



EIEF Working Paper 23/07

September 2023

Large shocks travel fast

By

Alberto Cavallo

(Harvard Business School)

Francesco Lippi

(Luiss University and EIEF)

Ken Miyahara

(University of Chicago)

Large shocks travel fast*

Alberto Cavallo

Harvard Business School

Francesco Lippi

Luiss University and EIEF

Ken Miyahara

University of Chicago

August 28, 2023

Abstract

We leverage the inflation upswing of 2022 and various granular data sets to identify robust price setting patterns following a large supply shock. We show that the frequency of price changes increases dramatically after a large shock. We setup a parsimonious New Keynesian model and calibrate it to fit the steady state data before the shock. The model features a significant component of state-dependent decisions, implying that large cost shocks incite firms to react more swiftly than usual, resulting in a rapid pass-through to prices—large shocks travel fast. Understanding this feature is crucial for interpreting recent inflation dynamics.

JEL Classification Numbers: E5

Key Words: sticky prices, misallocation, price dispersion, cost of inflation.

*We thank, without implicating, several colleagues and coauthors who generously provided feedback on a preliminary draft of the paper. In particular we wish to thank Fernando Alvarez, Isaac Baley, Andres Blanco, Xavier Gabaix, Erwan Gautier, Hugo Hopenhayn, Anil Kashyap, Herve Le Bihan, Claudio Michelacci, Virgiliu Midrigan, Anton Nakov, Luigi Paciello, and Facundo Piguillem. Tomas Pacheco provided excellent research assistance. Lippi acknowledges financial support from the ERC grant 101054421-DCS.

1 Introduction

The frequency of price adjustment is a key determinant of the speed with which shocks transmit through the economy in New Keynesian models. For example, [Gopinath and Itskhoki \(2010\)](#) document that a higher frequency of price changes is associated with a higher pass-through of exchange rate shocks, attributing the cross-sectional differences to primitives affecting the curvature of the profit function (see also [Devereux and Yetman \(2010\)](#)). Many empirical studies use micro data to document a large cross-sectional heterogeneity in the frequency of price adjustments, but there have been few attempts to study what causes changes in this frequency over time.¹ In this paper we focus on a key determinant of the frequency of price changes, and therefore the speed of shock propagation: the *size* of the shock. We document that large shocks result in an increased frequency of price adjustment, leading to higher rates of cost pass-through and boosting the propagation speed of the shock. In short we demonstrate that large shocks travel fast.

The cause of the result is the state-dependence of the firm’s price setting decisions: the probability that a firm resets its price depends on the benefits of the adjustment. A large cost shock shrinks the profit margins, inciting firms to adjust prices. Such a swift “extensive margin” response is the hallmark of state-dependent models where the probability of price adjustment depends on the firm’s state. We argue that time-dependent models, which are widely used by central bankers today, are ultimately unfit to study inflation dynamics after large shocks like the ones experienced in 2022.²

We base the analysis on a rich granular dataset on price setting behavior by retailers from several countries, before and after the large energy shocks that occurred in 2022. We interpret the data using a generalized Ss model, following the seminal work of [Caballero and Engel \(1993a,b\)](#). This model nests a broad class of State-Dependent (SD) models, such as [Golosov](#)

¹[Cavallo \(2019\)](#) found that online competition led to an increase in the frequency of price changes for US multi-channel retailers between 2008 and 2017. See [Gautier et al. \(2022\)](#) for recent extensive evidence on the frequency of price changes across the industries of European countries.

²The extensive margin is also emphasized by [Gagnon et al. \(2012\)](#), who argue that models that abstract from variation in the timing of price changes miss an important channel of macroeconomic adjustment.

and Lucas (2007) or Nakamura and Steinsson (2010), as well as several Time-Dependent (TD) models, such as the well-known Calvo (1983) model. We parametrize the model by mapping it to three cross-sectional moments of price setting behavior, measured before 2022, a period over which aggregate shocks were arguably small. These moments, shown to characterize the economy’s response to a small aggregate shock by Alvarez, Lippi, and Oskolkov (2022), are: (i) the frequency of price changes, (ii) the standard deviation of the size of price changes and (iii) the kurtosis of the size of price changes. The quantitative exercise clearly reveals that the price setting decisions feature state-dependence: the probability that a price is adjusted increases as the distance between the actual and ideal price grows. This result confirms the seminal findings of Eichenbaum, Jaimovich, and Rebelo (2011) and the recent ones by Gautier et al. (2022); Karadi et al. (2023); Dedola et al. (2023) on the empirical relevance of state-dependent pricing.

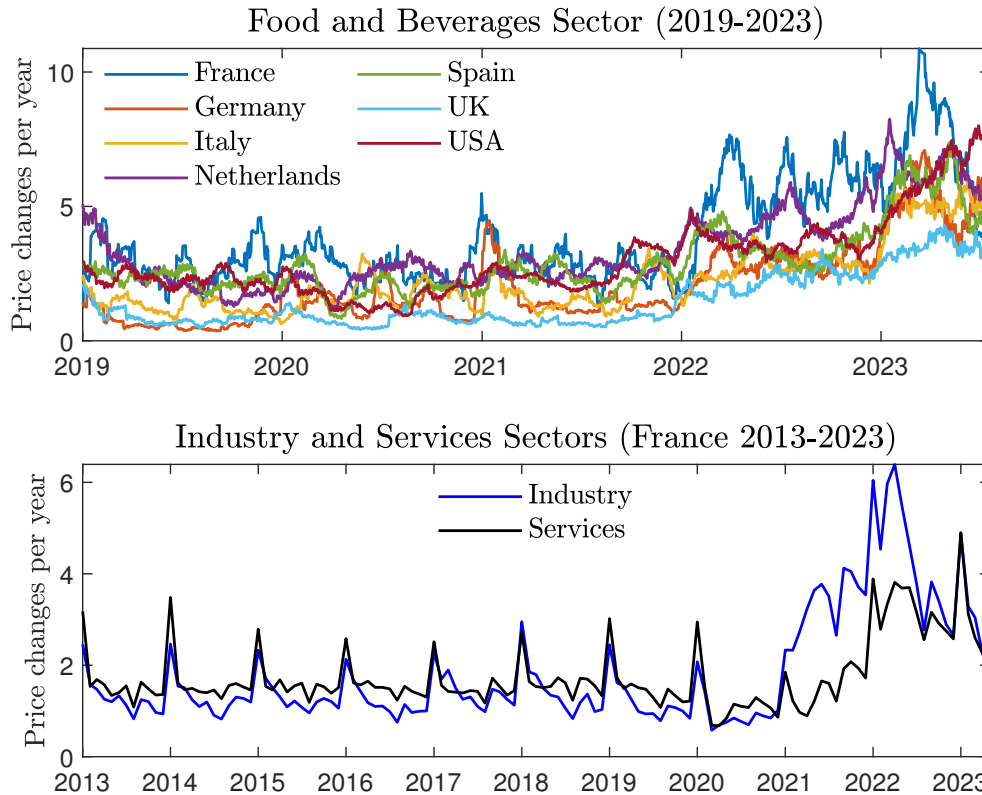
It has been shown that the distinction between TD and SD models may be irrelevant in the presence of small aggregate shocks, see e.g. Gertler and Leahy (2008) and Alvarez, Lippi, and Passadore (2016).³ The main message of this paper is that in the presence of a large shock these models are quite different: firms are responsive to a large shock in SD models, while pricing behavior is unresponsive to economic conditions in TD models.⁴ Our main finding is that large cost shocks induce firms to react faster than in normal times. This can be seen in Figure 1, which uses data on the frequency of price changes in the food and beverages sectors in several countries (top panel), and in the manufacturing and services sector in France (bottom panel). After several years of overall stationary behavior, the frequency of price changes increased significantly in 2022 and 2023, more than doubling across all sectors and countries. We show that such a pattern can be matched by a generalized Ss model with

