time-dependent model (middle panel), while they lie in between in the linear hazard case (right panel). The left and the right panels also illustrate a general feature of state-dependent models: The probability of price adjustment increases with the size of the price gap (non-decreasing hazard). In our subsequent empirical implementation, we provide evidence showing that the hazard function is indeed increasing in the data, supporting state-dependence in price setting.

One can formally capture these features (and then, in the next section, tie them to selection) by decomposing (the positive-gap) inflation into a mean and a covariance term (Costain and Nakov, 2011b):

$$\begin{aligned} \pi_t^- &= \int_{x\geq 0} -x\Lambda(x)h(x)dx \\ &= -\bar{x}^-\,\bar{\Lambda}^- + \int_{x\geq 0} -\left(x-\bar{x}^-\right)\left(\Lambda(x)-\bar{\Lambda}^-\right)h(x)dx \end{aligned} \tag{4}$$

where $\bar{x}^- = \int_{x\geq 0} xh(x)dx$ is the average positive price gap, and $\bar{\Lambda}^- = \int_{x\geq 0}\Lambda(x)h(x)dx$ is the average probability of price adjustment among products with positive price gaps. This decomposition implies that the positive price-gap inflation is the sum of (a) the product of the opposite of the average positive price gap and the average frequency of price decreases and (b) a covariance term between the opposite of the gap and the probability of price decreases.

The covariance term is a potential measure of the state-dependence in price setting and expresses the extent to which prices with larger misalignments are adjusted with higher probability. The term disappears in constant-hazard time-dependent models (Calvo, 1983). It is straightforward to see why: in these models $\Lambda(x)$ is constant; therefore $\left(\Lambda(x)-\bar{\Lambda}^-\right)=0$ for each x. In contrast, the covariance term is necessarily non-zero (negative for positive gaps, and positive for negative gaps) in state-dependent price-setting models with increasing hazard: The larger the absolute value of the gap, the higher the probability of price adjustment.

2.2 Adjustment to shocks: Narrow selection and the gross extensive margin

Given this decomposition of inflation, one can decompose the inflation *response* to aggregate shocks into several economically meaningful adjustment margins that we connect to selection. We quantify these margins later and use the results to discriminate among models of price setting. To start, we follow Caballero and Engel (2007) in measuring price flexibility by focusing on the impact of a marginal aggregate shock on inflation,

$$\frac{\partial \pi}{\partial m} = F(0)\frac{\partial \pi^+}{\partial m} + \left[1 - F(0)\right]\frac{\partial \pi^-}{\partial m} \tag{5}$$

where π^+ and π^- are inflation among the group of products with negative and positive gaps, respectively.[6]

Then, from equation (4), the marginal impact of an aggregate shock on the positive-gap group becomes

$$\begin{aligned}
\frac{\partial \pi^-}{\partial m} &= \underbrace{\bar{\Lambda}^-}_{\text{intensive}} + \overbrace{\underbrace{-\bar{x}^-\frac{\partial \bar{\Lambda}^-}{\partial m}}_{\text{gross extensive}} + \underbrace{\int_{x\geq 0} -x\left(\frac{\partial \Lambda(x)}{\partial m} - \frac{\partial \bar{\Lambda}^-}{\partial m}\right)h(x)dx}_{\text{narrow selection}}}^{\text{broad selection / extensive margin}} \\
&= \underbrace{\bar{\Lambda}^-}_{\text{intensive}} + \overbrace{\underbrace{\bar{x}^-\int_{x\geq 0}\Lambda'(x)h(x)dx}_{\text{gross extensive}} + \underbrace{\int_{x\geq 0} x\left(\Lambda'(x) - \int_{x\geq 0}\Lambda'(x)h(x)dx\right)h(x)dx}_{\text{narrow selection}}}^{\text{broad selection / extensive margin}} \qquad (6) \\
&= \underbrace{\bar{\Lambda}^-}_{\text{intensive}} + \overbrace{\underbrace{\bar{x}^- E[\Lambda'(x)|x\geq 0]}_{\text{gross extensive}} + \underbrace{cov(\Lambda'(x), x|x\geq 0)}_{\text{narrow selection}}}^{\text{broad selection / extensive margin}} \qquad (7)
\end{aligned}$$

The second line in the equation uses the insight derived in Caballero and Engel (2007) that $\partial\Lambda(x)/\partial m = -\Lambda'(x)$ for all x where Λ is differentiable. The result holds because the aggregate shock shifts all the gaps uniformly by one unit. Consequently, $\partial\bar{\Lambda}^-/\partial m = -\int_{x\geq 0}\Lambda'(x)h(x)dx$. In words, the marginal change in the average frequency of price-decreases equals the average slope of the hazard function over the positive price gaps. The expression is analogous in the negative-gap group, and the margins of adjustment for the overall inflation rate are the sums of the respective margins in the negative-gap and the positive-gap inflation rates weighted by their respective shares.

The expression needs to be adjusted in case $\Lambda(x)$ is non-differentiable at the inaction thresholds, as is the case in standard (S,s) models (Caballero and Engel, 2007). If $\Lambda(x)$ shifts from 0 to 1 at point x^u, the

[6] Notably, we do not allow the aggregate shock to affect the threshold ($x = 0$) between the groups with negative and positive gaps; that is, we allocate products into groups based on *ex-ante* gaps, which exclude the impact of the aggregate shock. We do this for multiple reasons. First, it simplifies the algebra because we do not need to worry about any shift between negative- and positive-gap groups caused by the aggregate shock. Second, as we further clarify below, using our definition, the extensive margin stays inactive in the time-dependent Calvo (1983) model. This is advantageous because we stay consistent with previous definitions (e.g., Caballero and Engel, 2007), which distinguish state-dependent models from time-dependent models through the existence of an active extensive-margin effect.

decomposition becomes

$$\frac{\partial \pi^-}{\partial m} = \underbrace{\bar{\Lambda}^-}_{\text{intensive}} + \overbrace{\underbrace{\bar{x}^- h(x^u)}_{\text{gross extensive}} + \underbrace{\left(x^u - \bar{x}^-\right) h(x^u)}_{\text{narrow selection}}}^{\text{broad selection / extensive margin}}. \tag{8}$$

It is straightforward to see that in such an (S,s) model, selection is generically non-zero. It is positive when the average positive gap is strictly below the upper inaction threshold (and symmetrically the average negative gap is above the lower inaction threshold), which holds in standard menu cost models (Golosov and Lucas, 2007).

The expression decomposes the impact of the shock into three terms. The first is the intensive-margin effect, which is driven by the increase in gaps brought about by the aggregate shock uniformly over the whole distribution. Its magnitude takes a particularly easy form: it equals the average frequency of price adjustment in the positive-gap group ($\bar{\Lambda}^-$). This is intuitive: all firms that would have changed their prices now change their prices marginally more by a term $-\partial \bar{x}/\partial m = 1$. Overall, the intensive-margin effect $\bar{\Lambda} = F(0)\,\bar{\Lambda}^+ + \left[1 - F(0)\right]\bar{\Lambda}^-$, is equal to the frequency of price changes (Caballero and Engel, 2007).

The intensive margin is the only margin of adjustment present in constant-hazard time-dependent models (e.g., Calvo, 1983). Consequently, the complementary second and the third adjustment margins are active in state-dependent models. We refer to them jointly as the "broad selection" effect or the "total extensive-margin" effect (Caballero and Engel, 2007). Our analysis separates this effect into two adjustment sub-margins: the "gross extensive margin" effect and the "narrow selection" effect. We argue that they are conceptually different and measuring them separately can help us distinguish between state-dependent pricing models with very different aggregate implications, as illustrated in our subsequent analysis.

Next to the intensive-margin term, the gross-extensive-margin effect in our decomposition captures the impact of the shock on the relative share of price changes between the positive-gap and the negative-gap groups. A policy tightening reduces the share of price increases, and increases the share of price decreases. The strength of the effect depends on the average size of the gap and the average impact on the frequency of price changes in the positive-gap group. This latter term is positive within the relevant class of models with the increasing hazard property, so the gross-extensive-margin effect will amplify the impact of the aggregate shock. The overall impact of the gross extensive margin also leads to amplification: a tightening leads to fewer price increases and more price decreases. Formally, $F(0)\,(-\bar{x}^+)\,\partial\bar{\Lambda}^+/\partial m + \left[1 - F(0)\right](-\bar{x}^-)\,\partial\bar{\Lambda}^-/\partial m \le 0$. In a constant hazard model $\partial\bar{\Lambda}^-/\partial m = \int_{x\le 0}\Lambda'(x)h(x) = 0$ as well as $\partial\bar{\Lambda}^+/\partial m = \int_{x\ge 0}\Lambda'(x)h(x) = 0$ so the gross extensive margin is absent.

The third term is what we call the narrow selection effect. It measures the contribution arising from any shift in the distribution caused by the unusual "selected" position of new price changes. Mathematically, as expressed by equation (7) above , it is given by the covariance of the marginal probability of adjustment in response to a shock and the size of the new price adjustments. Notably, what thus matters for narrow selection is not whether prices far from their optimal levels are *regularly* adjusted with higher probability (state dependence), but rather the size of the *extra* price changes that are triggered or canceled by the aggregate shock.

Figure 2 illustrates this point through the example of a policy tightening. Tightening reduces the optimal price of each firm, raising each firm's gap, or, equivalently, shifts the adjustment hazard to the left as

Figure 2: Selection

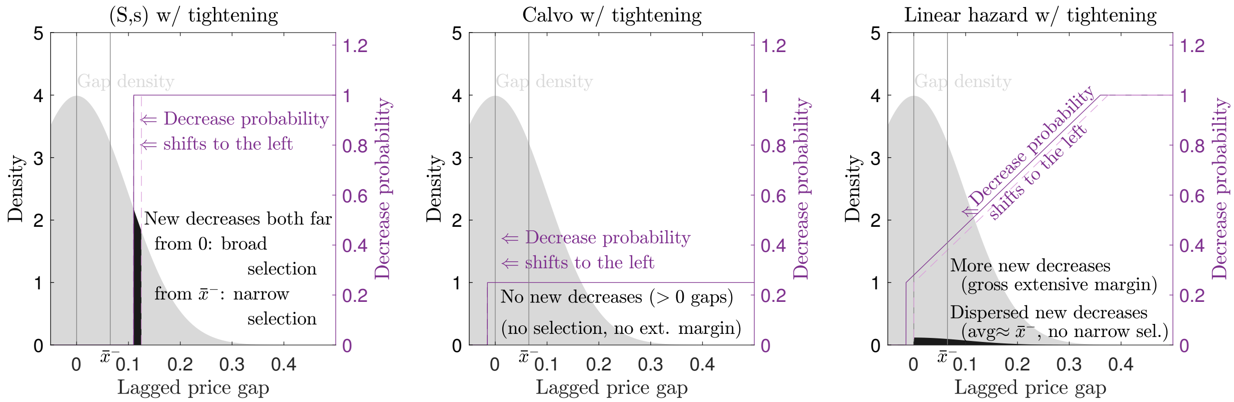

Note: The figure shows the density of the ex-ante price gaps, the price-decrease hazards and the distribution of new price changes as a function of gaps in a state-dependent menu cost model (left panel), a time-dependent model (middle panel) and a linear hazard model (right panel). The figure illustrates the differences between the adjustment margins in the frameworks. First, the distribution of ex-ante price changes do not change in the time-dependent model, so both gross extensive margin and narrow selection are absent. Second, there are new price decreases in both state-dependent models ((S,s) and linear hazard) so the gross extensive margin is active in both cases. Third, in the (S,s) model adjusting prices are far from the average gap (implying large narrow selection), while they are at the average gap in the linear hazard model (implying no narrow selection).

