Inflation and misallocation in New Keynesian models

By Alberto Cavallo[1], Francesco Lippi[2] and Ken Miyahara[3] [4]

Abstract

The New Keynesian framework implies that sluggish price adjustment results in a distorted allocation of resources. We use a tractable model to identify these unobservable distortions, using granular data that depict the price-setting behavior of firms. We propose a method to estimate welfare costs for the period preceding 2022, and during the subsequent high inflation period. Using granular data from PriceStats, as well as data from the ECB PRISMA project, we find that these welfare costs are sizeable. In the low inflation environment prevalent before 2022, the efficiency cost is quantified in about 2 percentage points of GDP in the Euro Area. Moreover, we estimate that the recent inflationary shock has temporarily increased these costs, in the order of an additional 2.7 percentage points of GDP.

1 Introduction

The inflation surge that followed the sizeable increases in energy prices in Europe, pictured in Chart 1, has revived interest in inflation and its welfare costs. After more than three decades of stable prices, in 2022 inflation peaked at about 10% in the Euro area and the US. The inflation spike was associated with an unusually large increase in the frequency of price revisions. In the food-and-beverage sector, where the average number of regular price adjustments per year hovered between 1 to 3 before 2022, the number of price changes almost doubled in 2022 and 2023 (see Chart 2). These facts raise several questions about the mechanism by which cost shocks transmit to consumer prices, the welfare costs of inflation, the future path of inflation, and the appropriate policy actions.

[1] Harvard Business School.

[2] Luiss University and EIEF.

[3] Luiss University and EIEF.

[4] First draft, May 2023. We thank, without implicating, our discussant Sharon Kozicki, an anonymous ECB reviewer and 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, Luigi Paciello, and Facundo Piguillem. Tomas Pacheco provided excellent research assistance. Lippi acknowledges financial support from the ERC grant 101054421-DCS.

Chart 1

The 2019-2023 dynamics of energy prices

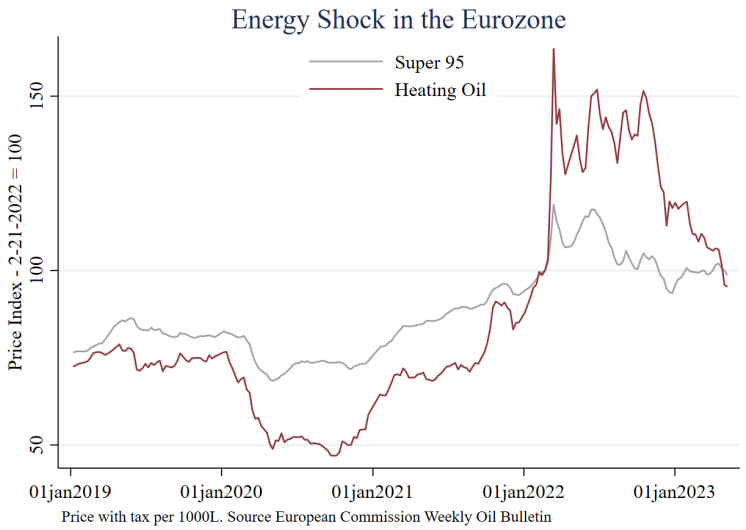

Sources: European Commission Weekly Oil Bulletin
Notes: Price with tax per 1000L.

In this paper we address two narrow and well-defined questions. The first one concerns the mechanism: understanding inflation dynamics requires a model of the price-setting decisions of a large number of individual firms, and their aggregation. We will use a simple model, validated against a rich granular dataset, to argue that the workhorse New Keynesian model of monetary policy in use at most central banks, built around the assumption of a constant average frequency of repricing by firms, misses a key aspect of the inflation dynamics. We will show that such model fails to capture the sizeable different speed of passthrough of large versus small aggregate shocks. We present a simple tractable model that improves upon the benchmark and yields smaller inflation forecast errors. In other words, our preferred model features a “highly non-linear” Phillips curve.