³The equivalence, established by proposition 1 and 2 in Alvarez, Lippi, and Passadore (2016), states that the propagation of a small aggregate shock in a TD model coincides with the propagation in a SD model provided these models have the same frequency and kurtosis of price changes. See Klenow and Kryvtsov (2008) for an early analysis of shocks in state- versus time- dependent models.

⁴The size of the shock is measured relative to the standard deviation of price changes. Small cost shocks, such as those normally considered in VAR analyses, change marginal costs by less than 1%, which corresponds to less than 1/10 of the standard deviation of price changes. We will consider marginal costs shocks that are in the order of 10%, which corresponds to more than 1/2 of the standard deviation of price changes.

a non-negligible degree of state dependence.

Figure 1: Frequency of price changes



Note: The top panel uses data from PriceStats (see Section 3). The data source for the bottom panel is the Banque de France Monthly Business Survey (see Dedola et al. (2023)); we map the original data (share of firms changing prices per month) into number of price changes per year.

Other papers have shown that firms tend to adjust fast following a large shock: [Bonadio, Fischer, and Sauré \(2019\)](#) offer empirical evidence on the swift adjustment of currency invoicing after the large 2015 Swiss appreciation. [Karadi and Reiff \(2019\)](#) use a menu cost model to analyze the large increase of the VAT tax in Hungary in 2004 and 2006, showing that the aggregate price level responded quickly to the shock. The closest to our analysis is [Alvarez and Neumeyer \(2020\)](#), who study price dynamics in Argentina after large increases in the cost of utilities in 2014 and 2016 and a sharp exchange rate devaluation in 2018. They show that such large cost shocks make the economy behave like one with flexible prices. They use a menu cost model, specialized to an economy with a very high average inflation,

to interpret the result as a manifestation of state-dependent behavior. Our paper differs from the above in terms of the data, the context, and the model used to interpret the facts. Our data cover several countries with a low average inflation. On the modeling side, we employ a parsimonious generalized Ss model, characterized by 3 parameters, which allows us to fit the price-setting patterns observed before 2022, accurately capturing the shape of the size-distribution of the price changes (hence the standard deviation and the kurtosis) as well as the frequency of price changes. [Alvarez and Neumeyer \(2020\)](#) uses a canonical menu cost model where the distribution of price changes features no small adjustments.⁵

The paper is organized as follows. The next section outlines the model. [Section 3](#) derives the model's predictions for observables, compares them to the data, and presents a steady-state calibration of the model. [Section 4](#) uses the model to analyze the propagation of a large energy shock, and discusses the issue of asymmetric impulse responses (cost increases vs decreases). [Section 5](#) concludes.

2 A generalized setup for NK models

This section presents a New Keynesian setup that nests several models. The firm's price setting decision is described by a generalized hazard function, as in [Caballero and Engel \(2007\)](#), relating the probability of price adjustment to the firm's state, measured by the distance between the firm's current price from the profit maximizing price. The notion of a generalized hazard function was developed in seminal papers by [Caballero and Engel \(1993a,b\)](#), a derivation from first principles based on random adjustment costs was provided in [Caballero and Engel \(1999\)](#) and [Dotsey et al. \(1999\)](#), and later revisited using information theoretical foundations by [Woodford \(2009\)](#) and [Costain and Nakov \(2011\)](#).

⁵Price adjustments only occur at the boundary of the inaction region, and hence kurtosis is 1.

Household Preferences. Time is continuous and $t \in [0, \infty)$ and the preferences of the representative agent are given by the discounted present value

$$\int_0^\infty e^{-\rho t} U(C(t), H(t)) dt \quad \text{where} \quad U(C, H) \equiv \frac{C^{1-\epsilon}}{1-\epsilon} - \alpha H, \quad (1)$$

where $C(t)$ is an aggregate of $i \in [0, 1]$ goods and $H(t)$ is the labor supply. Let $c_i(t)$ be the consumption of product i at time t . The Dixit-Stiglitz consumption composite $C(t)$ is

$$C(t) \equiv \left[\int_0^1 A_i(t)^{\frac{1}{\eta}} c_i(t)^{\frac{\eta-1}{\eta}} di \right]^{\frac{\eta}{\eta-1}} \quad (2)$$

where $A_i(t)$ is a preference shock associated with good i at time t , which acts as a multiplicative shifter for the demand of each good. Labor supply is an aggregate of the labor that each firm hires so that $H(t) = \int_0^1 h_i(t) di$. Households maximize preferences in [equation \(1\)](#) subject to a budget constraint choosing consumption and labor for each good i and time t . There is a distortionary labor subsidy to allow the flexible price equilibrium to coincide with the efficient allocation.

The supply side. We consider an economy populated by a unit mass of firms, indexed by $i \in [0, 1]$, and each of them produces one good. For firm i to produce $c_i(t)$ of the good i at time t requires both labor (h_i) and energy (m_i) inputs, according to the production function

$$c_i(t) = \left(\frac{h_i(t)}{Z_i(t)} \right)^{1-\beta} m_i(t)^\beta \quad (3)$$

where $\beta \in (0, 1)$ is the energy share in production and firm i 's marginal cost of production at time t is: $mc_i(t) = KE(t)^\beta (W(t)Z_i(t))^{1-\beta}$ where $W(t)$ is the nominal wage, $E(t)$ is the price of the energy input, $Z_i(t)$ is a firm-specific process affecting productivity, and $K \equiv \beta^{-\beta}(1-\beta)^{\beta-1}$ is a constant. The technology exhibits constant returns to scale. We assume that $Z_i(t) = \exp(\sigma z_i(t))$ where $\{z_i\}$ are standard Brownian motions independent across i . We assume that $A_i(t) = Z_i(t)^{\eta-1}$ so the (log of) marginal cost and the preference

shock are perfectly correlated.⁶

We consider the profit maximization problem for a firm in steady state using the generalized hazard function of Caballero and Engel (1999). The setup embeds a broad class of sticky-price models, including well known cases such as the canonical Golosov and Lucas (2007), the pure Calvo (1983) model and the hybrid Calvo-Plus model by Nakamura and Steinsson (2010). The state of the firm x is given by its “price gap”, defined as the price currently charged by the firm relative to the price that maximizes current profits:

$$P_i^*(t) = \frac{\eta}{\eta - 1} \text{mc}_i(t)$$

namely the marginal cost times the constant markup $\frac{\eta}{\eta-1}$, implied by the CES demand system. Note that $P^*(t)$ depends on time because productivity is stochastic and because the marginal costs can change due to, for instance, an energy cost shock.