a function of the *ex-ante* gap that makes up the horizontal axis. The black shaded areas on the figures show the new adjusters among the group with positive gaps. The outcome is very different in the three pricing models. In the time-dependent model (middle panel), there are no new adjusters within the relevant group, both the gross-extensive-margin and the selection effects are missing. By contrast, in the (S,s) state-dependent model (left panel), both channels are present. First, there are new adjusters contributing to the gross extensive margin. Second, the new adjusters have large gaps; therefore, they are far both from their optimal prices and from the average gap when the aggregate shock hits. Consequently, when they decrease their prices, they will decrease them by a lot so as to release their accumulated price pressures. Thus, the idiosyncratic price pressures amplify the impact of the aggregate shock and make the aggregate price level more flexible; that is, narrow selection is high. In the linear hazard model (third panel), only the gross-extensive margin is present. It is there, because, similarly to the (S,s) model, the aggregate shock generates new decreasers. However, differently from the (S,s) model, narrow selection is absent. The reason is that new decreasers are a representative subset of the population with positive gaps. Therefore, they share the same average gap. [7] In Section 4.4 we quantify the relative contributions of the different channels in calibrated

[7] The example in the middle panel of Figure 2 can help clarify some of the implications of allocating prices into groups based on ex-ante gaps in our decomposition. As the tightening shifts the price-decrease hazard to the left, there will be new price decreases among products with marginally *negative* ex-ante gaps (those whose ex-post gap became positive due to the impact of the aggregate shock). According to our definition, which restricts attention to positive ex-ante gaps, the impact of these new decreases does not influence the gross extensive margin; instead, they are allocated to the intensive-margin effect. Why? Because the shift from price increases toward price decreases in this constant-hazard Calvo (1983) model is just a mechanical side effect of the adjustment on the intensive margin, whereby each adjusting firm reduces its prices by marginally more, and, therefore, some firms that would have increased their prices are now going to decrease them. Importantly, grouping based on ex-ante gaps has no influence on the strength of selection (only the share between the intensive- and the gross-extensive-margin effect): the new adjusters after a marginal aggregate shock have zero gaps, and therefore zero impact on selection, which is defined as the product of the gap size (zero) and the mass of new adjusters. Another feature of the definition is that the observable shift between price increases and price decreases after an aggregate shock cannot be directly used to assess the

quantitative price-setting models and confirm the qualitative conclusions derived above. But first, we turn to measure selection in the data.

3 Data

Our analysis builds on two data sets. The main data set is a U.S. supermarket scanner data set. A secondary data set is the micro price data underlying the calculation of the U.S. producer-price index at the Bureau of Labor Statistics.

The supermarket scanner data set is collected by marketing company IRi (Bronnenberg et al., 2008), and includes weekly total sales ($) as well as quantities sold at the barcode level for 31 food and health-care product categories (for example, carbonated beverages) in major supermarkets across 50 U.S. metropolitan areas between 2001-2012. It covers 15% of the consumer expenditure survey.[8]

Given total sales revenue and quantities, we calculate posted prices in the IRi data as $P_{psw} = \frac{TR_{psw}}{Q_{psw}}$, where TR is the total revenue and Q is the quantity sold for each product p in store s in week w. We conduct some straightforward data cleaning,[9] and conduct two further steps to construct the price series for our analysis. First, we use the modal-price filter of Kehoe and Midrigan (2015) to construct a weekly series of reference prices, P^f_{psw}, that is insulated from the impact of temporary sales. Our subsequent analysis concentrates on reference prices but our results are robust to the inclusion of sales data.[10] Second, we construct monthly observations from the weekly data set. This step helps us concentrate on lower frequency developments in prices, which are more relevant for business cycle fluctuations. The monthly price P_{pst} is defined as the mode of the weekly prices over a calendar month, choosing the highest if there are multiple modes. Picking the mode makes sure we do not create artificial prices and price changes, as could happen if we used averages. Monthly expenditure is the sum of weekly sales.

To cross check the external validity of our scanner data, we also calculate an expenditure-weighted supermarket price index. Our aim is to create a chain-weighted index similar to the consumer price index, but utilize the wider and deeper information available in our data set. In particular, we use annual revenue weights $\omega_{psy} = TR_{psy}/\sum_p \sum_s TR_{psy}$, where subscript y reflects the year of the observation month. We measure inflation for posted and reference prices ($i = p, f$) as $\pi^i_t = \sum_s \sum_p \omega_{pst} \left(p^i_{pst} - p^i_{pst-1}\right)$, where $p^i_{pst} = \log P^i_{pst}$. We define sales-price inflation as the difference between posted- and reference-price inflation

strength of the gross extensive margin. The reason is that the change in the observable share of price increases versus price decreases is contaminated by the intensive-margin channel. Therefore, the measurement of the gross-extensive-margin effect requires an indicator of the gap, as in our empirical implementation below.

[8] The number of observations in our sample is almost 2.65 billion, spread over 168,000 unique products in 3187 unique stores in 169 supermarket chains.

[9] First, prices are rounded to the nearest penny, as fractional prices reflect the impact of promotional sales during the week, not actual posted prices. (The rounding influences 9% of the prices in our sample.) Second, the product identification numbers of *private-label* products are not unique throughout our sample. In particular, the identification number changes for a subset of products in January 2007, January 2008, and January 2012. When constructing price spells, we follow a conservative approach and assume that all private-label goods were replaced with new goods on these dates. We disregard the price and expenditure changes for these goods on these three dates. Third, in each year, our data only include products that are available over the whole year. This way, our analysis excludes entering and exiting products, which might exhibit idiosyncratic pricing behavior, for example, motivated by learning about the products' demand function at introduction (Argente and Yeh, 2022) or during a clearance sale at exit. These idiosyncratic factors would further reduce selection, so excluding these prices is a conservative choice, raising the probability that we would find selection in the data. This step makes us drop 17% of the products (18% of annual expenditure).

[10] To construct reference prices, we start with a running 13-week two-sided modal price. Then, we iteratively update it to align the timing of reference-price changes with the actual price changes. Online Appendix D describes details of the algorithm.

$\pi_t^s = \pi_t^p - \pi_t^f$. We seasonally adjust the series using (12) monthly dummies.

Figure 3: Posted, reference, and sales-price inflations (year-on-year)

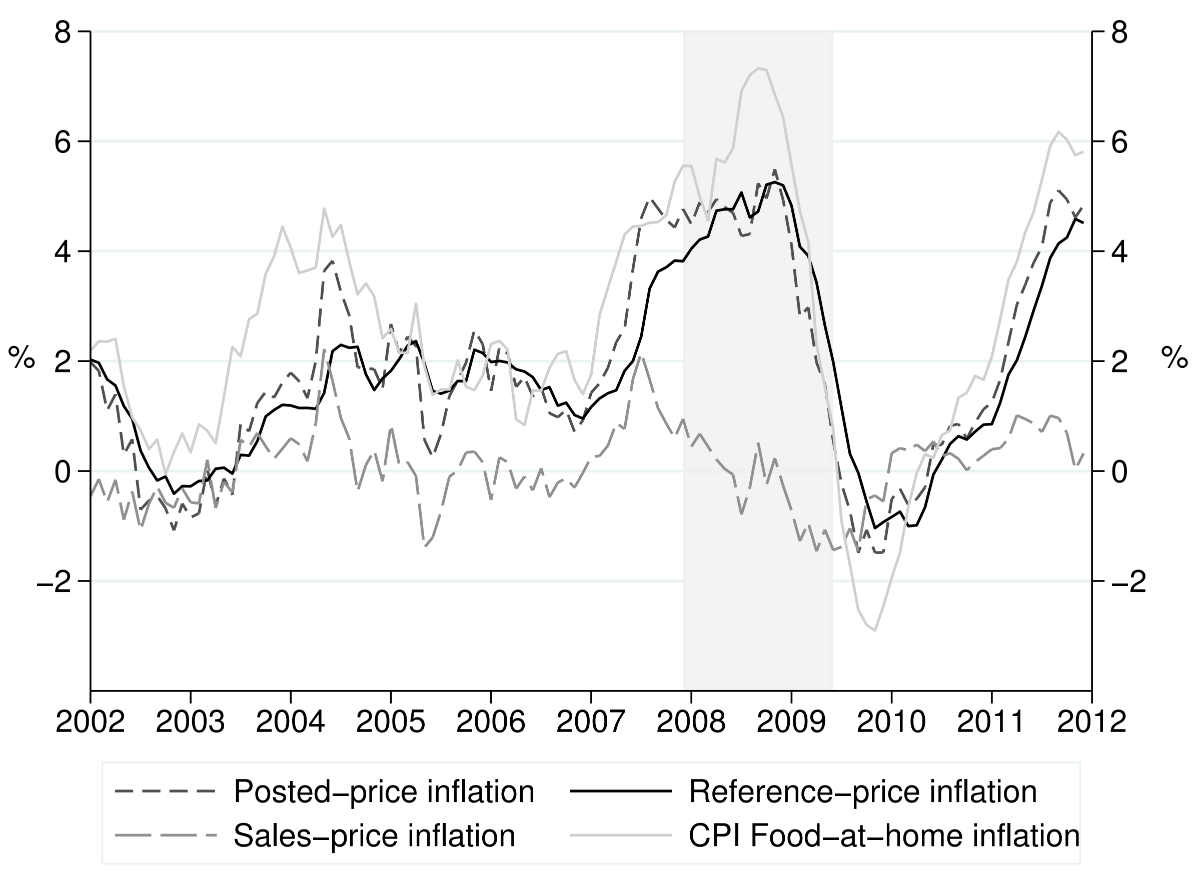

Note: The figure depicts the evolution of year-on-year posted-, reference- and sales-price inflation rates over our sample period, comparing them to the food-at-home subindex of the CPI. The reference-price inflation, which is used in the preceding analysis, closely tracks the food-at-home inflation.

Figure 3 shows the evolution of the year-on-year posted-, reference- and sales-price inflations ($\pi_t^i = \log P_t^i - \log P_{t-12}^i, i = p, f, s$) and compares them to the evolution of the year-on-year changes in the food-at-home consumer-price subindex published by the Bureau of Labor Statistics. We see that the fluctuations in the price index calculated from our sample closely track those of the food-at-home index.[11]

Table 1 lists some unconditional moments of the (sales-filtered) price-change distribution for our supermarket prices. We can observe inflation of around 1.8% in our sample. The inflation rate in reference prices is very close to that of the posted prices, 1.75% versus 1.84%, while average sales-price inflation is only 0.05%

[11] Our inflation measure somewhat underestimates the overall food-at-home inflation over the period. The annualized inflation in posted prices in 1.84%, and in reference prices it is 1.75%. The CPI food-at-home inflation over the same period is 2.7%. The main reason is that our index relies on the price development of existing products and – differently from the CPI – it ignores the endogenous replacement of old, exiting products with new, entering products, which tend to be more expensive and of higher quality.

and not shown. Posted prices change very frequently: every month more than one-third of them change.[12] But the majority of price changes are temporary and only 11% of reference prices change each month. Out of these reference-price changes, 6.6% are increases and 4.2% are decreases. Price changes are large when they occur: the average increases are 12.5%, while the average decreases are 15.1%.[13] The kurtosis of the standardized reference price change distribution is 2.7, below the kurtosis of the Gaussian distribution.[14]

Table 1: Average moments

Annualized inflation		Frequency		Reference frequency		Reference size		Standardized reference	
Posted	Reference	Posted	Reference	Increase	Decrease	Increase	Decrease	Skewness	Kurtosis
1.84%	1.75%	36.2%	10.8%	6.6%	4.2%	12.5%	-15.1%	-0.13	2.71

Note: The table lists some relevant moments of posted and reference prices. It confirms that posted price changes are very frequent, but they are mostly driven by temporary sales. The probability of reference-price changes in a month is around 11%. The size of a non-zero reference-price change is large, over 10%, and the kurtosis of the standardized non-zero reference-price change is 2.7, below the Gaussian kurtosis.