Chart 2

Frequency of price changes (Food and Beverages)

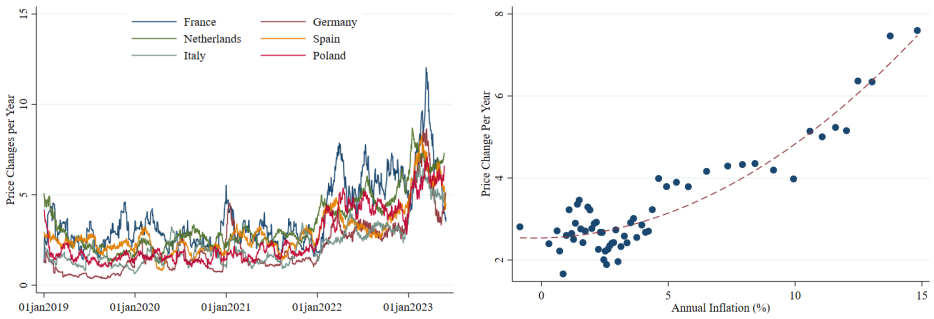

Sources: PriceStats data for regular price changes.
Notes: The left graph shows the frequency of price changes for regular prices (excluding sales) for six selected countries. The right graph is a binned scatterplot of the frequency and annual inflation rates for nine countries: France, Germany, Greece, Ireland, Italy, Netherlands, Spain, Poland, and the UK. Frequency computed using microdata. Annual inflation rates computed by PriceStats.

The second question concerns the welfare costs. A classic approach in macroeconomics studies the welfare cost of inflation using a powerful public finance idea, namely that inflation acts as a distortionary tax on the demand of real balances, see e.g., Bailey (1956); Friedman (1968). A microfounded money demand model can then be used to quantify these welfare costs, as in e.g., Lucas (2000). In this paper, we focus instead on the welfare costs that arise in New Keynesian (NK) models, by far the dominant framework employed by academics and central bank researchers in recent decades, see e.g., Woodford (2003); Gali (2008); Walsh (2010). In the NK framework, the welfare costs are made of two elements, both related to the assumption of sticky prices. First, since prices deviate from their efficient level, such wedges impose a welfare cost to consumers and workers. This is what the literature refers to as "misallocation", which we denote by the variable χ. Second, the costs associated with the price-management activities are wasteful, much like the resources that agents waste to protect themselves from inflation in the money demand models cited above. We denote this welfare cost by the variable φ. Both χ and φ will be measured as a proportion of total GDP so that their magnitude has a straightforward interpretation.

As neither measure of welfare cost is directly observable to researchers, their assessment requires a model, providing us with an explicit mapping between these objects and the observable data moments.[5] In the first part of the paper, we set up such a model, drawing on Caballero and Engel (1999,2007), and parametrize it using a granular data set for the food and beverages sector for several European countries, see Cavallo (2018). A founding principle of our analysis is to identify a model that is broadly consistent with the recent observed price-setting behavior. There are two reasons why this is important. First, the credibility of the analysis on retail price inflation requires that the model is consistent with the facts about price-setting behavior by retailers. Second, the welfare costs vary substantially across models, in spite of the fact that these models reproduce the same mean frequency and mean size of price changes. For instance, the welfare costs of misallocation in the well-known Calvo model, at the core of most NK analyses, are two times larger than the misallocation produced by a staggered-price adjustment model a la Taylor (1980), and six times larger than the misallocation produced by a menu cost model a la Golosov and Lucas (2007).

Matching the model fundamentals to the price-setting patterns observed in the granular data will lead us to reject the Calvo model, because of its impossibility to account for the significant increase in the frequency of price-setting observed in the data, shown in Chart 2, and because of its failure to fit other features of price-setting behavior, such as the bi-modal size distribution of the price changes depicted in Chart 3. This result yields an important policy lesson: large shocks travel faster than small shocks. Failure to acknowledge this fact will lead the policymaker to a wrong inference about inflation dynamics. Relatively to our simple model, the use of a textbook Calvo model leads to underpredicting inflation in the early periods after the shock, and to underestimating its slowdown later (see Chart 10).

[5] See Zbaracki et al. (2004) for an attempt at measuring such costs directly. We relate to their findings with ours in Section 4.

Chart 3

Distribution of Size of Price Adjustments 2020 (Food and Beverages)

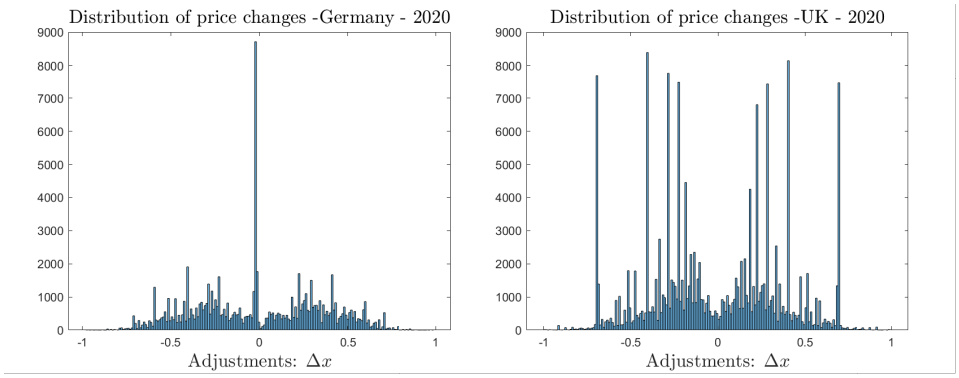

Source: PriceStats microdata for regular price changes.