More precisely, the price gap $x_i(t)$ for firm i is the time t wedge between the actual price $P_i(t)$ and the desired price $P_i^*(t)$:

$$x_i(t) \equiv \log P_i(t) - \log P_i^*(t) \tag{4}$$

Absent pricing frictions the gap is identically zero, i.e. each firm charges the optimal price $P_i(t) = P_i^*(t)$. If the price is not adjusted, the price gap changes due to trend inflation (increasing nominal costs) and the idiosyncratic productivity shocks, so that the law of motion of the price gap for each i is

$$dx_i(t) = -\mu dt + \sigma dz_i(t), \tag{5}$$

where z_i is a standard Brownian motion and μ is the growth rate of nominal wages per unit

⁶This assumption, also used in Woodford (2009); Bonomo et al. (2010); Midrigan (2011), allows the problem to be described by a scalar stationary state variable, the price gap x . This is used to write the dynamic programming problem of the firm as well as to keep the expenditure shares stationary across goods in the presence of permanent idiosyncratic shocks.

of time i.e. the inflation rate.⁷

The price-setting problem. The firm's value function $v(x)$ solves the following problem:

$$\rho v(x) = \frac{\eta(\eta-1)}{2} x^2 - \mu v'(x) + \frac{\sigma^2}{2} v''(x) + \min_{\ell \geq 0} \{ \ell \cdot (v(x^*) - v(x)) + (\kappa \ell)^\gamma \} \quad (6)$$

where the quadratic term represents a second order expansion of the profit function around the profit-maximizing price, and x^* is the profit maximising reset price-gap that satisfies $v'(x^*) = 0$.⁸ At every moment the firm chooses an optimal effort rate ℓ for price resetting, so that with probability ℓ per unit of time the firm is able to adjust its price, i.e., to control x .⁹ The optimal choice balances the benefit of a price reset, given by $v(x^*) - v(x)$, with the cost of effort, given by the convex power function $(\kappa \ell)^\gamma$, with $\gamma > 1$ and $\kappa > 0$. The firm's problem is equivalent to profit maximization, aiming to minimize the deviations between realized profits and frictionless profits. We note that the units of the cost function are expressed as a fraction of forgone (steady state) profits.

Policy Rules. The price setting behavior implied by [equation \(6\)](#) is summarized by a *generalized hazard function* (GHF), $\Lambda : \mathbb{R} \rightarrow \mathbb{R}_+$. The function gives the probability (per unit of time) that a firm with price gap x will change its price. The first order condition gives the generalized hazard function $\Lambda(x) = \ell^*$, that is the optimal value of ℓ for a given value of x , namely:

$$\Lambda(x) = \kappa^{\frac{\gamma}{1-\gamma}} \left(\frac{v(x) - v(x^*)}{\gamma} \right)^{\frac{1}{\gamma-1}}, \quad \text{for each } x \in (-\infty, \infty) . \quad (7)$$

As in [Caballero and Engel \(2007\)](#) the policy rule implies that price changes occur prob-

⁷This is can be rationalized by assuming the money supply grows at an exogenous rate μ .

⁸The term $\frac{\eta(\eta-1)}{2}$ in the flow cost is related to the curvature of the profit function, see Proposition 1 in [Alvarez, Lippi, and Souganidis \(2023\)](#).

⁹We could enrich the model by assuming the firm can also change its price by paying the menu cost $\Psi > 0$. We ignore this possibility for simplicity but we note that little is lost in generality, as proven in Proposition 3 in [Alvarez, Lippi, and Oskolkov \(2022\)](#).

abilistically, with a probability that is governed by the GHF. Compared to the workhorse Calvo (1983) model, where the adjustment probability is constant, a generalized hazard function $\Lambda(x)$ allows it to depend on the state x , the firm’s desired adjustment. A large number of models are nested by this framework, including the canonical Calvo model with a constant hazard $\Lambda(x) = 1/\kappa$ as $\gamma \uparrow \infty$, and a canonical Ss model as $\gamma \downarrow 1$, where the hazard is flat (near zero) over a range of x and diverges at a critical threshold point. Intermediate cases cover the so called Calvo-Plus model by Nakamura and Steinsson (2010) and the random menu cost problem of Dotsey and Wolman (2020).

The value function $v(\cdot)$ and the generalized hazard function $\Lambda(\cdot)$ have a minimum at x^* and are increasing in $|x - x^*|$. Intuitively, the optimally chosen probability of adjustment is increasing in the distance between x and the optimal reset gap x^* (near zero due to a small drift and a symmetric return function): a larger distance increases the incentives to adjust, as shown in the first panel of Figure 2 for three model calibrations designed to fit Germany, the UK and the US (described below). The hazard functions in the figure show that the adjustment probability increases in the (absolute value) of $|x - x^*|$. This feature emerges for all calibrations and is the hallmark of state dependence. Our findings complement previous evidence on the relevance of state-dependent pricing, see e.g. Eichenbaum et al. (2011); Gautier et al. (2022); Karadi et al. (2023); Dedola et al. (2023).

Aggregation of the firm’s decisions. Focusing on a steady state, the hazard $\Lambda(x)$ and the law of motion of price gaps in equation (5) uniquely determine the steady state distribution of firm’s price gaps. Let us define this distribution by the density function $f : \mathbb{R} \rightarrow \mathbb{R}_+$ that solves the Kolmogorov forward equation:

$$\Lambda(x) \cdot f(x) = \mu f'(x) + \frac{\sigma^2}{2} f''(x), \quad \text{for each } x \neq x^*. \quad (8)$$

with boundary conditions $\lim_{x \downarrow x^*} f(x) = \lim_{x \uparrow x^*} f(x)$; $1 = \int_{-\infty}^{\infty} f(x) dx$, and $\lim_{x \rightarrow \infty} f(x) = \lim_{x \rightarrow -\infty} f(x) = 0$, see Alvarez and Lippi (2022) for details. The density function f , pictured

in the second panel of [Figure 2](#) (in dotted blue), depicts the steady state distribution of price gaps.

3 Price setting behavior: data vs theory

This section derives the model’s prediction for the frequency and the distribution of the size of price changes. We calibrate the model’s steady state using data for price setting behavior observed before 2021. We show that the GHF model is able to account for several key patterns observed in the data.

Frequency of price changes. The frequency of price adjustments N is given by

$$N = \int_{-\infty}^{\infty} f(x)\Lambda(x) dx \tag{9}$$

The equation counts the total number of price adjustments in a time period (say a year), by weighting the rate at which firms adjust at each x , $\Lambda(x)$, with the density of firms $f(x)$.