Our second main data set is the micro price data set underlying the construction of the U.S. PPI. These data have been described in detail in several papers that analyze the PPI micro price data, such as Nakamura and Steinsson (2008); Goldberg and Hellerstein (2009); Bhattarai and Schoenle (2014); Gorodnichenko and Weber (2016); Gilchrist et al. (2017). For details on the PPI data we refer the reader to these papers; we discuss below some of the key features and how they contribute to the robustness of our results.

The PPI data are a monthly data set of transaction-based prices. Goods in manufacturing and services are the focus of the PPI data. The data set contains prices for approximately 28,000 firms and more than 100,000 goods every month, belonging to 540 six-digit NAICS product categories. These NAICS categories are very narrow categories, such as "Tortilla Manufacturing" (NAICS code 311830). Within this setting, goods produced by each firm are uniquely identified according to their "price-determining" characteristics. These characteristics include the type of buyer, the type of market transaction, the method of shipment, the size and units of shipment, the freight type, and the day of the month of the transaction. Goods remain in the data on average for 70 months, while firms exist for approximately 7 years. Sales are very rare in the PPI data, as documented by Nakamura and Steinsson (2008), so we work with unfiltered data. However, we check for the importance of product substitutions. Whenever we identify a rare product substitution through a so-called base price change, we assume a price has changed.

The main appeal of the PPI microdata is that (i) the data aim to map out the "entire marketed output of U.S. producers" in a representative fashion (BLS, 2020). Such producer data rather than retail data are

[12]The frequency of increases and decreases between months $t-1$ and t are defined as

$$\xi^{\pm}_{t-1,t} = \sum_i \bar{\omega}_{it} I^{\pm}_{it-1,t}, \tag{9}$$

where $I^{+}_{it-1,t}$ and $I^{-}_{it-1,t}$ are indicators that take the value 1 if the price of item i (a product in a particular store) increased or decreased, respectively, between month $t-1$ and month t and 0 otherwise. The weight $\bar{\omega}_{it}$ is measured as the annual expenditure share of item i in the calendar year of t.

[13]The average size of price increases and decreases is defined as

$$\psi^{\pm}_{t-1,t} = \frac{\sum_i \bar{\omega}_{it} I^{\pm}_{it-1,t}(p_{it} - p_{it-1})}{\xi^{\pm}_{t-1,t}}. \tag{10}$$

[14] We standardize reference-price changes at the item level and only include items with at least 5 reference-price changes over the price spell. The restriction is valid for 45% of our sample. We calculate the kurtosis using annual expenditure weights.

particularly suitable for studying how all kinds of firms set prices and how they react to aggregate shocks; (ii) the data range from 1981 to 2015, providing a long time series that spans multiple business cycles and provides variation in monetary policy and credit frictions; and (iii) the producer-price data cover a wide range of sectors in manufacturing but also three-quarters of services in the U.S. economy, lending more general validity to our analysis.

4 Empirical hazard function and selection

This section generates empirical hazard functions and gap densities from the micro price data. We then use them to obtain estimates of the selection effect, which imposes discipline on the choice of pricing models.

4.1 A price-gap proxy

Estimating the hazard function and the gap density requires a proxy for the price gap $x_t = p_t - p_t^*$, the distance of the price from its unobserved optimal level. In a wide class of state-dependent models, the price gap is the relevant idiosyncratic state variable driving the incentive to adjust prices (Golosov and Lucas, 2007; Alvarez et al., 2016). The further a price is from its optimum, the stronger is the incentive to adjust it.

Our baseline price-gap proxy is the competitor-price gap, which measures the distance of the price from a suitably adjusted average price of close competitors. In Appendix **??**, we show that our results are robust to using two alternative gap proxies: the competitor-*reset*-price gap, which measures the distance from the average price of those competitors that changed their prices in the particular month. and the reset-price gap, which measures the distance of a price from its own reset price (Bils et al., 2012). We first lay out the details of the construction of the competitor-price gap.

One of the primary concerns of firms in making their price-setting decisions is figuring out how far the price of an item is from the average price of their competitors. Our data allow us to answer this question at barcode-level granularity after taking into account detailed fixed effects. Formally, we define the competitor-price gap for product p in store s in month t as $\tilde{x}_{pst} = p^f_{pst} - \bar{p}^f_{pt}$, where p^f_{pst} is the logarithm of the reference price and $\bar{p}^f_{pt}$ is the average reference price of the same product across stores excluding store s. We deal with persistent heterogeneity across stores (i.e., chains, locations) by subtracting the average store-level gap α_s and reformulating the price gap as $x_{pst} = \tilde{x}_{pst} - \alpha_s$.[15]

Using the competitor-price gap as our baseline price-gap proxy is motivated by models with strategic complementarities in price setting. In models with strategic complementarities, the average price of competitors affects the optimal reset price over and above the nominal marginal cost (Ball and Romer, 1990; Woodford, 2003). Strategic complementarities are a relevant feature of micro-level price setting (Gopinath and Itskhoki, 2010; Beck and Lein, 2020), and an important factor in explaining the macro-level estimates of monetary non-neutrality (Woodford, 2003; Gertler and Leahy, 2008; Nakamura and Steinsson, 2010). Many frameworks may justify our measure and they will share as their key modeling ingredients (i) an oligopolistic market structure between price-setters (Ellickson, 2013) and (ii) price-adjustment frictions. The first ingre-

[15] In our sample, the average product is sold in 104 stores in an average month. A large number of products, however, are sold in only a few stores. Table **??** in the appendix shows that our results are robust if we restrict our sample to product-months for which 50 stores are available.

dient introduces complementarity with competitor prices into the optimal price of a firm (see, for example, Atkeson and Burstein, 2008; Auer and Schoenle, 2016; Pennings, 2017). The second ingredient, pricing frictions, introduces a dynamic aspect that in particular hampers future competitor price cuts (as analyzed, for example, in Mongey, 2021). As a result, they reinforce the incentive of price-setting firms to move toward their competitors' prices to raise suboptimal markups (if below) or increase suboptimal market share (if above). Campbell and Eden (2014) have effectively used competitor prices in their empirical analysis aimed at distinguishing time- and state-dependent pricing elements. Appendix **??** shows that our results are robust to using an alternative proxy, the competitor-reset price, which is valid under even more general conditions.

4.2 Empirical hazard function and price-gap density

Figure 4: The size and frequency of subsequent reference-price changes as a function of the competitor-price gap and its density

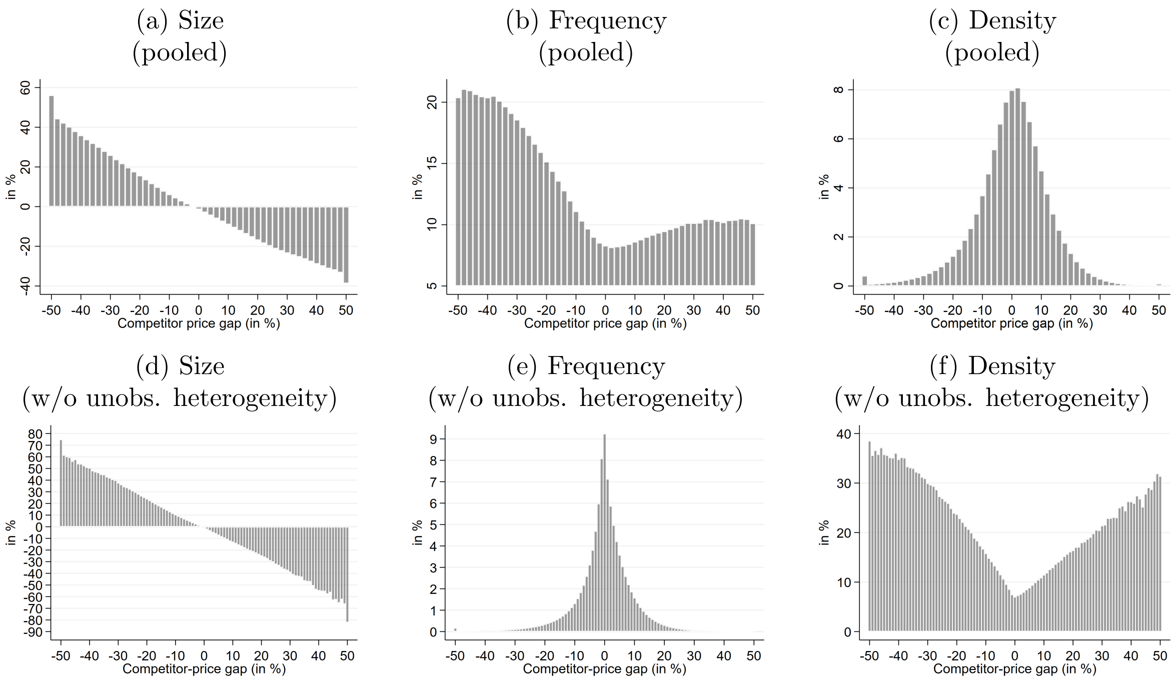

Note: The panels in the first row show the frequency and size responses of subsequent price change as a function of the price gap and the price-gap density in the baseline IRi supermarket data set pooled across products and time. The second row shows the same moments after controlling for unobserved heterogeneity across products, stores, and time. The figures show that (i) the gap is a relevant proxy as the size of average subsequent adjustments has a very tight, close to (minus) one-to-one relationship with the gap (first column), (ii) the frequency of subsequent price adjustment increases with the absolute size of the gap (second column) and (iii) the density of the competitor-price gaps has fat tails (third column). The frequency of adjustment as a function of the gap (generalized hazard function) is close to (piecewise) linear, especially after controlling for unobserved heterogeneity (panel e).

What do our key empirical objects look like? Figure 4 shows (i) the size and (ii) the probability of price adjustment as a function of the price gap and (iii) the density of the competitor-price gap x_{pst} in two ways. In the first row of panels, moments are calculated after simply pooling the data across products and time. The second row takes into account unobserved heterogeneity across products, stores, and time by deducting estimated product-store and time fixed effects from each respective moment y_{pst}.[16] The associated moments now reflect variation within homogeneous groups and are much closer to structural moments used in standard theoretical frameworks.

The figures in the first column of Figure 4 show the average size of non-zero reference-price adjustments conditional on the lagged competitor-price gap. They reveal a tight, nearly (negative) one-to-one relationship between the gap and the reference-price change. This tight relationship validates our proxy by confirming that distance from competitor prices is indeed a relevant component of the theoretical price gap, and actual reference-price changes aim to close this gap, on average. The second column of Figure 4 shows the probability of price changes as a function of the price gap, the generalized price-adjustment hazard. The probability is increasing with the distance away from zero in a reverse J-shaped pattern that is consistent with state-dependent pricing models. This relationship suggests again that our proxy indeed captures an important component of the unobserved theoretical price gap (as argued also by Gagnon et al., 2012; Campbell and Eden, 2014).

As a novel contribution, the second panel in the bottom row shows the empirical hazard function after taking into account unobserved heterogeneity. Such an exercise reveals important features of the structural hazard function, which are less apparent using pooled data. First, the hazard function is close to piecewise linear, especially in the relevant range (between -20% and +20%, where 95% of the mass of price gaps is concentrated). Second, the hazard function is much steeper than the pooled data would suggest (e.g., around 30% instead of 10% for a gap of over 50%). Still, the hazard function stays quite flat, as the probability of adjustment stays well below 50% even for gaps of 50%. Third, the asymmetry between the probability of adjustment with a negative gap versus a positive gap is still present, but smaller than in the pooled data.[17] Finally, the densities in the third column show a sizable dispersion of price gaps even after we control for temporary sales and permanent differences in the price level of stores with fat tails.