In the second part of the paper, we use our model to infer the magnitude of the welfare costs caused by the presence of sticky prices, and analyze how these costs change following a large inflationary shock such as the recent large increase in energy prices. We supplement the granular data for the food and beverages industry with some descriptive statistics for the Euro Area, taken from the Price-setting Microdata Analysis (PRISMA) network, see e.g., Gautier et al. (2022), and descriptive statistics for supermarket data of the Euro Area, drawn from Karadi et al. (2023). We develop two conceptually distinct exercises. First, we measure the welfare costs χ and φ in a steady state. In particular, we gauge these costs using data from the low inflation period prevalent before 2022. The results suggest that the welfare costs in this low inflation environment range at about 1.9 percentage points of GDP in the Euro Area. This suggests that the resources lost every year due to the sticky price frictions are not negligible. Second, we analyze the dynamics of these welfare costs following the large energy price increases recorded in 2022. This exercise is a canonical impulse-response analysis, studying how inflation and welfare costs evolve from their steady state levels following a large cost shock. The exercise allows us to quantify the welfare costs that arise above and beyond the steady state costs. Our preliminary estimates suggest that the recent inflationary shock triggered a temporary increase in the welfare costs, on the order of an additional 2.7 percentage points of GDP.

Structure and overview of contents. The paper is organized as follows. Section 2 presents the New Keynesian setup that guides our analysis of the price-setting activity of firms and will (later) be used to quantify the welfare costs. The model is inspired by the seminal work of Caballero and Engel (1999, 2007) and nests several well-known cases such as the Calvo (1983) model or the menu cost model of Golosov and Lucas (2007).

Section 3 describes the model's predictions for the frequency and the size-distribution of price changes and compares them with cross-sectional facts observed from the low inflation period before 2022. This part of the analysis relies on a granular dataset for European countries provided by PriceStats. The 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. For our purposes, the dataset offers two key advantages over traditional data sources such as Consumer Price Index (CPI) and Scanner Data. Firstly, the daily price data collection with uncensored spells allows for an accurate identification of sales and price changes (Cavallo, 2018). Second, the data is available without any lags, allowing us to study the recent period of high inflation in real time.[6]

Mapping the model to the observables allows us to select a data-consistent structural model of price-setting. We show that a main feature of the selected price-setting model is a sizeable component of state-dependent decisions. This means that the firms' responsiveness to the shocks depends on the size of the shocks[7]. Following a large shock, such as the recent energy shock, firms react faster than in normal times. This is important to understand the dynamics of inflation, as also noted by Alvarez and Neumeyer (2019); Karadi and Reiff (2019). This finding also differs markedly from the time-dependent models widely used at central banks, such as the workhorse model of Calvo (1983). We show that the selected model can qualitatively replicate the response of the frequency of price changes after a large cost shock, such as the ones recently observed, as in e.g., Chart 2. This result is important to quantify the deadweight losses triggered by the shock, as firms must engage in above-normal repricing activities that are costly to them (as captured by φ).

Section 4 quantifies the welfare cost of misallocation in NK models. We first derive a mapping between the theory-based measure of misallocation, χ and φ, and a set of observable moments from the size distribution of price changes. For instance, we show that the welfare cost of misallocation, χ, is proportional to the product of the variance times the kurtosis of price changes. This provides a direct mapping to quantify the welfare costs of inflation, that we implement using three different granular datasets. As noted above, this result also highlights that different models can lead to estimates of misallocation that differ by an order of magnitude (e.g., the welfare costs in the Calvo model are six times larger than in the canonical menu cost model) as captured by the kurtosis of price changes.

We then use PriceStats data, as well as data provided from the ECB PRISMA project (Gautier et al., 2022), to estimate the welfare cost of inflation in the low inflation environment before 2022 to be about 1.9 percentage points of GDP in the Euro Area. The bulk of this welfare cost originates from the sizeable misallocation that is estimated in the data (the cost component χ). A smaller part, about 40 basis points of GDP, relates to the resources that are used for the price-management activity (the cost component φ). We also estimate that the recent inflationary shock has triggered a temporary increase of the welfare costs, in the order of an additional 2.7 percentage points of GDP. This temporary surge is attributed to a roughly equal

---

6 Our data was provided by PriceStats, a private company related to The Billion Prices Project (see Cavallo and Rigobon (2016)). This dataset is a comprehensive collection of retail prices obtained from the websites of large, multichannel retailers. It is generated using a technology known as web scraping, which automatically scans the code of publicly available webpages daily to gather and store relevant data.