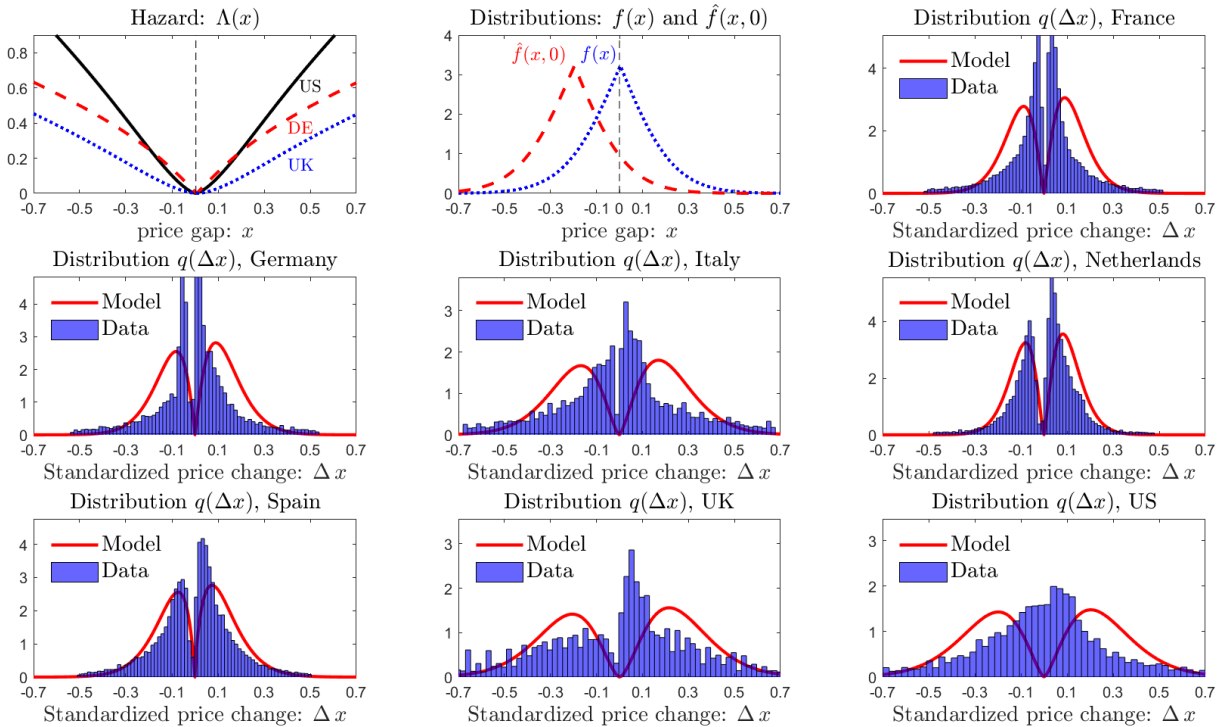
Distribution of the size of price changes. Recall that upon any price change the firm resets its gap from x to the optimally chosen return point x^* , hence the size of the adjustment is $\Delta x = x^* - x$. This occurs with probability $\Lambda(x)$ per unit of time. Recall also that at steady state, there is a density $f(x)$ of firms with price gap x . The distribution of the size of price changes has the following density:

$$q(\Delta x) \equiv \frac{\Lambda(x)f(x)}{N}. \tag{10}$$

Below we present a calibration of the model parameters that yields a distribution of price changes $q(\Delta x)$ that approximately matches the data. The panels of [Figure 2](#) show data from the European countries and the US which align to the predictions of the model, in particular the model ability to replicate a small mass of tiny price changes and the bimodality

of the distribution, a feature that is visible for most countries and that is crucial for state dependence, as we further discuss below.¹⁰

Figure 2: Hazard function $\Lambda(x)$, distribution $f(x)$, and distribution $q(\Delta x)$



Note: The first panel shows the calibrated hazard function (expressed as a monthly probability) for the US, Germany and the UK. The second panel shows the steady state distribution of price gaps before and after a large shock; it uses the model calibration for the food and beverages sector of all countries in our sample. The rest of the panels show the distribution of the size of the price changes, $q(\Delta x)$, standardized at the COICOP level along with the calibrated model-implied distribution. See Table 1 for more information on the calibration.

A brief description of the dataset. We base our analysis on granular data on price setting behavior, as in Cavallo (2018). These data contain detailed information on the frequency and size of daily price changes for a large number of firms, and provide the necessary information to solve the inverse inference problem mentioned above. Our data was provided

¹⁰We remark however that the calibrated model distribution and the data histogram use two different methods to handle unobserved heterogeneity in the size of productivity shocks: the histograms use data standardized at the COICOP level that retain unobserved heterogeneity within the same-COICOP firms. The model calibration uses a firm-level control method that provides a more accurate estimate of the shape (kurtosis) of the distribution.

by PriceStats, a private company related to the The Billion Prices Project (see [Cavallo and Rigobon \(2016\)](#)). It is generated using a technology known as web scraping, which automatically scans the code of publicly available webpages to gather and store relevant data. The dataset consists of product-level information, such as product id, price, category, and sale status, collected on a daily basis from various retailers’ websites. The data is uncensored and detailed, covering the entire lifespan of all products sold by these retailers, and provides prices that are similar to those obtained in offline stores ([Cavallo et al., 2018](#)). The data is comparable across countries, collected using identical techniques for similar categories of goods over the same time period. We use a subset of data from several European countries and the US, from January 1st 2019 to July 22nd 2023. We use a v-shape filter to identify sales and perform the analysis using regular prices as it is standard in the literature.

We focus on the “Food and Beverages” category which has experienced one of the highest rates of inflation during this period and has the largest weight in the goods CPI basket for most countries. Importantly, other sectors, such as industry and services, also record a similar increase in the frequency of price revisions, as depicted in the bottom panel of [Figure 1](#). We stress that the higher steady-state frequency of price changes in the “Food and Beverages” category, as compared to other sectors, does not tilt our results towards estimating larger state dependence. Obviously, a time-dependent model can also capture a higher steady state frequency (e.g., through a smaller κ), but it cannot capture a *change* of the frequency after an aggregate shock.

In [Table 1](#) we present summary statistics of price setting behavior for several European countries, the US and comparable statistics from related studies. The standard deviation of price changes varies from 15% to 31% across countries. This highlights the significant role of idiosyncratic productivity shocks. The kurtosis measure is very similar across all countries and ranges between 2 and 3.