4.3 Selection

The equations derived in Section 2.2 allow us to estimate the strength of the selection effect relying on the empirical hazard and density function estimates. We do so in two ways. First, in this section, we simply assume that the empirical moments represent their theoretical counterparts.[18] We then compute the importance of the three adjustment margins based on this assumption. Second, in the next section, we build a flexible quantitative price-setting framework and estimate its parameters so as to match the empirical counterparts of both the hazard function and the price-gap density. Based on the theoretical gap density and the hazard function, we again repeat our decomposition exercise and benchmark it against counterfactuals given by other price adjustment models.

[16] In particular, we first take the residual of the panel regressions $y_{pst} = \alpha_{ps} + \alpha_t$, where α_{ps} is a product-store and α_t is a time fixed effect. Then, second, we increase each residual by the average fixed effects ($\bar{\alpha}_{ps} + \bar{\alpha}_t$) across all observations. This makes sure that the mean of the filtered moment stays equal to the mean of the original. y is, respectively, the gap x_{pst} for the density, $I_{pst,t+1}$ for frequency, where $I_{pst,t+1}$ is an indicator function that takes a value 1 if product p in store s changed from month t to $t+1$, and $\Delta p_{pst,t+1}$ for size, where $\Delta p_{pst,t+1}$ is the size of a non-zero price change of product p in store s between month t and $t+1$.

[17] Some features (below 1 probability at high gaps, asymmetry, and positive hazard at 0) are in line, but some others (piecewise linearity) are not in line with the results of Luo and Villar (2021), who estimate flexible hazard functions by matching price-setting moments of U.S. consumer-price inflation.

[18] Our proxy measures only a component of the theoretical price gap. The assumption requires that the evolution of this component and the probability of the price adjustment as a function of this component both respectively represent the evolution of and the probability of adjustment to the theoretical price gap, on average.

Table 2: Relative strength of the adjustment channels based on empirical moments

	Intensive margin	Gross extensive margin	Selection effect
Relative contributions	73.4%	26.5%	0.2%

Note: The table presents the relative contributions of the various adjustment margins. It uses the formulas described in Section 2 and the empirical estimates of the hazard function, and density as presented in the second row of Figure 4.

We decompose the inflation response to a permanent money shock into the intensive, gross-extensive-margin, and the narrow selection effects under the assumptions outlined in Section 2. Table 2 shows the result of the empirical decomposition: The columns contain the relative contributions of the adjustment channels to the overall immediate impact on inflation of a marginal permanent money shock. As Auclert et al. (2022) argue, the immediate effect is informative about the full impulse-response function in a wide class of discrete-time state-dependent pricing models (with the dynamics close to equivalent to a Calvo (1983) model up to a single scaling factor). We conduct the underlying computation using the same level of discretization as presented in panels (e) and (f) of Figure 4, calculating the derivative of the hazard function over each grid-point as a centered finite difference approximation.

The table shows that the gross extensive margin increases aggregate price flexibility noticeably, at a magnitude that is around one-third of the intensive-margin effect, the latter being the only one present in the time-dependent Calvo (1983) model. Additionally, the impact of the narrow selection effect is minuscule: two orders of magnitude smaller than the gross-extensive-margin effect. The results imply monetary non-neutrality that is only 36% milder than that in the time-dependent Calvo (1983) model. This is very far from the state-dependent model of Golosov and Lucas (2007) with strong selection (Alvarez et al., 2021) as we show next.

4.4 A structural analysis

This section presents a flexible quantitative price-setting model and estimates its parameters so as to match the empirical hazard function with its model counterpart. The model fits the price-change distribution and the price-gap density well, even though they are not directly targeted. Crucially, we use the model to calculate the theoretical decomposition of the intensive, gross-extensive, and narrow selection margins and show that the model-based estimates are close to those using the empirical moments calculated in the previous section. We also show that narrow selection would be sizable in a counterfactual exercise assuming a fixed menu cost model (Golosov and Lucas, 2007). This finding suggests that focusing on its measurement to distinguish between different state-dependent models, as we do in our paper, is quantitatively meaningful.

4.4.1 The model

The model is a standard quantitative dynamic stochastic general equilibrium price-setting model with random price-adjustment costs (Dotsey et al., 1999; Alvarez et al., 2021). The framework is flexible and nests the time-dependent Calvo (1983) model and the fixed menu cost model of Golosov and Lucas (2007) as special cases, as well as a continuum of intermediate cases. We describe the key features of the model here and direct the interested reader to Costain and Nakov (2011a) for details and derivations.[19]

In this setup, there is a continuum of differentiated goods (i), which are sold in a market with monopolistic competition. This market structure gives the producer of each good market power to set prices at a markup above marginal cost. The market power is determined by the elasticity of demand, which, in turn, is governed by the (constant) elasticity of the substitution parameter ε.

Production requires labor, and the product-specific productivity is subject to idiosyncratic shocks. As argued by Golosov and Lucas (2007), these shocks are necessary to explain the large absolute size of price changes. We assume that productivity follows a random walk, with an idiosyncratic shock $z_t(i)$ with standard deviation σ_z ($A_t(i) = A_{t-1}(i) + z_t(i), z_t(i) \sim N(0, \sigma_z^2)$) (this is different from Costain and Nakov, 2011a, who assume persistent, but not permanent shocks). All the relevant firm-level information is incorporated into the price gap, defined as the distance of its (log) price from its (log) optimal price $x_t(i) = p_t(i) - p_t^*(i)$. In particular, its expected present discounted profit is a function of the price gap, and it is maximized when the price gap is zero. The price gap fluctuates as idiosyncratic shocks hit the optimal price, and the firm does not necessarily reset it to zero because adjusting the product price ($p_t(i)$) is costly.

Firms face price-adjustment costs, κ, which are drawn every period independently from the distribution $G(\kappa)$. If the firm is willing to pay the cost, it can reset its price and close its price gap. As in Costain and Nakov (2011a), we use the generalized hazard function, instead of the distribution of the menu costs as a primitive of the model. This choice comes without loss of generality assuming a non-decreasing hazard function.[20] In our baseline model, we assume that the hazard function is the form expressed by equation (1), i.e., it is piecewise linear, with potentially different slopes over negative (c) and positive gaps (b) and a potentially positive intercept (a). We contrast our baseline with two key alternatives: the time-dependent Calvo (1983) case when the hazard function is a constant; and the strongly state-dependent fixed menu cost case of Golosov and Lucas (2007), where the hazard function is a step function.

4.4.2 Estimation

This section estimates key structural parameters in the model by matching the empirical hazard function as well as the frequency and the size of price changes with their model counterparts. Our analysis also sets some parameters to levels used in the literature following Woodford (2009), with one difference. We set the elasticity of the substitution parameter (ε) to 3. This is the parameter used by Midrigan (2011), and it implies markup levels relevant for supermarkets.

The parameters we estimate are (i) the standard deviation of the idiosyncratic shocks (σ_z), and the three parameters of the linear hazard function: the slopes over the negative and positive gaps (b, c) and its intercept a. We estimate these parameters by targeting three empirical moments by their simulated counterparts in the stationary equilibrium: the shape of the generalized hazard,[21] and the frequency and

[19] We thank Anton Nakov for posting his code

[20] As has been shown by Caballero and Engel (2007) and Alvarez et al. (2021), all non-decreasing hazard functions can be an outcome of a suitably chosen menu cost distribution. We only consider non-decreasing hazards.

[21] The estimation algorithm minimizes the squared difference between the empirical and theoretical hazard functions, weighted by the price-gap density. The empirical hazard, calculated as the probability of a price change *next month* as a function of the current price gap, is matched with a simulated hazard, which similarly expresses the probability of a price change in the next month as a function of the current gap.

size of the price changes.[22] We also check how the model matches some untargeted moments, such as the standardized price-change distribution and price-gap density.

4.4.3 Results

The estimation produces a tight fit of the targeted moments, but also an acceptable fit of untargeted moments as Figure 5 illustrates. The left panel shows that the linear hazard model matches the targeted shape of the empirical next-period generalized hazard very closely, and the fit is especially tight in the region where most of the mass is concentrated. Notably, the framework attributes the smoothness of the empirical hazard around zero gap to idiosyncratic shocks hitting between period $t-1$, when the price gap is measured, and period t when the probability of adjustment is measured (the theoretical hazard has a sharp kink at 0). As Table 3 shows, the theoretical hazard assigns a positive probability of adjustment at zero gap (positive intercept a) and its slope is higher for negative gaps (over 80%) than for positive gaps (below 50%).

Figure 5: The matched hazard function and the untargeted standardized price-change distribution

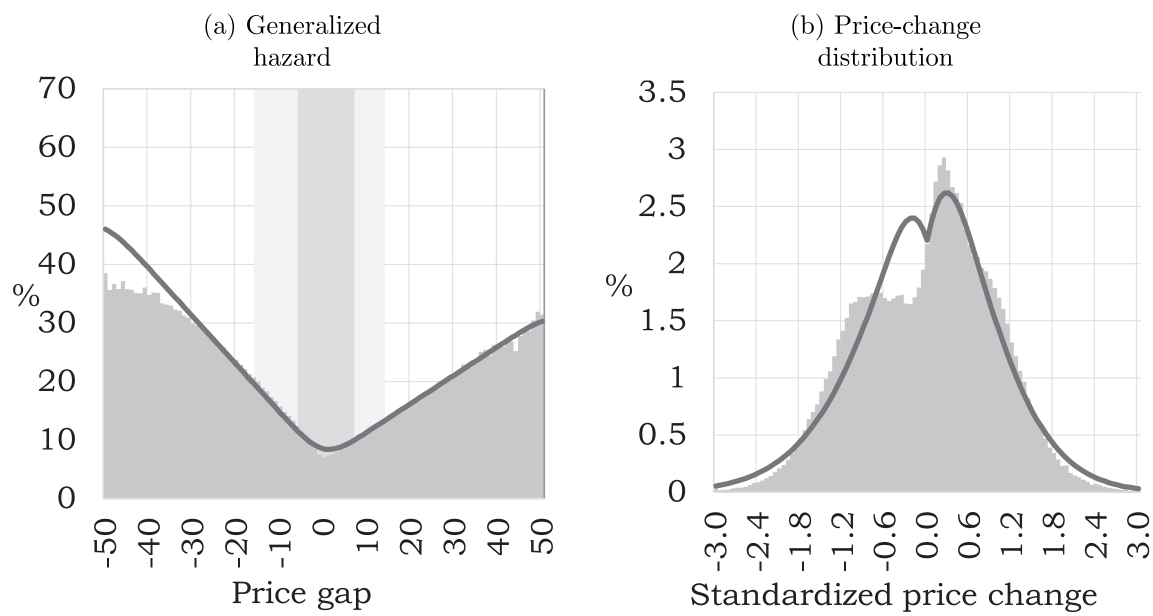

Note: The figure shows the fit of the baseline linear hazard model. The left panel shows that it is able to match the targeted generalized hazard very closely, especially in the region where most of the mass resides (67%: dark shaded area, 90%: light shaded area). The right panel shows that the model captures well the shape of the standardized reference-price change distribution: it is bimodal, asymmetric and has fat tails.

The right panel of Figure 5 shows that the model fits the untargeted standardized[23] reference-price change distribution well. It matches the (i) bimodality, (ii) the asymmetry and the (iii) fat tails of the distribution

[22]For internal consistency of our quantitative exercise, the frequency and size measures we match here are derived from (unweighted, truncated at $+-50\%$) generalized hazard and density estimates. In particular, frequency is measured as $\sum_j \Lambda_j f_j$, and size as $\sum_j |x_j|\Lambda_j f_j / \sum_j \Lambda_j f_j$, where Λ_j is the height of the generalized hazard, f_j is the relative share of products in the price-gap bin j, and x_j is its midpoint. These measures are not equal to the (weighted, untruncated) frequency (13.8% versus 14.1%) and size (15.2% versus 11.8%) measures reported in Section 3, but they are of a similar order of magnitude.

[23]Standardization at the item level (see footnote 14) guarantees that any unmodelled heterogeneous cross-sectional volatility does not bias the estimated shape.