7 See Gagnon (2009); Alvarez et al. (2019a); Karadi et al. (2023) for an extensive documentation of the importance of state dependent pricing in several countries including the US and Europe. The idea of using large shocks to discuss model selection has been used by other authors, such as Gopinath and Itskhoki (2010); Alvarez, Lippi, and Paciello (2016); Bonadio, Fischer, and Sauré (2019).

change in both cost components, each by approximately 1.5 percentage points of GDP. Provided no new large shocks arrive, this additional cost component is expected to vanish, and inflation is expected to return to the baseline.

1.1 Related literature

As mentioned, several classic contributions on the welfare cost of anticipated inflation consider that inflation acts as a distortionary tax on the demand for real balances (Bailey, 1956; Friedman, 1968). Several papers have taken those ideas seriously and developed carefully designed models to quantify the deadweight losses caused by a stationary inflation rate, as in e.g., Aiyagari, Braun, and Eckstein (1998); Lucas (2000); Lagos and Wright (2005). A common finding is that the welfare costs of moderate steady inflation are not negligible. The results differ, depending on the specifics of the money demand aggregates that are used and other details of the modeling strategy, but the estimates are aligned in placing the order of magnitudes of the deadweight losses caused by a moderate inflation between 1 to 3% of annual consumption.[8]

Our paper is not the first one to quantify the misallocation caused by sticky prices. As mentioned, measuring misallocation is complex because it requires the identification of the gap between actual prices and efficient ones, where the latter are not directly observable from the data. The literature has followed different routes to address this problem. Nakamura, Steinsson, Sun, and Villar (2018) and Sheremirov (2020) use US price data and proxy misallocation using observations on the degree of price level dispersion, and the size of price changes. Relatedly, a paper by Adam, Alexandrov, and Weber (2023) assumes that efficient prices follow (product-specific) trend inflation and uses this assumption to identify changes in the inefficient price dispersion in the UK data. A common feature of these papers is to estimate how observed changes in inflation map into an increased cost of misallocation. But the level of the misallocation cost itself cannot be measured. In this paper we use an alternative approach. We use recent results by Baley and Blanco (2021) and Alvarez, Lippi, and Oskolkov (2022) to construct a mapping that allows us to infer the price gaps using observable moments on the size and timing of price changes. This allows us to estimate the level of the cost of misallocation, as well as its evolution following a large inflationary shock. A similar approach is used by Blanco, Boar, Jones, and Midrigan (2022) using the CPI data for the UK.

Finally, other papers focus on the effects of inflation surprises and their distributional effects, namely identifying winners and losers after an inflation surprise, see e.g., Bach and Stephenson (1974); Doepke and Schneider (2006). Other interesting and related analyses of the distributional effects of inflation can be found in Argente and

[8] For instance, Aiyagari et al. (1998) model the costs of inflation as the resources that the households use to protect themselves from the inflation tax. They use a simple model and several empirical datasets to quantify the steady state costs of moderate inflation. An inflation rate of about 10 percent causes welfare losses that are estimated to be between 1% to 2% of total consumption. A similar conclusion is reached by Lucas (2000) who uses a money demand model and quantifies the benefits of reducing inflation from 10 percent to zero at about 1% to be equivalent to an increase in real income of about 1%. A higher value, between 3 to 5% of consumption, is estimated by the paper of Lagos and Wright (2005) based on a search-theoretic model.

Lee (2020), who focus on the dynamics of prices for rich and poor households during the great recession of 2008. While such distributional effects are important, these measurements do not offer a direct assessment of welfare costs, as the deadweight loss of a redistributive policy is not directly measured by those statistics. We focus here on the deadweight losses associated with both anticipated and unanticipated inflation but our analysis is mute about the distributional effects.

2 A generalized setup for NK models

This section presents a New Keynesian setup that describes the firm's price-setting decisions. The motivation for introducing a formal model is that it will allow us to relate the observed price-setting behavior to the fundamental costs and benefits of the price-management activity. We follow a flexible framework proposed by Caballero and Engel (2007) that describes the firm's key decision in terms of the probability of price adjustment. The economics are simple: the more a firm is willing to adjust its price, the more resources must be assigned to that task. If no resources are used, then the prices stay constant. This view is aligned with empirical studies that measure the amount of resources dedicated to price management activities, such as Zbaracki et al. (2004). The firm's behavior is related to its price deviation from profit maximization, denoted by the variable x. The firm's choice variable will be described by a function, $\Lambda(x)$, giving the probability (per unit of time) that the price will be adjusted. Intuitively, it will be shown that larger deviations of x from its ideal value increase the probability that a price change is observed. 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).