The measurement of kurtosis. To measure kurtosis in [Table 1](#) we follow [Alvarez, Lippi, and Oskolkov \(2022\)](#) in accounting for unobserved heterogeneity. Kurtosis is a scale-free

Table 1: Price Setting Behavior before 2022

	Price-Setting Statistics					Model Calibration		
	Mean of Δx	St. Dev. of Δx	Kurtosis of Δx	Frequency N	# obs. (thousands)	σ	γ	κ
Food and Beverages data (PriceStats data, period 2019-21)								
France	-0.002	0.15	2.4	2	503	0.21	2.65	0.10
Germany	0.011	0.16	2.6	1.2	177	0.18	2.95	0.20
Italy	0.007	0.26	2.2	1.5	93	0.32	2.36	0.19
Netherlands	0.009	0.13	2.3	2.3	330	0.20	2.50	0.07
Spain	0.011	0.16	2.9	2.1	350	0.23	3.52	0.14
UK	0.007	0.31	2.1	0.8	95	0.28	2.23	0.39
US	0.012	0.31	2.2	2.1	282	0.45	2.37	0.16
Average	0.010	0.27	2.3	1.9	-	0.37	2.51	0.16
Euro area CPI data (PRISMA data, period 2005-19, Gautier et al. 2022)								
	0.025	0.10	2.8 ^(a)	1.0	-	0.10	3.25	0.20
Euro area supermarket data (IRI data, period 2013-17, Karadi et al. 2023)								
	-	0.13	3.2	1.1	-	0.14	4.17	0.28

Notes: The PriceStats data uses a sample of daily changes in regular prices (excluding sales). The statistics are computed over the distribution of the size of price changes after dropping price changes larger than 1.50 log points in absolute value and products with less than 3 price spells for the period 2019-2021. Kurtosis is computed using [equation \(11\)](#), a correction for unobserved heterogeneity discussed in the text. Frequency is annualized. The Average statistics are weighted by 2019 consumption at current USD. The statistics from the Price-setting Microdata Analysis (PRISMA) network are from Table 7 in [\(Gautier et al., 2022\)](#). These data covers the period from 2005 to 2019. Standard deviation and kurtosis figures were obtained in private exchanges with the authors. (a) The matched value of kurtosis for the PRISMA data of 4.1 is corrected for heterogeneity using a multiple from [Alvarez et al. \(2021\)](#) who perform the correction for French CPI data. The statistics from [Karadi et al. \(2023\)](#) are from their Table 2 and correspond to the average of 4 Euro Area countries; Germany, France, Italy and the Netherlands between 2013 and 2017. The parameters σ, γ, κ are calibrated using a GMM estimator to match the frequency and the kurtosis of price changes. The drift of price gaps is $\mu = 2\%$. The additional parameters are set to standard values: $\eta = 6, \rho = 0.05$.

statistic that can be largely affected by the mixing of distributions with different scales. For instance mixing normal distributions with different variances can yield a kurtosis much higher than 3 in spite of the fact that each of the underlying distribution is normal. In particular we assume that each product i within a category has a product-specific scale b_i , such that $\Delta x_{it} = b_i \Delta x$, kurtosis can be computed accounting for unobserved heterogeneity using

$$Kurt(\Delta x_t) = \frac{\mathbb{E}[(\Delta x_{it})^4]}{\mathbb{E}[(\Delta x_{it})^2(\Delta x_{is})^2]}, \quad (11)$$

where the empirical counterparts are estimated with

$$\begin{aligned} \widehat{\mathbb{E}}[(\Delta x_{it})^4] &= \frac{1}{\#I} \sum_{i \in I} \frac{1}{\#T(i)} \sum_{t \in T(i)} (\Delta x_{it})^4, \\ \widehat{\mathbb{E}}[(\Delta x_{it})^2(\Delta x_{is})^2] &= \frac{1}{\#I} \sum_{i \in I} \frac{1}{\#T(i)(\#T(i) - 1)} \sum_{t, s \in T(i), t \neq s} (\Delta x_{it})^2 (\Delta x_{is})^2, \end{aligned}$$

where $\#I$ are the number of products, $T(i)$ is the set of times where product i changes prices and $\#T(i)$ are the number of price changes for product i . The unobserved heterogeneity in our data also informs the standardization of price changes in the histograms of [Figure 2](#) since our model assumes equal size of idiosyncratic shocks across all firms (see footnote 10).

Calibration. We calibrate the model to match three cross sectional moments measured before 2022: the standard deviation and the kurtosis of the size of price changes, as well as the frequency of price changes. Our choice of moments is guided by theory: the frequency of price changes gives general information about price flexibility and thus the level of the cost of price management κ . Moreover, when analysed together with the standard deviation of price changes, they allow us to infer the size of idiosyncratic productivity shocks affecting firms using the fact that $\sigma^2 = N \cdot Var(\Delta x)$ as this relationship holds for a wide variety of models when $\mu \approx 0$, see [Alvarez, Lippi, and Oskolkov \(2022\)](#). In turn, kurtosis has been shown to reveal key information about state dependence and the response of an economy to aggregate

shocks. In our model, this is captured by the curvature of the hazard function controlled by γ . Kurtosis encodes information about state dependence and the “selection” effect: the idea that the observed price changes come from firms who need it the most and not from a random sample. A large kurtosis indicates a large mass of late adjusters, which implies weaker state-dependence and more persistent effects of a cost shock. The Calvo model, with a kurtosis of 6, does not feature selection nor state dependence since the adjusters are a random sample each period. In contrast, a pure menu cost model, with a kurtosis of 1, features selection and maximal state dependence since all price changes share the same (absolute) size, and they stem from the sample of firms that are the furthest from their optimal price.

We use standard values for the parameters of the elasticity of substitution and intertemporal preferences: $\eta = 6$ (which implies a markup of 20%), a time discount $\rho = 0.05$ and an inflation rate of $\mu = 2\%$ consistent with inflation at steady state. Then the three moments discussed above are used to calibrate $\{\sigma, \kappa, \gamma\}$. We use the frequency and the variance of price changes to calibrate σ^2 and we use a GMM estimator to calibrate κ and γ from the effort cost function. In short, the calibration consists of choosing these two parameters to match the kurtosis and the frequency of price changes. The calibrated parameters are shown in the last three columns of [Table 1](#). We find a non-negligible degree of state dependence, represented by (a small value of) the parameter γ , identified by a relatively small kurtosis of price changes. This finding is robust across countries. The first panel of [Figure 2](#) depicts this fact: the calibrated shape of the hazard function is similar across countries, displaying a comparable curvature (the parameter γ) and somewhat differing levels (the parameter κ) to match the different mean frequency of price changes.

[Figure 2](#) shows that the calibrated model is able to capture some key features of the data: the fact that the distribution of price changes $q(\Delta x)$ is bimodal, with a small mass of tiny price changes. Intuitively, when x is close to the optimal return gap x^* , the losses from the suboptimal price are low and hence the incentives to change prices are small. The small mass of tiny price changes is a major difference compared to the prediction of the Calvo model

where the constant hazard implies a mode near zero, i.e. that the most frequently observed price change has a tiny size. This prediction is counterfactual and is a telltale of the fact that price setting behavior displays state dependence: prices are adjusted only when the benefits of doing so are large enough. Indeed, we note that the bimodality of the distribution of the size of price changes cannot be produced by a time-dependent model. The bimodality implies that the GHF is not constant i.e. it implies state dependence.¹¹

4 The propagation of an aggregate cost shock

In [Figure 1](#) we observed that the frequency of price adjustments rose quickly after a large energy shock. In this section, we provide a thought experiment that rationalizes this fact using the model. Namely, we consider an economy at steady state and study the dynamics that follow an aggregate shock to the firms’ marginal cost. We use the calibration to the food and beverages sector of all countries in our sample to study the propagation of large and small shocks. We will illustrate that, under a calibrated GHF model, large and small shocks have different implications for pass-through and the frequency of price adjustments.