Table 3: Estimated parameters

Parameters	Baseline linear hazard	Calvo (1983) time dependent	Golosov and Lucas (2007) (S,s)
Review cost ($\kappa, \%$)			1.3
Std. dev. of idiosyncratic shocks ($\sigma_z, \%$)	3.8	4.1	3.0
Hazard intercept (a, %)	6.5		
Hazard slope (b,$x < 0$, %)	81.8		
Hazard slope (c,$x \geq 0, \%$)	48.7		

Note: The table shows the estimated parameters of the baseline linear hazard model, the time-dependent Calvo (1983) model and the strongly state-dependent (S,s) Golosov and Lucas (2007) model.

not only qualitatively, but also quantitatively reasonably well (kurtosis in the data is 2.7 and in the model it is 3.5). In the online Appendix B.1, we show that the model also matches the standardized price-gap distribution well.

Table 4: Relative strength of adjustment channels

	Empirical moments	Baseline linear hazard	Calvo (1983) time dependent	Golosov and Lucas (2007) (S,s)
Intensive margin (%)	73.4	71.3	100.0	22.7
Gross extensive margin (%)	26.5	28.7	0.0	38.7
Narrow selection (%)	0.2	0.0	0.0	38.6

Note: The table shows the relative strength of the adjustment channels using the empirical moments (hazard and price-gap distributions) and in the baseline linear hazard model, the time-dependent Calvo (1983) model and the strongly state-dependent Golosov and Lucas (2007) model. The results show that the baseline linear hazard model implies a decomposition similar to the one using the empirical moments. It assigns no role to narrow selection and finds that the gross extensive margin raises the impact of the intensive-margin effect by 36%. The table also shows that the time-dependent Calvo (1983) model would assign no role to the gross extensive margin, while the strongly state-dependent Golosov and Lucas (2007) model assigns a strong and approximately equal role to both the gross-extensive-margin and the narrow selection effects.

Given this model setup, we compute the decomposition of the various adjustment channels under our baseline linear hazard model and contrast this decomposition with the decompositions under various counter-factual alternatives. The second column of Table 4 shows the decomposition that uses the theoretical hazard function and price-gap distribution from the main linear hazard model. We find that the decomposition is qualitatively similar to the one that relies on the empirical moments. It assigns no role to narrow selection and a significant role to the gross extensive margin. The quantitative magnitude of the estimates is slightly different: according to the model the gross extensive margin raises the impact of the intensive-margin effect by around 20% rather than 26% based on the empirical moments.

Crucially, we also use the flexible theoretical model to assess the strength of the various adjustment margins in counterfactual scenarios with a time-dependent constant hazard function (Calvo, 1983) or a strongly state-dependent hazard function (Golosov and Lucas, 2007). These exercises impose discipline on the choice of price-setting models. To do so, we recalibrate the parameters of the respective models to match the frequency and the size of price changes. Because the model solutions provide the endogenous price-gap distribution, we can then use that distribution along with the hazard function to decompose the impact of the various adjustment margins based on the equations in Section 2.2.

What do we learn? The third column of the table shows that the time-dependent model would imply no adjustment on either the gross-extensive or the narrow selection margins. By contrast, the last column shows that in the strongly state-dependent Golosov and Lucas (2007) model, the total extensive margin is powerful and raises the impact effect relative to a time-dependent benchmark (the intensive-margin effect) almost five-fold. Furthermore, both narrow selection and the gross extensive margin contribute to this state-dependent impact almost equally. This finding justifies using narrow selection to distinguish between different state-dependent models. In the online Appendix B.2, we show the decomposition for a wider range of popular state-dependent models (Woodford, 2009; Gertler and Leahy, 2008). In all these models, narrow selection plays a strictly positive role in the adjustment differently from our baseline linear hazard model.

5 Dynamic shocks and selection

This section presents our second approach to assessing the contributions of different adjustment margins. We explicitly identify aggregate shocks and measure how their interaction with product-level price gaps affects price adjustment decisions. We find no effect of this interaction, while both the shock and the price gap are positively associated with price-adjustment decisions, implying that the intensive and gross extensive margins are active but narrow selection is not.

This section first discusses the properties of our baseline identified shock: the aggregate credit shock. The online Appendix C.2 contains a discussion of the response to a monetary policy shock. Second, we present our panel-regression specification, our estimation results and a battery of robustness tests.

5.1 Dynamic impact of credit shocks

Our analysis gauges financial conditions using the excess bond premium measure, which is a time series of corporate bond spreads purged from the impact of firms' idiosyncratic default probabilities (Gilchrist and Zakrajšek, 2012). We identify credit shocks with standard exclusion restrictions. We first use a local projection approach to establish the dynamic properties of the credit shock in the economy (Jordà, 2005). We implement the identifying restrictions by including the contemporaneous values of the excluded variables as controls (Plagborg-Moller and Wolf, 2021).

In particular, our estimation includes a series ($h = 0, \ldots, 24$ months) of regressions of the form:

$$x_{t+h} - x_t = \alpha_h + \beta_h \mathrm{ebp}_t + \Gamma_h \Phi(L) X_t + \Psi_h \Theta(L) \mathrm{ebp}_{t-1} + u_{t,h}, \tag{11}$$

where ebp_t denotes the excess bond premium and X_t bundles the one-year Treasury rate, the core consumer price index and the industrial production. The regressions include the contemporaneous values and 1-12 months lags of all variables.

The key object of interest is the coefficient β_h on the credit shock. We plot it for all variables of interest for $h = 0, 1, \ldots, 24$ along with 95% confidence bands using the Newey and West (1987) heteroscedasticity and autocorrelation-consistent standard errors (Stock and Watson, 2018). Figure 6 shows the impulse responses over the sample from 2001 to 2012. The figures show that even though the credit shock (a) is accompanied by a quick easing of monetary policy (b), it is associated with a sizable decline in industrial production (c) and the core CPI (d). The figures also show that the credit shock was accompanied by a sizable gradual decline in the CPI food-at-home subcomponent (e) and the decline is well captured by our supermarket reference-price index (f).[24] The decline does not reach its peak before 24 months.

Figure 6: Impulse responses of key macroeconomic variables to a credit shock, 2001-2012

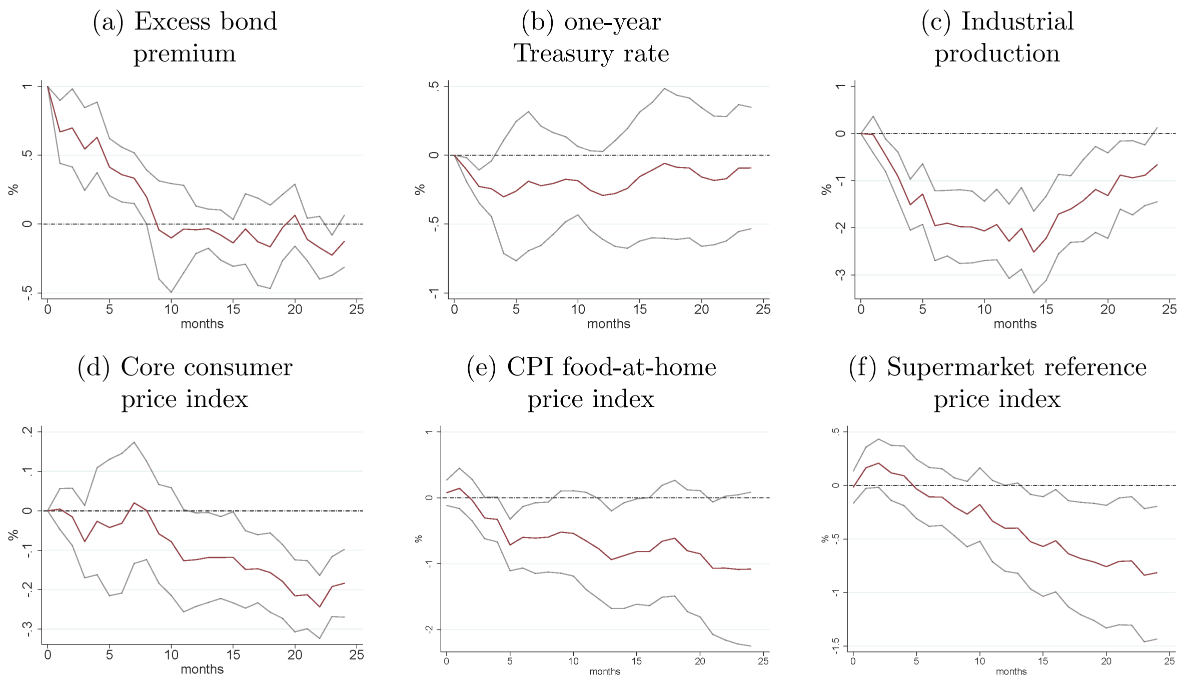

Note: The panels show impulse responses to an identified credit shock over the sample 2001-2012 in a local-projection framework, and 95% confidence bands using Newey-West standard errors. The figures show that the credit tightening causes a sizable drop in activity and various price indexes despite sizable monetary policy easing.

5.2 Credit shocks and selection

This section measures the strength of narrow selection conditional on a credit shock. Our main finding is that narrow selection appears to be absent in the data even though both the price gap and the aggregate shocks each have direct effects on price setting.

To arrive at these results, our analysis gauges the extent to which idiosyncratic price-adjustment pressures interact with the credit shock. We approach this estimation in a linear-probability panel-data setting. The linear-probability specification is motivated by the linearity of the empirical hazard function and its theoretical counterpart (see panel (e) of Figure 4 and panel (a) of Figure 5). Section 5.4.1 systematically validates the assumption using a non-parametric specification. Our main object of interest is the impact of the interaction of the aggregate shock and the price-gap measures on the probability of price adjustment. Our specification controls for the impact of the price gap in normal times (by including the lagged level of the gap directly), the direct impact of the shock (by including the aggregate shock directly), and time-

[24] We concentrate on reference prices because although sales-price inflation responds significantly to the credit shock in our baseline regression (not shown), it is only a feature of the large credit shock of the Great Recession, and it disappears if this period is excluded from the analysis (or if only monetary policy shocks are analyzed, which are smaller on average).

dependence (by including the age of the price); includes both product-store and (12) calendar-month fixed effects; and employs two-way clustering across product category and time. We use a random 10% subsample of our IRi supermarket data set.[25]

Formally, our specification is

$$I^{\pm}_{pst,t+h} = \beta^{\pm}_{xih} x_{pst-1} \widehat{\mathrm{ebp}}_t + \beta^{\pm}_{xh} x_{pst-1} + \beta^{\pm}_{ih} \mathrm{ebp}_t + \gamma^{\pm}_h T_{pst-1} + \Gamma^{\pm}_h \Phi(L) X_t + \alpha^{\pm}_{psh} + \alpha^{\pm}_{mh} + \varepsilon^{\pm}_{psth} \tag{12}$$

where $I^{\pm}_{pst,t+h}$ is an indicator of a reference price increase or decrease, respectively, of product p in store s between periods t and $t+h$. The choice of horizon ($h = 24$) is motivated by the peak price-level impact of the shock (see the previous section).[26]

The specification controls for the direct effect of the one-month lagged price gap $x_{pst-1} = p_{pst-1} - p^*_{pst-1}$. The lagged measure is predetermined and unaffected by the contemporaneous aggregate shock. Its independence from the aggregate shock simplifies the interpretation of the empirical results. The expected sign of the coefficients β_{xh} is negative for price increases (and positive for price decreases), as a larger positive gap indicates that the current price is too high (less incentive to increase prices) and a larger negative gap indicates that the current price is too low (more incentive to increase prices).