Next, we summarize the key model ingredients[9]. We consider a setting where firms are hit by idiosyncratic productivity shocks, so that firm's i profit-maximizing price, P_i^*, is given by a constant markup over marginal costs, mc_i:

$$P_i^*(t) = \frac{\eta}{\eta-1} mc_i(t) \qquad (1)$$

where $\eta > 1$ is the price elasticity of demand, assumed to be constant. Note that $P_i^*(t)$ depends on time because the marginal costs can change over time due to the productivity shocks. Marginal costs are also affected by aggregate shocks, such as a generalized increase in energy prices.

The assumption of sticky prices, the hallmark of New Keynesian economics, creates a wedge between the actual price $P(t)$ and the desired price $P_i^*(t)$. We refer to this gap as the "price gap" and denote it by $x_i(t)$ for firm i at time t. It is given by

$$x_i(t) \equiv \log P_i(t) - \log P_i^*(t) \qquad (2)$$

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

[9] For a detailed illustration of the underlying theoretical setup see Caballero and Engel (1999) and Caballero and Engel (2007), and Alvarez et al. (2022) for a simplified version.

inflation, given by μ, and the idiosyncratic productivity shocks, which are assumed to follow a driftless Brownian motion $\sigma\mathcal{B}(t)$, where σ is the standard deviation of the productivity innovations per unit of time[10].

We describe the firm's price-setting decision as the solution to a minimization problem: the firm chooses its price to minimize the expected present value of the non-zero price gaps, discounted at the rate ρ. The solution to this problem involves balancing two costs: on the one hand, a price gap $x(t)$ implies that the firm's profits are below the maximum level by the amount: $\frac{\eta(\eta-1)}{2}x(t)^2$, where for notation convenience we drop the i index. The quadratic term is obtained from a second-order expansion of the profit function around the profit-maximizing price. The firm would like to "keep x small," i.e., to adjust its price P to track P^*, but since price-setting is costly, this cannot be done in every period. We assume that at each point in time, the firm can choose the probability of price resetting per unit of time, $\ell(t)$, by spending resources $(\kappa\ell(t))^\gamma$ with $\kappa > 0, \gamma > 1$. At each time $t = \tau_j$ where the effort is successful, the price is reset, so that x is reset at the ideal level x^* by a price change of size $\Delta x_{\tau_j} = x^* - x(t)$.[11] This means that the price gap obeys the law of motion $x(t) = x(0) + \int_0^t d\mathcal{B}(s) + \sum_{\tau_j<t}\Delta x_{\tau_t}$.

Formally, the firm solves:

$$v(x) = \mathbb{E}\left[\int_0^\infty e^{-\rho s}\min_{x^*,\ell\geq 0}\left(\frac{\eta(\eta-1)}{2}x(s)^2 + \left(\kappa\ell(s)\right)^\gamma\right)ds \mid x(0) = x\right] \qquad (3)$$

The key element of this problem is the effort rate ℓ for price resetting that each firm chooses at each point in time. As highlighted by Caballero and Engel (1999), this allows us to describe the optimal firm policy through a generalized hazard function (GHF): $\ell^*(t) = \Lambda\big(x(t)\big)$. This function gives the probability that the price will be adjusted given the firm's current price gap $x(t)$.[12]

Chart 4 illustrates the main properties of the firms' optimal price-setting decisions as summarized by the generalized hazard function $\Lambda(x)$. First, the function $\Lambda(x)$ has a minimum at x^*, where it is equal to zero. This is intuitive: when $x = x^*$ the firm is pleased with the current price gap, and there are no incentives to adjust prices. Second, the probability of adjustment is increasing in the distance between and x and the optimal reset gap, $x^* \approx 0$.[13] This is intuitive: a larger value of x increases the benefit of adjusting the price, leading the firm to pay more attention to this task.

---

10 Technically the variable x follows the diffusion $dx = -\mu dt + \sigma d\mathcal{B}$.

11 The optimal return point x^* is the profit-maximizing reset price gap that satisfies $v'(x^*) = 0$. See Appendix A of Alvarez and Lippi (2014). We note that the units of the cost function are expressed as a fraction of forgone (steady state) profits. Given the CES demand system, to express these units in terms of the revenues (and output) they must be divided by η.

12 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 (2011b).