We consider an economy in a steady state that is hit by an unexpected once-and-for-all shock to marginal cost that displaces the distribution of states δ percentage points “to the left”, as shown by the red distribution in the second panel of [Figure 2](#). The firms would then like to increase their prices to “close their price gap”. This incentive shapes the dynamic response of the price level and the frequency of price changes after the shock. For simplicity we develop the dynamic analysis assuming that firms use the steady state decision rule, $\Lambda(x)$.¹² We let $\hat{f}(x, t) : \mathbb{R} \times [0, \infty) \rightarrow \mathbb{R}_+$ be the distribution of price gaps t periods after the shock. Intuitively, the distribution depends on the time elapsed because the initial shock places the distribution away from the invariant. The distribution $\hat{f}(x, t)$ solves a Kolmogorov

¹¹Notice that the [equation \(8\)](#) for $\mu = 0$ shows that f is unimodal at 0 and strictly convex. Since q is the product of f and Λ , a constant GHF would imply a unimodal q .

¹²This assumption is often used in the literature, see [Baley and Blanco \(2021\)](#). Solving the full problem, with feedback from the aggregates to the individual firm’s decision can be done using the methods in [Alvarez, Lippi, and Souganidis \(2023\)](#).

forward equation with initial condition $\hat{f}(x, 0) = f(x + \delta)$, for each $x \in \mathbb{R}$. We will describe the transition of these variables back to steady state for small and large shocks (see Section 4 of [Alvarez and Lippi \(2022\)](#) for details).

In the second panel of [Figure 2](#) we plot the distribution of price gaps f , centered around the vertical line $x = x^*$, and the distribution of firms' price gaps after the shock $\hat{f}(x, 0)$. We note that since inflation is small, namely 2%, the value of $x = x^*$ is about zero. We see that the distribution after the shock places most firms' prices in a region far from their desired prices (at around -20% price gap). In this region the probability of adjustment, $\Lambda(x)$, is larger, so a large shock triggers an increase in number of price adjustments.

Using $\hat{f}(x, t)$ we can describe the path for the price level and the frequency of price changes after an aggregate shock, as

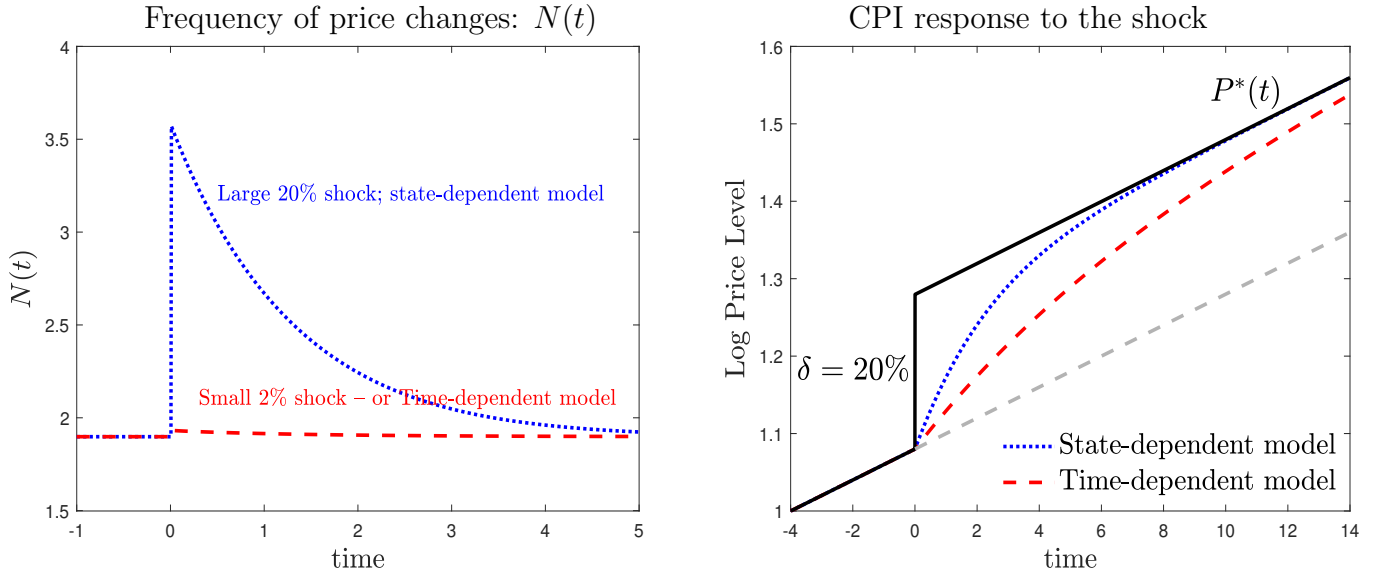
$$P(t) \approx P(0) + \int_{-\infty}^{\infty} x \cdot \hat{f}(x, t) dx - (X_{ss} - \delta), \quad (12)$$

$$N(t) = \int_{-\infty}^{\infty} \Lambda(x) \cdot \hat{f}(x, t) dx \quad (13)$$

where $X_{ss} = \int_{\mathbb{R}} x \cdot f(x) dx$. The right hand side of [equation \(12\)](#) can be written more succinctly as $P(0) + X(t) - (X_{ss} - \delta)$ where $X(t)$ denotes the mean price gap at time t . Notice that the change in the price level is computed as deviations from $X_{ss} - \delta$ because that is the aggregate price gap right after the shock occurs, so that $P(0)$ on impact (before any price change happens) does not jump.

The left panel of [Figure 3](#) displays a sharp increase in the frequency of price changes after a large shock, reaching a magnitude comparable to the average frequency of 4.3 price changes per year in 2022-2023 for all countries in our sample. In the model this is due to many firms lying in a region far from their desired gap x^* , i.e. a region where the hazard, $\Lambda(x)$, is relatively high. This yields a persistent increase in the frequency of price adjustments. These results are robust in our sample of countries and appear as long as the GHF is far from constant. Obviously this effect is not present in the Calvo model, where the frequency

Figure 3: Propagation of Large vs Small Cost Shock



Note: The model uses the calibration for the food and beverages sector of all countries in our sample, see [Table 1](#). The calibration matches a frequency of $N = 1.9$ price changes a year, the kurtosis and standard deviation of price changes of 2.3 and 27% respectively. Time is depicted in months after the shock.

is constant by assumption.