Our analysis identifies the credit shock by exclusion restrictions analogous to the local projections shown in Section 5.1. We measure financial conditions by the excess bond premium (ebp) (Gilchrist and Zakrajšek, 2012). The direct impact of the shock on the probability of price adjustment is measured by $\beta^{\pm}_{ih}$.[27] The coefficient has a negative expected sign for price increases (and positive expected sign for price decreases), as a credit tightening implies fewer increases and more decreases, and a credit easing implies more increases and fewer decreases. The credit shock measure ($\widehat{\mathrm{ebp}}$) in the interaction term purges the excess bond premium by regressing it on the current month industrial production, the core consumer price level, and the one-year Treasury yield and up to six lags of its own level and the other aggregate variables. This is equivalent to identifying the credit shock using a Cholesky factorization and ordering the credit shock last.[28]

Our focus is the coefficient of the interaction term between the aggregate shock and the price gap, $beta^{\pm}_{xih}$. It governs the extent of narrow selection. What results should we expect if we follow economic intuition? For price increases I^{+}_{pst} the expected sign is negative: after a credit tightening ($\widehat{\mathrm{ebp}}_t \leq 0$): the more positive the price gap ($x_{pst-1} \leq 0$), the larger is the decline in the probability of price increases; and vice versa after a credit easing ($\widehat{\mathrm{ebp}}_t < 0$): the more negative the price gap ($x_{pst-1} \leq 0$), the larger the expected increase in the probability of price increases. In contrast, after the price decreases I^{-}_{pst}, the expected sign is positive.[29]

[25] We limit the size of our IRi sample because of computational constraints. The estimates are invariant to further sample increases and to drawing a second 10% sample. We subsequently show the robustness of our results in the full PPI sample and non-linear probability models in Section 5.4, and, in the online appendix, to monetary policy shocks and when we use alternative price-gap proxies.

[26] On average, the probability of a reference-price increase or decrease at this horizon is 54% and 24%, respectively.

[27] The regression controls directly for the impact of the current-period activity, prices, and Treasury yields, which is equivalent to an identification with Cholesky factorization and ordering the credit shock last (Plagborg-Moller and Wolf, 2021).

[28] The credit shock is a generated regressor. However, as the interaction term is insignificant in our regressions, we do not need to adjust our standard errors (Wooldridge, 2010).

[29] The interaction term of the current aggregate shock and the *lagged* price gap ignores the impact of contemporaneous idiosyncratic shocks, which, together with the contemporaneous aggregate shock, also affect the contemporaneous optimal price (p^*_{pst}), and, thus, the size of the price adjustment. The relevance of the lagged gap in explaining a sizable variation in the contemporaneous gap and its close relationship with the size of the price change (Figure 4), however, makes its interaction term with the contemporaneous aggregate shock a valid object of interest to test the presence of narrow selection. The impact of unobserved idiosyncratic shocks on price selection is the focus of the analysis of Dedola et al. (2021), who find a statistically significant, but economically small impact using Danish producer-price data.

Our specification controls for some time-dependent features of price setting. We measure the age of the price as the number of months since the last reference-price change (T_{pst}) and include its logarithm as a control variable. Finally, $\Gamma_h^{\pm}\Phi(L)X_t$ denotes aggregate controls. We include contemporaneous values and up to six lags of industrial production as a measure of economic activity, the core consumer price index (CPI) as a measure of the price level, and the one-year Treasury yield as a monetary policy indicator.

To show the robustness of our analysis, we present two additional variants of our main specification. First, we analyze a specification with time fixed effects (which absorb the direct effects of the aggregate shock). Second, we analyze a specification with separate coefficients for positive and negative gaps.

Our main results are twofold. Our first result presents strong evidence for broad selection in the data as the probability of price adjustment increases with the price gap, and the aggregate credit shock affects the probability of price adjustment through the gross extensive margin. Table 5 illustrates these results in our baseline regressions for the competitor-price-gap proxy in response to a credit shock. Column (1) is the main specification for price increases and Column (4) for price decreases. Columns (2) and (5) show specifications with time fixed effects while (3) and (6) show results for positive and negative gaps separately.

The coefficients show that the effects are generally economically meaningful. First, the probability of a price increase for a product at the first quartile of the price-gap distribution (-0.08) is 26 percentage points lower relative to the third quartile (0.07); for the price decrease, it is 23 percentage points higher. Second, a one-standard-deviation credit tightening (30 basis points) reduces the probability of price increases by around 1 percentage point and increases the probability of a price decrease by the same amount.

Our second main finding is that there is no evidence of narrow selection: conditional on the aggregate shock, the chance of new price adjustments is independent of our measure of the price gap. None of the coefficients on the interaction term are statistically significantly different from zero.

5.3 Implications

These regressions are complementary to the analysis of Section 4, which uses unconditional moments. They explicitly assess the reactions conditional on identified aggregate shocks. They confirm that the (aggregate credit) shock indeed activates the gross extensive margin in line with predictions of the linear hazard model presented in Section 4.4. In line with the same model, they suggest that narrow selection is absent.

A relevant feature of the regressions is that they measure cumulative price-change probabilities over a 24-month period until the peak impact of the shock. This horizon gives time for the credit shock to gradually pass through to prices. Over such an extended period, however, unmeasured idiosyncratic shocks and time aggregation can also potentially influence the results. In this section, we therefore also assess the impact of this feature relying on the data and simulations of the models presented in Section 4.4.

Panel (a) of Figure 7 shows the empirical cumulative probability of a price change over a period of $h = 1, 3, 12, 24$ months as a function of the previous month's competitors' price gap. The panel shows that the cumulative generalized hazard functions stay V-shaped over time. This is in line with broad selection also at longer horizons: the probability of a price change increases with the size of the initial price gap. Notably, the size of the cumulative price changes conditional on adjustment stays very close to one-to-one over longer horizons (not shown). Such evidence for broad selection is inconsistent with time-dependent models like Calvo (1983), which ignores product-dependent pressures as a factor in price-adjustment probability.

Panels (b-d) show equivalent cumulative generalized hazard functions for different horizons simulated in

Table 5: Estimates, scanner data, competitor-price gap, credit shock

	(1)	(2)	(3)	(4)	(5)	(6)
	Price increase $\left(I^{+}_{pst,t+24}\right)$			Price decrease $\left(I^{-}_{pst,t+24}\right)$		
Gap (x_{pst-1})	-1.75^{***}	-1.75^{***}		1.55^{***}	1.55^{***}	
	(0.06)	(0.06)		(0.06)	(0.06)	
Shock (ebp_t)	-0.03^{***}		-0.04^{***}	0.03^{***}		0.03^{***}
	(0.01)		(0.01)	(0.01)		(0.01)
Selection ($x_{pst-1}\hat{\text{ebp}}_t$)	-0.00	-0.00		0.01	0.01	
	(0.04)	(0.04)		(0.05)	(0.04)	
Age (T_{pst-1})	0.02^{***}	0.02^{***}	0.02^{***}	0.00^{**}	0.01^{***}	0.01^{***}
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Pos. gap (x^{+}_{pst-1})			-2.26^{***}			2.29^{***}
			(0.13)			(0.10)
Neg. gap (x^{-}_{pst-1})			-1.44^{***}			1.10^{***}
			(0.07)			(0.06)
Pos. sel. ($x^{+}_{pst-1}\hat{\text{ebp}}$)			0.04			-0.04
			(0.06)			(0.05)
Neg. sel. ($x^{-}_{pst-1}\hat{\text{ebp}}$)			-0.03			0.04
			(0.06)			(0.07)
Product x store FE	✓	✓	✓	✓	✓	✓
Calendar-month FE	✓	✗	✓	✓	✗	✓
Time FE	✗	✓	✗	✗	✓	✗
N	$16.1M$	$16.1M$	$16.1M$	$16.1M$	$16.1M$	$16.1M$
within R^2	18.5%	16.6%	18.9%	17.3%	16.4%	18.2%

Note: The table shows estimation results from a linear-probability panel model using scanner data. The regressions are run separately using an indicator with value 1 for reference-price increases (Columns 1-3) and an indicator with value 1 for reference-price decreases (Columns 4-6). The regressions include product-store and calendar-month fixed effects, control for the age of the price (time since last change), and use standard errors with two-way clustering. The baseline regressions (Columns 1 and 4) show that even though the price gap and the aggregate shock significantly influence the price-change probabilities, which is evidence for state-dependence and an active gross extensive margin, their interaction remains insignificantly different from zero, suggesting narrow selection is absent. The results stay robust to a specification with time fixed effects (Columns 2 and 5) and a specification with separate coefficients for positive and negative gaps (Columns 3 and 6).
*Standard errors in parentheses; *: significant at 10%, **: significant at 5%, ***: significant at 1%.*

our calibrated models. Panel (b) shows that the linear hazard model is broadly consistent with the data. It somewhat overestimates the probability of price adjustment at 12-to-24-month horizons, and implies hazard function shapes that are more non-linear than those found in the data. The fixed menu cost (S,s) model similarly overestimates the probability of adjustment at long horizons, but it also implies hazards that are too steep at short horizons. Notably, the shapes of the cumulative hazards of the linear and the (S,s) models become similar over time.

The similarity shows up in Table 6, which uses the cumulative generalized hazard functions to decompose

Figure 7: Cumulative hazard functions in the data, and three models

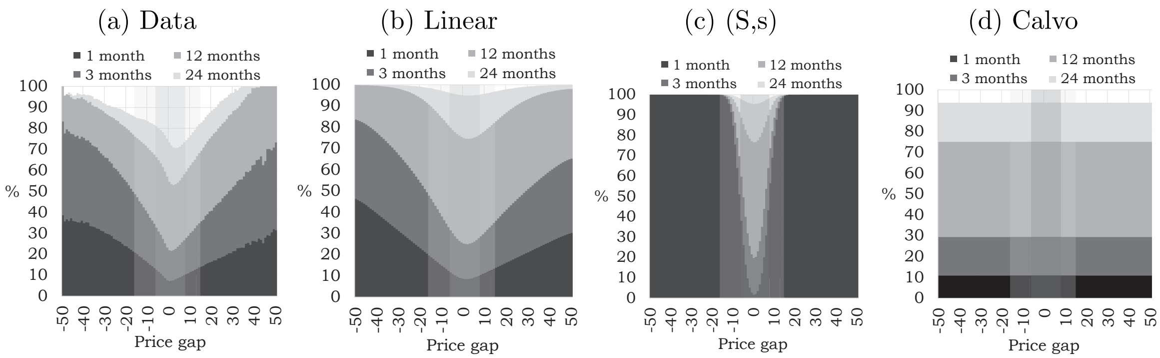

Note: The panels show the cumulative hazard functions ($h = 1, 3, 12, 24$ months) in the data (a), in the linear hazard model (b), in the Golosov and Lucas (2007) (S,s) model (c) and in the Calvo (1983) time-dependent model (d). The first panel shows that the hazard functions retain their V shapes over longer horizons, which is inconsistent with time-dependent models (d). The linear hazard model (b) stays broadly consistent with the data. The difference in the shapes between the models becomes less pronounced with the horizon.

the cumulative inflation impact into adjustment channels. Relying on the empirically consistent linear hazard in the model, the second and third columns of the table confirm that the gross extensive margin stays active at both 12-month and 24-month horizons, and the narrow selection effect remains small. These effects are captured well by the linear hazard model, but are generally missed by the Calvo (1983) model, as the last 2 columns show. Importantly, however, the fixed menu cost (S,s) model also captures the long-horizon decomposition in the data quite well as columns 5 and 6 show. We conclude that the regressions presented in this section provide further evidence against time-dependent models and support the linear hazard framework presented in Section 4. However, the regressions relying on 24-month cumulative price changes cannot successfully distinguish between various state-dependent models. For that, one would need aggregate shocks with faster pass through than either credit or monetary policy shocks (e.g., value-added tax shocks). We leave this for future research.