13 The optimal value of x^* depends on the inflation rate. At zero inflation $x^* = 0$ and remains roughly constant at around zero until inflation enters the two-digit region. See Alvarez et al. (2019b) for the theory and an empirical illustration.

Chart 4

The firm decision rule for price changes: $\Lambda(x)$

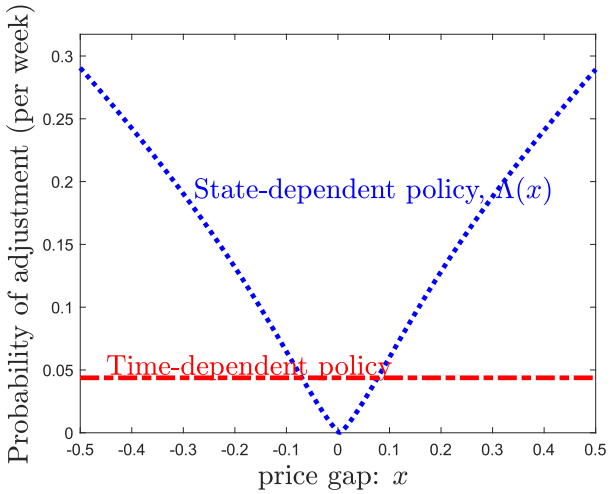

Notes: The model uses the calibration for the food and beverages sector of euro area countries in Table 2. The calibration matches a frequency of N = 2.4 price changes a year, the kurtosis and standard deviation of price changes of 2.4 and 15% respectively.

It is helpful to compare the generalized hazard function $\Lambda(x)$ with the workhorse Calvo (1983) model, where the adjustment probability is assumed to be constant, as depicted by the horizontal line in Chart 4. The key difference is that price-setting decisions in our model depend on the firm's desired adjustment, x. Such state dependence is appealing theoretically, see e.g. Barro (1972); Sheshinski and Weiss (1977); Dixit (1991); Golosov and Lucas (2007), and has been found to be relevant empirically, see e.g. Fougere et al. (2007); Dias et al. (2007); Eichenbaum et al. (2011); Gautier et al. (2022), Gautier and Le Saout (2015), Karadi et al. (2023).[14] We will show below that a key implication of this framework is that state dependence is essential to understand the propagation of large aggregate shocks. Intuitively, after a large shock many firms find themselves with a price that is far away from where it should be. This leads the firms to dedicate more resources to resetting their prices, causing an increase in the frequency of price changes, akin to what was shown in Chart 2. This increased activity will be important to understand the consequences of the large inflationary shock for misallocation, as measured by the dispersion of the price gaps, as well as to measure the amount of resources that are used to "keep prices right," a wasteful activity that is reminiscent of the shoe-leather cost of inflation.

The distribution of the price gaps, described by the density $f(x)$, is important for several questions. In a steady state, $f(x)$ is uniquely determined by the hazard

[14] Several authors have employed the generalized hazard function in applications and empirical work. For recent applications see e.g., Costain and Nakov (2011a); Carvalho and Kryvtsov (2018); Sheremirov (2020); for empirical work see e.g., Berger and Vavra (2018); Petrella et al. (2018), and for related theoretical work Baley and Blanco (2021). Many models are nested by this framework, including the canonical Calvo model with a constant hazard $\Lambda(x) = 1/\kappa$ as $\gamma \uparrow \infty$ the Golosov and Lucas (2007) model with x bounded by the adjustment thresholds where the hazard is flat (almost zero) over a range of x and then spikes up. 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).

function $\Lambda(x)$ and the law of motion of price gaps.[15] The distribution contains information on the inefficiencies at the steady state. For instance, as shown by Gali (2008), the consumer's welfare losses triggered by the presence of the non-zero price gaps are proportional to the variance of the price gaps, $Var(x)$. Intuitively, an economy where the firms have small values of x is preferable to one where the variance of x is large.

Chart 5

The cross-section distribution of price gaps: $f(x)$

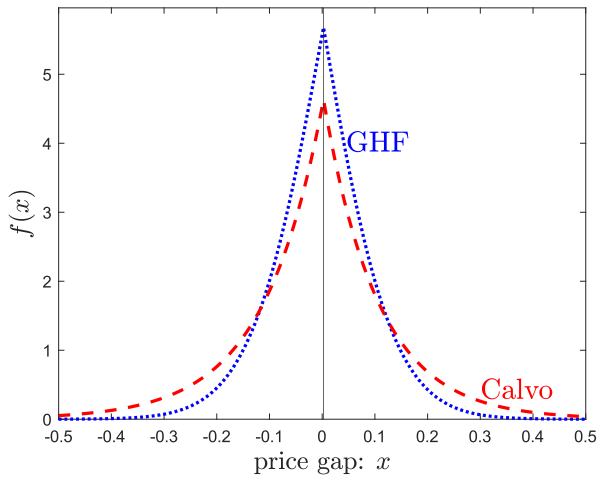

Notes: The model uses the calibration for the food and beverages sector of euro area countries in Table 2. The calibration matches a frequency of N = 2.4 price changes a year, the kurtosis and standard deviation of price changes of 2.4 and 15%, respectively.