Inflation dynamics. The right panel of [Figure 3](#) shows that the propagation of a large shock features a faster pass-through than one predicted by a Calvo model matching the same frequency (and size) of price changes. This result mirrors the observation about frequency. Upon a large shock the calibrated model is characterized by more firms adjusting prices upwards resulting in faster passthrough. Inspecting the slope of the price path, the figure illustrates that failing to account for the large increase in the frequency of price revisions, as would occur in a time-dependent model, leads to a substantial initial underestimation of inflation followed by a subsequent overestimation.

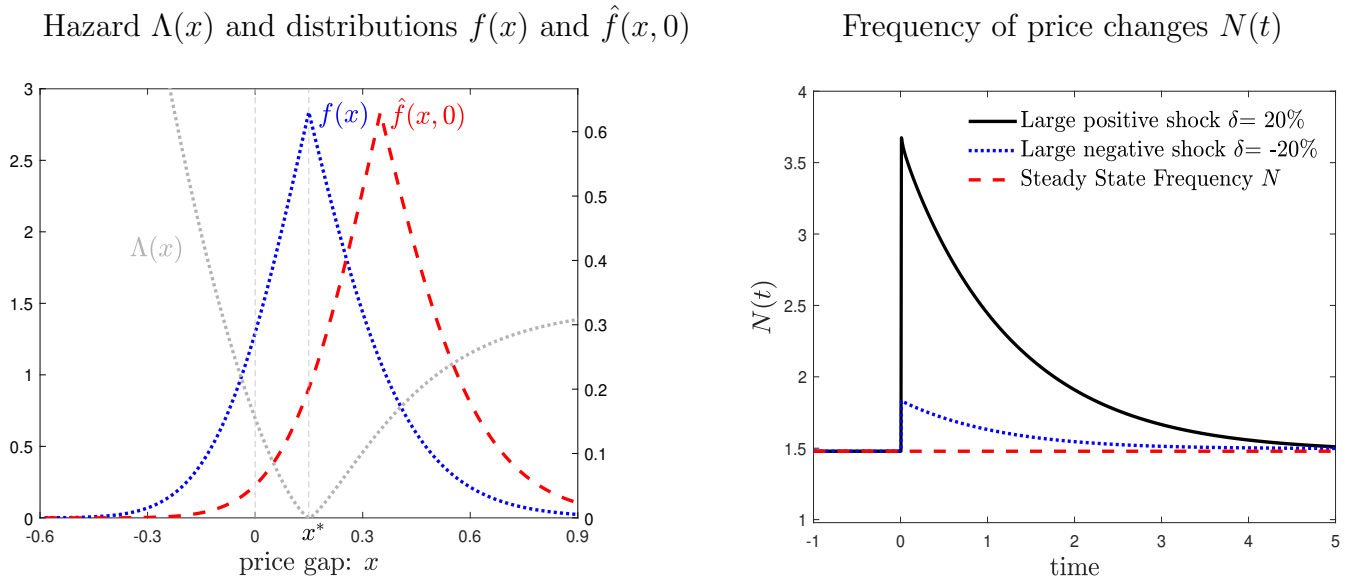
We note that one feature that is not properly captured by our simple model is the half-life of the shock, which is faster than observed in the data. Enriching the model to have strategic complementarities in price setting would slow down the propagation of the shock, while retaining the state dependence of the firms' decision and thus preserve a sizable response

along the extensive margin, as shown in Alvarez, Lippi, and Souganidis (2023). Such an extension, however, would come at the cost of a substantially more involved model and we leave it for future work.

In summary, the state-dependent model features dynamic responses to large cost shocks that resemble the data on the increase in the frequency of adjustments for the recent surge in inflation. Furthermore, as we have shown, the forecast of the frequency of price changes and the path of inflation can be very different depending on the model that the analyst is using. The implications of a purely time-dependent model are counterfactual because they miss the extensive margin response documented in Figure 1.

Asymmetric responses to large shocks. Starting in the last quarter of 2022, there has been a rapid decrease in energy prices, prompting the question of how swiftly firms will respond by adjusting consumer prices downward – a crucial consideration for predicting inflation dynamics.

Figure 4: Asymmetries: Impulse response to negative vs. positive cost shock



Note: The model uses the calibration for the food and beverages sector of all countries in our sample, see Table 1, and the exact profit-loss function. The left panel has a dual scale: the left axis gives the density of price gaps; the right axis gives the hazard rate of a price gap (as the probability per month).

Many features of the economic environment may induce an asymmetric response to increasing and decreasing cost shocks, such as sizable trend inflation or variable markups. In this section we will focus on the asymmetry of the firms' profit function. The asymmetry implies that having a price that is $x\%$ above the ideal price is better than having a price that is $x\%$ below the ideal. Such an asymmetric profit function arises in several standard models, such as in a CES demand system. Intuitively, having a "high price" is preferable to having a "low one" because in the latter case the firm is selling many units at close or below marginal cost. Thus, low prices give rise to a higher incentive to adjust negative price gaps.

The left panel of [Figure 4](#) depicts the above observation by means of the generalized hazard function. This GHF is obtained by solving the firm problem using the exact profit function, rather than the second order approximation (which is symmetric by construction) described in [equation \(6\)](#). Recall that the GHF is an increasing function of the incentive to adjust prices $v(x) - v(x^*)$. Therefore the left panel of [Figure 4](#) shows that the firm's incentive to adjust prices is smaller when the prices are above the target (decrease prices) compared to when prices are below (increase prices). The asymmetric profit function also affects the optimal reset price chosen by firms. The left panel shows that the reset price is located above the ideal price ($x^* > 0$), in order to hedge against productivity shocks that could quickly lead to negative gaps.

We study the propagation of a cost shock decrease of 20%, a negative shock henceforth, and we calibrate using food and beverages data for all countries, differing only in using the profit-loss function instead of its second-order approximation. The right panel of [Figure 4](#) shows the asymmetric response of the frequency of price changes to a large positive vs a large negative shock to marginal costs. Unlike a positive shock (thick line), a negative shock (dotted line) leads to a less pronounced rise in the frequency of price changes due to lower probabilities of adjustments for positive price gaps, stemming from lesser profit losses at higher prices. Although the passthrough of a negative shock is relatively slower than its positive counterpart, it continues to be faster than the passthrough predicted by a purely

time-dependent model.

In short, our model with asymmetries shows that prices are more flexible upon a cost increase than upon a cost decrease. This result might contribute to understanding the recent price setting behavior where energy prices have started to subside. It is also consistent with the reduction in the frequency of price changes for some countries (see [Figure 1](#)), implying a slower passthrough of the cost decrease to consumers.

5 Conclusions

We use a tractable New Keynesian model, calibrated using granular data from several countries, to study the propagation of a large supply shock. The model matches some key features of the data, including the (pre-shock) bi-modal distribution of the size of price changes and the significant increase in the frequency of price changes following a large shock. The findings underscore the state-dependent nature of price setting decisions, where the probability of a price adjustment increases as the gap between the actual and ideal price widens.