Table 6: Relative strength of adjustment channels using cumulative hazards

	Empirical		Linear		(S,s)		Calvo	
	12m	24m	12m	24m	12m	24m	12m	24m
Intensive margin	89.7	94.5	89.7	98.0	93.6	95.9	100	100
Gross extensive margin	8.8	3.8	8.22	1.6	5.1	1.0	0.0	0.0
Selection	1.5	1.7	2.1	0.4	1.4	0.25	0.0	0.0

Note: The table shows the decomposition of the adjustment channels using cumulative hazards over 12 and 24 months based on the empirical hazard functions (first and second columns) and the linear hazard, the Golosov and Lucas (2007) (S,s) and the time-dependent Calvo (1983) models. The results show that the linear hazard model generates decompositions that are in line with the data, but long horizons diminish the difference between the various price-setting models.

5.4 Robustness

This section shows that our regression results remain robust when we use a non-linear specification and when we use producer-price (PPI) microdata instead of retail data. The online Appendix C.2 presents a battery of further robustness checks. In particular, we show that our results are robust when we consider an alternative gap measure, the reset-price gap, when we use monetary policy shocks instead of the credit shocks, and when we conduct our estimation at the category level. Narrow selection remains absent in all cases.

5.4.1 Non-linearity

A potential concern with the analysis may be found in the linear relationship imposed by our baseline specification between the price-change probabilities and the price gap, as well as the price gap conditional on the size of an aggregate shock. If these relationships are non-linear, our rejection of the presence of narrow selection could be a rejection of the linearity assumption. This section conducts a robustness check using a non-parametric approach and rules out such concerns. Online Appendix C.3 shows the robustness of our results in two additional non-linear specifications: a multinomial- and an ordered-probit model.[30]

Our main non-parametric approach begins by assigning price gaps into 15 approximately equal-sized bins.[31] The bin that serves as a reference group includes items with small price gaps, in particular price gaps between $-1\% \leq x_{pst-1} < 1\%$.

As before, our analysis runs separate regressions for positive and negative price changes with both product-store and time fixed effects and two-way clustering. Instead of the size of the price gap, we now include bin-dummies both as direct regressors and in the interaction terms. Our specification excludes the reference group from the direct regressors and the interaction terms, so the estimated coefficients are all relative to the reference group.

Figure 8 depicts the estimated coefficients. The red lines show the impact on price-increase probabilities and price-decrease probabilities of the price-gap groups relative to the group with small gaps. The panels show that the size of the gap has a large impact on the probability of price adjustment in the next 2 years: For example, if the negative gap exceeds 20%, the probability of a future price decrease goes up by almost 50 percentage points. The figure also shows that the relationship between the gaps and the adjustment probabilities is monotone and close to linear, justifying our linear baseline specification. The blue lines show the impact of the gap-bin-credit-shock interaction terms on the price-change probabilities. The panels show that the additional impact brought about by the aggregate shock does not significantly vary with the gap, so narrow selection is undetectable. The panels also show that the insignificant results in our baseline specifications are not the consequence of imposing linearity because there is no detectable interaction between the aggregate shock and the gaps at any gap levels.

[30] These specifications explicitly take into account some fundamental aspects of probabilities, most notably that (a) the sum of the probabilities of price increases, price decreases and no price changes equals one and (b) that probabilities are non-negative.

[31] The bins are not equal-sized to maintain symmetry. In particular, we generate 5 equal-sized bins with negative gaps and 5 equal-sized bins with positive gaps. Then we merge the largest negative and the smallest positive bins and each consecutive negative bin and positive bin. Because the price-gap distribution is approximately symmetric, we obtain 15 approximately equal-sized bins.

Figure 8: Non-linear specification, 15 groups

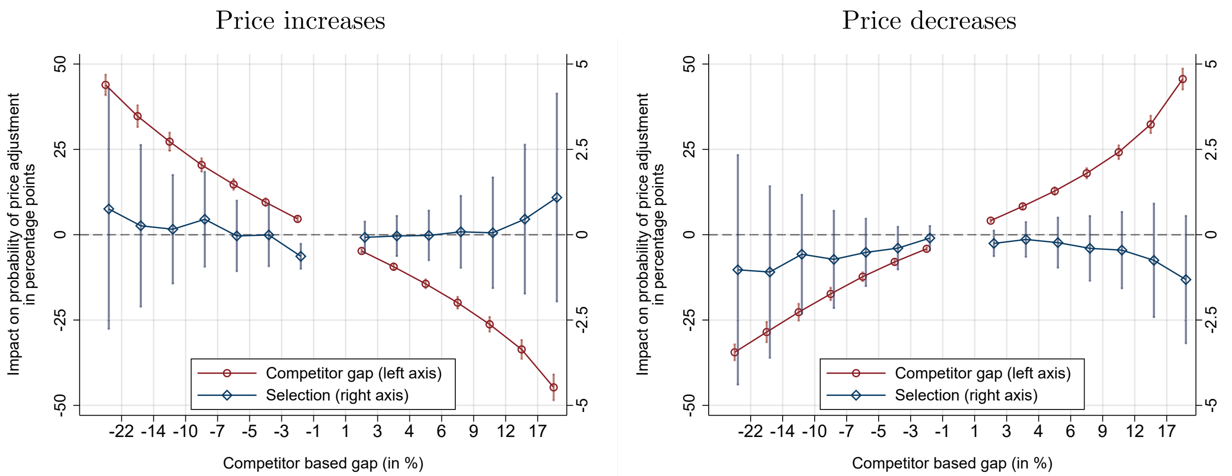

Note: The figures depict the impact on price-increase (left panel) and price-decrease (right panel) probabilities. The red lines show the estimated differential impact of 15 equal-sized groups with various price gaps and the blue lines the corresponding price-gap-credit-shock interaction terms, always relative to the group with gaps close to zero. The specification includes product-store and time fixed effects. Standard errors are clustered along categories and time. The vertical lines show 95% confidence bands. Both figures show that while the gap itself significantly influences the price-adjustment probability, it has no significant interaction with the aggregate shock at any gap size, indicating that narrow selection is absent.

5.4.2 Producer-price microdata

Our results also hold when we use producer-price microdata rather than the (retail) IRi scanner data. We obtain producer-price data from the microdata underlying the U.S. producer-price index. It is less granular than the IRi microdata. We now define the price gap at the six-digit NAICS level, which is still fairly detailed (such as Tortilla Manufacturing (NAICS code 311830)). While this relatively lower granularity reduces the quality of the relevant price-setting proxies, the data also span a much longer time period (1981-2012) and a wider set of sectors than the IRi microdata.

Consistent with the baseline results we find that the impulse response of the PPI to the credit shock is economically intuitive. Despite a monetary easing, industrial production and prices fall following a credit shock. We estimate specification (11) with the PPI on the left-hand side to show these results (see Figure 9).

Moreover, our price-gap proxy is informative in a fashion similar to that of the IRi retail data, establishing the general validity of our proxy measures. Figure 10 shows the density of competitor-price gaps, together with the relationship of the gaps and the frequency and size of subsequent price changes. The figure shows that there is a clear negative relationship between the size of the price changes and the proxies for the gaps. The proxies clearly capture a relevant part of the theoretical price gaps, even though the proxies are not as well measured as in the more granular scanner data. The frequency of price changes increases with the price gap as the gap becomes sufficiently large, even though it declines for small price gaps. This might reflect the role of heterogeneity in price-setting frequencies across sectors, as well as the presence of measurement error.

Figure 9: Impulse responses of key macroeconomic variables to a credit shock, 1985-2015

(a) Excess bond premium (b) One-year Treasury (c) PPI (d) Industrial Production

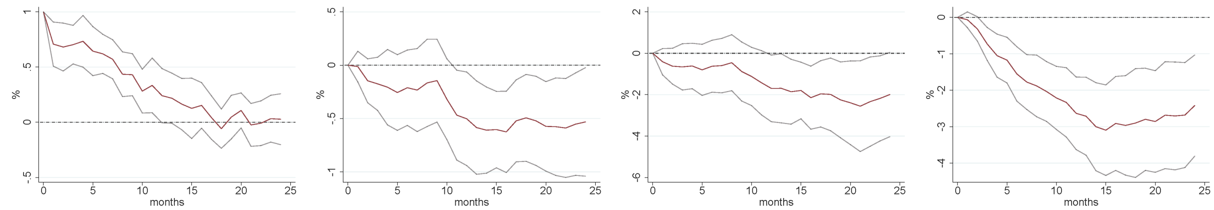

Note: The figures show impulse responses to an identified credit shock over the PPI sample 1985-2015 in a local-projection framework, and 95% confidence bands using Newey-West standard errors. The figures show that the credit tightening causes a sizable drop in activity and the price index despite sizable monetary policy easing.

Figure 10: Competitor-price-gap density and the subsequent frequency and size of price changes as a function of the gap, PPI data

(a) Size (b) Frequency (c) Density

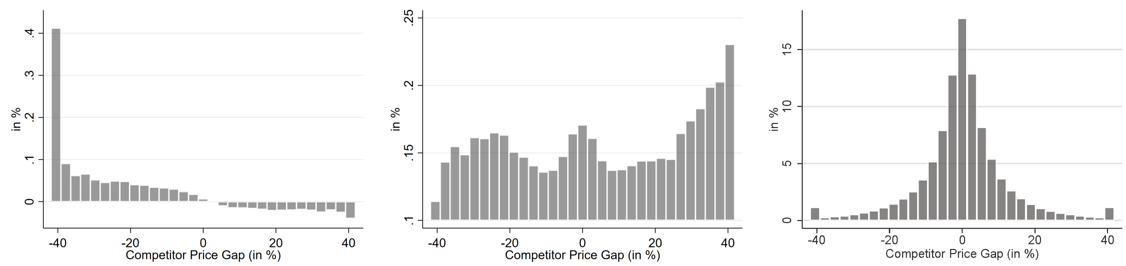

Note: The figures show the size and frequency responses of subsequent price changes as a function of the price gap and its density in the PPI data set. The figures show that the size of average subsequent adjustments is negatively related to the gap (panel a); the frequency of subsequent price adjustments increases with the absolute size of the gap, after the gap is large enough (panel b); and the density of the competitor-price gaps has fat tails (panel c).

Estimation results based on the PPI microdata are entirely consistent with the results of the baseline analysis. Table 7 shows the linear-probability panel estimations for our baseline and time fixed effect specifications using the PPI microdata for the competitor-price gap. There is evidence for state dependence because larger gaps and the aggregate shocks change the probabilities of price adjustment. At the same time, there is no evidence of selection: conditional on the aggregate shock, the new adjusters do not come from those prices that are further from their optimal levels. None of the interaction terms are statistically significant.

The impact of the price gaps and the aggregate credit shocks also remains economically significant. The probability of price increases between a product with a competitor-price gap at the third quartile and at the first quartile gets smaller by 23 percentage points, and the probability of price increases gets larger by 22 percentage points. Finally, a one standard deviation credit tightening reduces the price increase probability by around 0.7 percentage point and increases the price decrease probability by a similar amount.

Overall, these results indicate that the absence of a narrow selection effect generalizes from the supermarket sector to the wider U.S. economy.