Chart 5 illustrates two density functions produced in a steady state with a 2 percent inflation by two models, both featuring a standard deviation of the size of price changes equal to 15%, as in the Euro area data discussed below. The blue dotted density is the one generated by the hazard function $\Lambda(x)$ shown in Chart 4. The red dashed function is the density generated by the corresponding Calvo model. Although both distributions give rise to price-setting behaviors that look alike in the steady state (similar frequency and size of price changes), there are important differences. First, it is apparent that the Calvo model has "fatter tails." This observation, confirmed by a rigorous analysis of the model, implies that the welfare costs of misallocation are larger in the Calvo model compared to a state-dependent model fitting the same price-setting behavior. Second, we will show that despite the steady state similarities, these models imply very different behavior in response to a large aggregate shock.

These considerations suggest that it is of interest for several policy questions to estimate $f(x)$ as precisely as possible. Unfortunately, since price gaps are unobservable, the density f cannot be directly measured in the data. To address this challenge, we calibrate the model and identify f using the observed distribution of the *sizes of price changes*.

[15] Formally, the density $f(x)$ solves the Kolmogorov forward equation $\Lambda(x) \cdot f(x) = \mu f'(x) + \frac{\sigma^2}{2} f''(x)$, for each $x \neq x^*$, with boundary conditions $\lim_{x \downarrow x^*} f(x) = \lim_{x \uparrow x^*} f(x)$; $1 = \int_{-\infty}^{\infty} f(x)dx$, and $\lim_{x \to \infty} f(x) = \lim_{x \to -\infty} 0$.

3 Price-setting behavior: data vs. theory

This section presents statistics of price-setting behavior that allows calibrating the model's steady state to match the empirical evidence. We first use data from the years before 2022, before the large energy shocks hit Europe, to infer behavior in a steady state. We present a few critical predictions of the model about the frequency and the size distribution of the price changes. We argue that the GHF model is able to account for several key patterns observed in the data.

In Section 3.1 we use the model to study the response of prices to a large energy shock. This exercise serves two purposes. First, it provides a validation of the model by comparison of the predictions for e.g. the frequency of price changes with the actual data for 2022 and 2023. Second, it allows us to quantify the welfare costs following the large shocks, an issue that we will inspect in Section 4.2.

Overall, this section shows that a data-consistent model of price-setting implies that the economy's response to a large shock differs markedly from the response to a small shock. In particular, a large shock will give rise to a much faster passthrough from costs to prices, resulting in a period of temporarily-high inflation.

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 information required to solve the inverse inference problem of unobservable price gaps mentioned above. Our dataset was provided by PriceStats, a private company related to 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 web pages daily to gather and store relevant data. The dataset includes product details like price, category, and sale indicators. The data is uncensored, covering the entire lifespan of all products sold by these retailers, and provides prices similar to those obtained in offline stores (Cavallo 2017). We use a subset of data from retailers in 9 European countries: France, Germany, Greece, Ireland, Italy, the Netherlands, Spain, Poland, and the UK. The period ranges from January 1st, 2019 to May 1st, 2023. We focus on the "Food and Beverages" category, which has experienced one of the highest rates of inflation during this period in many countries.

This dataset offers several advantages over traditional data sources such as Consumer Price Index (CPI) and Scanner Data. Firstly, it provides daily price updates, free from unit values, time-averaging, and imputations, which are common issues in CPI and Scanner Data. This high-frequency data collection allows for a more accurate identification of sales and price changes (Cavallo, 2018). Another major advantage of this dataset is the uncensored price spells. Unlike with other data, prices here are continuously recorded from the day they are first offered to consumers until they are discontinued, offering a complete and unaltered view of the product's price life cycle. Furthermore, the data is comparable across countries, collected using identical techniques for similar categories of goods over the same time period. Finally, it offers real-time availability, providing up-to-date information

without processing delays. This makes it a potentially valuable tool for central banks and policymakers in real-time estimation of price stickiness and related statistics.