The main implication of our research is that large shocks lead to a rapid pass-through from costs to prices causing a temporary surge in inflation, i.e., we show that large shocks travel fast. Understanding the different firm-level reactions to large and small shocks is crucial to predict inflation dynamics following a large shock. Our findings highlight the limitations of time-dependent models, commonly used by central banks. By failing to account for the increased frequency of price changes, the time-dependent models are unable to capture the fast inflation run up after the shock, as well as the subsequent quick slowdown.

Future research should extend the analysis to more categories of goods, and explore the role of strategic complementarities in shaping the speed of pass-through of large shocks.

References

- Alvarez, F. and F. Lippi (2022). The analytic theory of a monetary shock. *Econometrica* 90(4), 1655–1680.
- Alvarez, F. and A. Neumeyer (2020, December). The Pass-Through of Large Cost Shocks in an Inflationary Economy. In G. Castex, J. Gali, and S. Diego (Eds.), *Changing Inflation Dynamics, Evolving Monetary Policy*, Volume 27 of *Series on Central Banking, Analysis, and Economic Policies*, pp. 7–48. Central Bank of Chile.
- Alvarez, F. E., A. Ferrara, E. Gautier, H. LeBihan, and F. Lippi (2021, November). Empirical investigation of a sufficient statistic for monetary shocks. Working Paper 29490, NBER.
- Alvarez, F. E., F. Lippi, and A. Oskolkov (2022, May). The macroeconomics of sticky prices with generalized hazard functions. *The Quarterly Journal of Economics* 137(2), 989–1038.
- Alvarez, F. E., F. Lippi, and J. Passadore (2016, June). Are state and time dependent models really different? *NBER Macroeconomics Annual* 31, 379–457.
- Alvarez, F. E., F. Lippi, and P. Souganidis (2023). Price setting with strategic complementarities as mean field games. *Econometrica* forthcoming.
- Baley, I. and A. Blanco (2021). Aggregate dynamics in lumpy economies. *Econometrica* 89(3), 1235–1264.
- Bonadio, B., A. M. Fischer, and P. Sauré (2019, 03). The Speed of Exchange Rate Pass-Through. *Journal of the European Economic Association* 18(1), 506–538.
- Bonomo, M., C. Carvalho, and R. Garcia (2010). State-dependent pricing under infrequent information: a unified framework. Staff Reports 455, Federal Reserve Bank of New York.
- Caballero, R. J. and E. M. Engel (1993a). Heterogeneity and output fluctuations in a dynamic menu-cost economy. *The Review of Economic Studies* 60(1), 95.
- Caballero, R. J. and E. M. Engel (2007, September). Price stickiness in ss models: New interpretations of old results. *Journal of Monetary Economics* 54(Supplement), 100–121.
- Caballero, R. J. and E. M. R. A. Engel (1993b). Microeconomic adjustment hazards and aggregate dynamics. *The Quarterly Journal of Economics* 108(2), 359–383.
- Caballero, R. J. and E. M. R. A. Engel (1999, July). Explaining investment dynamics in u.s. manufacturing: A generalized (s,s) approach. *Econometrica* 67(4), 783–826.
- Calvo, G. A. (1983, September). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12(3), 383–398.
- Cavallo, A. (2018). Scraped data and sticky prices. *Review of Economics and Statistics* 100(1), 105–119.

- Cavallo, A., W. E. Diewert, R. C. Feenstra, R. Inklaar, and M. P. Timmer (2018). Using online prices for measuring real consumption across countries. *AEA Papers and Proceedings* 108, 483–487.
- Cavallo, A. and R. Rigobon (2016). The Billion Prices Project: Using online prices for measurement and research . Working paper 22111, NBER.
- Cavallo, A. F. (2019). More amazon effects: Online competition and pricing behaviors. *ackson Hole Economic Symposium Conference Proceedings*.
- Costain, J. and A. Nakov (2011, 03). Price adjustments in a general model of state-dependent pricing. *Journal of Money, Credit and Banking* 43, 385–406.
- Dedola, L., E. Gautier, A. Nakov, S. Santoro, E. De Veirman, L. Henkel, and B. Fagan-dini (2023). Some implications of micro price setting evidence for inflation dynamics and monetary transmission. Occasional paper series 3XX, European Central Bank.
- Devereux, M. B. and J. Yetman (2010, February). Price adjustment and exchange rate pass-through. *Journal of International Money and Finance* 29(1), 181–200.
- Dotsey, M., R. G. King, and A. L. Wolman (1999, May). State-dependent pricing and the general equilibrium dynamics of money and output. *The Quarterly Journal of Economics* 114(2), 655–690.
- Dotsey, M. and A. L. Wolman (2020). Investigating nonneutrality in a state-dependent pricing model with firm-level productivity shocks. *International Economic Review* 61(1), 159–188.
- Eichenbaum, M., N. Jaimovich, and S. Rebelo (2011, February). Reference prices, costs, and nominal rigidities. *American Economic Review* 101(1), 234–62.
- Gagnon, E., D. Lopez-Salido, and N. Vincent (2012, June). Individual price adjustment along the extensive margin. *NBER Macroeconomics Annual* 27, 235–281.
- Gautier, E., C. Conflitti, R. P. Faber, B. Fabo, L. Fadejeva, V. Jouvanceau, J.-O. Menz, T. Messner, P. Petroulas, P. Roldan-Blanco, et al. (2022, June). New facts on consumer price rigidity in the euro area. Working paper series 2669, European Central Bank.
- Gautier, E., M. Marx, and P. Vertier (2022). How do gasoline prices respond to a cost shock? Working paper 861, Banque de France.
- Gertler, M. and J. Leahy (2008, 06). A Phillips Curve with an Ss Foundation. *Journal of Political Economy* 116(3), 533–572.
- Golosov, M. and R. E. J. Lucas (2007). Menu Costs and Phillips Curves. *Journal of Political Economy* 115, 171–199.
- Gopinath, G. and O. Itskhoki (2010, 05). Frequency of Price Adjustment and Pass-Through. *The Quarterly Journal of Economics* 125(2), 675–727.

- Karadi, P., J. Amann, J. S. Bachiller, P. Seiler, and J. Wursten (2023). Price Setting on the Two Sides of the Atlantic - Evidence from Supermarket Scanner Data. *The Journal of Monetary Economics* –(forthcoming), –.
- Karadi, P. and A. Reiff (2019, July). Menu costs, aggregate fluctuations, and large shocks. *American Economic Journal: Macroeconomics* 11(3), 111–46.
- Klenow, P. J. and O. Kryvtsov (2008, August). State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation? *The Quarterly Journal of Economics* 123(3), 863–904.
- Midrigan, V. (2011, July). Menu costs, multi-product firms, and aggregate fluctuations. *Econometrica*, 79(4), 1139–1180.
- Nakamura, E. and J. Steinsson (2010, August). Monetary non-neutrality in a multisector menu cost model. *The Quarterly Journal of Economics* 125(3), 961–1013.
- Woodford, M. (2009, October). Information-constrained state-dependent pricing. *Journal of Monetary Economics* 56, s100–s124.