Table 7: Estimates, PPI data, competitor-price gap, credit shock

	(1)	(2)	(3)	(4)
	Increases $\left(I^{+}_{pst,t+24}\right)$		Decreases $\left(I^{-}_{pst,t+24}\right)$	
Gap (x_{pst-1})	-0.23^{***}	-0.23^{***}	0.22^{***}	0.22^{***}
	(0.02)	(0.02)	(0.02)	(0.02)
Shock (ebp_t)	-0.023^{***}		0.021^{***}	
	(0.01)		(0.01)	
Selection ($x_{pst-1}\hat{\text{ebp}}_t$)	0.00	-0.00	-0.00	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)
Age (T_{pst-1})	0.035^{***}	0.035^{***}	0.01^{***}	0.01^{***}
	(0.00)	(0.00)	(0.00)	(0.00)
Product x store FE	✓	✓	✓	✓
Calendar-month FE	✓	✗	✓	✗
Time FE	✗	✓	✗	✓
N	$9.7M$	$9.7M$	$9.7M$	$9.7M$
Within R^2	4.4%	3.5%	4.3%	3.7%

Note: The table shows estimation results from the linear-probability panel model using PPI microdata. The regressions are run separately on an indicator with value 1 for price increases (Columns 1-2) and an indicator with value 1 for pride decreases (Columns 3-4). The regressions include product-store fixed effects, control for the age (time since last change) of the price, and use standard errors with two-way clustering. The baseline regressions (Columns 1 and 3) show that even though the price gap and the aggregate shock significantly influence the price-change probabilities, their interaction remains insignificantly different from zero, suggesting selection is absent. The results stay robust to a specification with time fixed effects (Columns 2 and 4).
*Standard errors in parentheses; *: significant at 10%, **: significant at 5%, ***: significant at 1%.*

6 Conclusion

This paper defines a measure of price selection and quantifies it in supermarket-scanner and producer-price microdata. The analysis is based on a decomposition of inflation responses that uses the hazard function and the density of price gaps. When we empirically estimate these objects, we find a linear hazard function that implies in our decomposition that 73.4% of the inflationary effect of an aggregate shock is driven by changes to the size of price adjustments (the intensive margin), 26.5% by shifts between price decreases and price increases (the gross extensive margin) and only 0.2% by endogenous selection of which prices adjust after the aggregate shock (the narrow selection effect). A structural model analysis that estimates a flexible price-setting model to match the empirically estimated hazard function and then uses the endogenously arising price-gap distribution to compute the decomposition arrives at the same result: narrow selection is absent. Regression analysis that systematically relates price responses to the interaction of shocks and price gaps aligns with our findings of no narrow selection.

To provide guidance for the choice of price-setting technologies, we run a counterfactual model exercise where we calibrate the model under Golosov and Lucas (2007) and Calvo (1983) price-adjustment technologies. A clear finding emerges: the time-dependent model implies no adjustment on either the gross extensive or the narrow selection margins. In contrast, the total extensive margin is powerful in the state-dependent Golosov and Lucas (2007) model and raises the impact effect of the extensive margin to 77.3%. Furthermore, both narrow selection and the gross extensive margin contribute to the impact effect nearly equally. Thus, the comparison to the model calibrated to the empirically relevant linear hazard allows us not only to discipline the choice between time-dependent and state-dependent models, but also among state-dependent models. Our evidence is generally consistent with state-dependent models with linear and flat adjustment hazards, which predict high monetary non-neutrality that is only around 36% milder than that in time-dependent models.

References

Alvarez, Fernando, Hervé Le Bihan, and Francesco Lippi (2016) "The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach," *American Economic Review*, Vol. 106, pp. 2817–51.

Alvarez, Fernando, Francesco Lippi, and Aleksei Oskolkov (2021) "The Macroeconomics of Sticky Prices with Generalized Hazard Functions," *The Quarterly Journal of Economics*, Vol. 137, pp. 989–1038.

Argente, David and Chen Yeh (2022) "Product Life Cycle, Learning, and Nominal Shocks," *The Review of Economic Studies*, Vol. 89, pp. 2992–3054.

Atkeson, Andrew and Ariel Burstein (2008) "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, Vol. 98, pp. 1998–2031.

Auclert, Adrien, Rodolfo D. Rigato, Matthew Rognlie, and Ludwig Straub (2022) "New Pricing Models, Same Old Phillips Curves?" NBER Working Papers 30264, National Bureau of Economic Research, Inc.

Auer, Raphael A. and Raphael S. Schoenle (2016) "Market Structure and Exchange Rate Pass-Through," *Journal of International Economics*, Vol. 98, pp. 60–77.

Baley, Isaac and Andrés Blanco (2021) "Aggregate Dynamics in Lumpy Economies," *Econometrica*, Vol. 89, pp. 1235–1264.

Ball, Laurence and David Romer (1990) "Real Rigidities and the Non-neutrality of Money," *The Review of Economic Studies*, Vol. 57, pp. 183–203.

Balleer, Almut and Peter Zorn (2019) "Monetary Policy, Price Setting, and Credit Constraints," CEPR Discussion Papers 14163.

Beck, Gunter W. and Sarah M. Lein (2020) "Price Elasticities and Demand-side Real Rigidities in Micro Data and in Macro Models," *Journal of Monetary Economics*, Vol. 115, pp. 200–212.

Berger, David and Joseph Vavra (2018) "Dynamics of the US Price Distribution," *European Economic Review*, Vol. 103, pp. 60–82.

Bhattarai, Saroj and Raphael Schoenle (2014) "Multiproduct Firms and Price-Setting: Theory and Evidence from US Producer Prices," *Journal of Monetary Economics*, Vol. 66, pp. 178–192.

Bils, Mark, Peter J Klenow, and Benjamin A Malin (2012) "Reset Price Inflation and the Impact of Monetary Policy Shocks," *The American Economic Review*, Vol. 102, pp. 2798–2825.

BLS (2020) "Producer Price Index: Frequently Asked Questions," https://www.bls.gov/ppi/ppifaq.htm.

Bonomo, Marco, Carlos Carvalho, Oleksiy Kryvtsov, Sigal Ribon, and Rodolfo Rigato (2023) "Multi-Product Pricing: Theory and Evidence from Large Retailers," *The Economic Journal*, Vol. 133, pp. 905–927.

Bronnenberg, Bart J, Michael W Kruger, and Carl F Mela (2008) "Database Paper – The IRi Marketing Data Set," *Marketing Science*, Vol. 27, pp. 745–748.

Caballero, Ricardo J and Eduardo MRA Engel (2007) "Price Stickiness in Ss models: New Interpretations of Old Results," *Journal of Monetary Economics*, Vol. 54, pp. 100–121.

Calvo, Guillermo A. (1983) "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, Vol. 12, pp. 383 – 398.

Campbell, Jeffrey R. and Benjamin Eden (2014) "Rigid Prices: Evidence from U.S. Scanner Data," *International Economic Review*, Vol. 55, pp. 423–442.

Caplin, Andrew S. and Daniel F. Spulber (1987) "Menu Costs and the Neutrality of Money," *The Quarterly Journal of Economics*, Vol. 102, pp. 703–726.

Carlsson, Mikael (2017) "Microdata Evidence on the Empirical Importance of Selection Effects in Menu-Cost Models," *Journal of Money, Credit and Banking*, Vol. 49, pp. 1803–1830.

Carvalho, Carlos and Oleksiy Kryvtsov (2021) "Price Selection," *Journal of Monetary Economics*, Vol. 122, pp. 56–75.

Costain, James and Anton Nakov (2011a) "Distributional Dynamics under Smoothly State-Dependent Pricing," *Journal of Monetary Economics*, Vol. 58, pp. 646 – 665.

——— (2011b) "Price Adjustments in a General Model of State-Dependent Pricing," *Journal of Money, Credit and Banking*, Vol. 43, pp. 385–406.

Dedola, Luca, M Strom Krisoffersen, and Gabriel Züllig (2021) "The Extensive and Intensive Margin of Price Adjustment to Cost Shocks: Evidence from Danish Multiproduct Firms," unpublished manuscript.

Dotsey, Michael, Robert G. King, and Alexander L. Wolman (1999) "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *The Quarterly Journal of Economics*, Vol. 114, pp. 655–690.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo (2011) "Reference Prices, Costs, and Nominal Rigidities," *American Economic Review*, Vol. 101, pp. 234–62.

Ellickson, Paul B (2013) "Supermarkets as a Natural Oligopoly," *Economic Inquiry*, Vol. 51, pp. 1142–1154.

Gagnon, Etienne, David López-Salido, and Nicolas Vincent (2012) "Individual Price Adjustment along the Extensive Margin," *NBER Macroeconomics Annual 2012, Volume 27*, pp. 235–281.

Gertler, Mark and Peter Karadi (2015) "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, Vol. 7, pp. 44–76.

Gertler, Mark and John Leahy (2008) "A Phillips Curve with an Ss Foundation," *Journal of Political Economy*, Vol. 116, pp. 533–572.

Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajšek (2017) "Inflation Dynamics during the Financial Crisis," *American Economic Review*, Vol. 107, pp. 785–823.

Gilchrist, Simon and Egon Zakrajšek (2012) "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, Vol. 102, pp. 1692–1720.

Goldberg, Pinelopi K. and Rebecca Hellerstein (2009) "How Rigid are Producer Prices?" Staff Reports 407, Federal Reserve Bank of New York.

Golosov, Mikhail and Robert E. Lucas (2007) "Menu Costs and Phillips Curves," *Journal of Political Economy*, Vol. 115, pp. 171–199.

Gopinath, Gita and Oleg Itskhoki (2010) "Frequency of Price Adjustment and Pass-Through," *The Quarterly Journal of Economics*, Vol. 125, pp. 675–727.

Gorodnichenko, Yuriy and Michael Weber (2016) "Are Sticky Prices Costly? Evidence from the Stock Market," *American Economic Review*, Vol. 106, pp. 165–199.

Hong, Gee Hee, Matthew Klepacz, Ernesto Pasten, and Raphael Schoenle (forthcoming) "The Real Effects of Monetary Shocks: Evidence from Micro Pricing Moments," *Journal of Monetary Economics*.

Jarociński, Marek and Peter Karadi (2020) "Deconstructing Monetary Policy Surprises: The Role of Information Shocks," *American Economic Journal: Macroeconomics*, Vol. 12, pp. 1–43.

Jordà, Òscar (2005) "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, Vol. 95, pp. 161–182.

Karadi, Peter and Adam Reiff (2019) "Menu Costs, Aggregate Fluctuations, and Large Shocks," *American Economic Journal: Macroeconomics*, Vol. 11, pp. 111–46.

Kehoe, Patrick and Virgiliu Midrigan (2015) "Prices are Sticky After All," *Journal of Monetary Economics*, Vol. 75, pp. 35–53.

Luo, Shaowen and Daniel Villar (2021) "The Price Adjustment Hazard Function: Evidence from High Inflation Periods," *Journal of Economic Dynamics and Control*, Vol. 130.

Midrigan, Virgiliu (2011) "Menu Costs, Multiproduct Firms, and Aggregate Fluctuations," *Econometrica*, Vol. 79, pp. 1139–1180.

Mongey, Simon (2021) "Market Structure and Monetary Non-neutrality," NBER Working Papers 29233, National Bureau of Economic Research, Inc.

Nakamura, Emi and Jón Steinsson (2008) "Five Facts about Prices: A Reevaluation of Menu Cost Models," *The Quarterly Journal of Economics*, Vol. 123, pp. 1415–1464.

——— (2010) "Monetary Non-neutrality in a Multisector Menu Cost Model," *The Quarterly Journal of Economics*, Vol. 125, pp. 961–1013.

Newey, Whitney and Kenneth West (1987) "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, Vol. 55, pp. 703–08.

Pennings, Steven (2017) "Pass-through of Competitors' Exchange Rates to US Import and Producer Prices," *Journal of International Economics*, Vol. 105, pp. 41–56.

Petrella, Ivan, Emiliano Santoro, and Lasse P. Simonsen (2019) "Time-varying Price Flexibility and Inflation Dynamics," EMF Research Papers 28, Economic Modelling and Forecasting Group.

Plagborg-Moller, Mikkel and Christian K. Wolf (2021) "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, Vol. 89, pp. 955–980.

Stock, James H and Mark W Watson (2018) "Identification and Estimation of Dynamic Causal Effects in Macroeconomics using External Instruments," *The Economic Journal*, Vol. 128, pp. 917–948.

Woodford, Michael (2003) *Interest and Prices: Foundations of a Theory of Monetary Policy*: Princeton University Press.

——— (2009) "Information-Constrained State-Dependent Pricing," *Journal of Monetary Economics*, Vol. 56, pp. S100–S124.

Wooldridge, Jeffrey M (2010) *Econometric Analysis of Cross Section and Panel Data*: MIT press.