In Table 1, we present summary statistics of price-setting behavior for several Euro area countries using data provided by PriceStats. We identify regular prices using the sales indicator provided by PriceStats. Additionally, we present aggregate statistics from related studies to complement our analysis. In particular, we use the statistic for the frequency of price changes from Gautier et al. (2022), who use data underlying the CPI. They report several statistics using data from the large-scale Price-Setting Microdata Analysis (PRISMA) network led by the ECB. The statistics of mean, standard deviation, and kurtosis of price changes were kindly provided by the authors upon request. We also use statistics from supermarket scanner data in 4 Euro Area countries from Karadi et al. (2023). The source of those data is IRI. To compute the standard deviation of the size of price changes for the supermarket scanner data, we use their reported measure for the mean absolute deviation (MAD) and the assumption that the distribution is close to normal so that $STD(\Delta x) \approx \sqrt{2} \cdot MAD(\Delta x)$.

Our choice of reported statistics is guided by theoretical insights. The frequency of price changes gives information about the cost of price management. Moreover, when analyzed with the standard deviation of price changes, they allow us to infer the size of idiosyncratic productivity shocks affecting firms. In turn, kurtosis has been shown to reveal important information about the response of an economy to aggregate shocks. In Alvarez et al. (2016), the authors remark that kurtosis encodes information about the “selection” of price changes: 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 relatively large mass of late price adjusters, which implies a more persistent effect of a cost shock. The Calvo model, with a kurtosis of 6, does not feature “selection” of price changes since the adjusters are a random sample each period.

Table 1 shows that the standard deviation of the size of price changes is similar across countries except for Italy and Ireland. The kurtosis measure is also similar across Euro area countries and ranges between 2.1 and 2.7, apart from Ireland, with a kurtosis of 1.6. The frequency of price changes in the food and beverages sector is larger than the aggregate data and differs across countries. For instance, the UK displays a frequency of 0.7 price changes per year, whereas the French one of 3.6 price changes per year.

Table 1

Price-setting Behaviour before 2022

Euro area CPI data (PRISMA data, period 2005-19, Gautier et al. 2022)					
	Mean (Δx)	**STD (Δx)**	**Kurtosis (Δx)**	**Frequency**	**$\hat{\mu}$**
	0.019	0.10	4.1	1	0.019
Euro area supermarket data (IRi data, period 2013-17, Karadi et al. 2023)					
	Mean (Δx)	**STD (Δx)**	**Kurtosis (Δx)**	**Frequency**	**$\hat{\mu}$**
	-	0.13	3.2	1.1	-
Euro area Food and Beverages data (PriceStats data, period 2019-21)					
	Mean (Δx)	**STD (Δx)**	**Kurtosis (Δx)**	**Frequency**	**$\hat{\mu}$**
France	-0.001	0.13	2.3	3.6	-0.002
Germany	0.008	0.13	2.5	1.4	0.017
Greece	0.003	0.10	2.7	2.0	0.005
Ireland	0.001	0.32	1.6	1.6	0.001
Italy	-0.002	0.23	2.1	2.1	-0.003
Netherlands	0.009	0.11	2.3	2.8	0.022
Spain	0.008	0.16	2.5	2.9	0.016
EA Average	0.004	0.15	2.4	2.4	0.007
Other countries Food and Beverage data (PriceStats, period 2019-21)					
	Mean (Δx)	**STD (Δx)**	**Kurtosis (Δx)**	**Frequency**	**$\hat{\mu}$**
Poland	0.013	0.21	2.2	2.8	0.023
UK	0.002	0.29	1.9	0.7	0.002

Notes: The PriceStats data uses a sample of changes in regular prices (excluding sales). The statistics are computed 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 a correction for unobserved heterogeneity proposed by Alvarez et al. (2022). The statistic for the frequency of price changes from the Price-setting Microdata Analysis (PRISMA) network is obtained from Table 7 in Gautier et al. (2022). The other statistics were kindly provided by the authors. We report the statistics after dropping outliers corresponding to the bottom and top 2.5% of the distribution of price changes. These data cover the period from 2005 to 2019. The statistics from Karadi et al. (2023) are taken from their Table 2 and correspond to the average of 4 Euro Area countries: Germany, France, Italy, and the Netherlands between 2013 and 2017.

Calibration. We calibrate the model to match the standard deviation, the kurtosis, and the frequency of price changes. We use the identity $\sigma^2 = N \cdot Var(\Delta x)$ since this relationship holds for a wide variety of models when $\mu \approx 0$, see Alvarez et al. (2022). We use standard values for the additional parameters of elasticity of substitution and intertemporal preference: $\eta = 6$ (which implies a markup of 20%) and a time discount $\rho = 0.05$. We choose an inflation rate of $\mu = 2\%$ consistent with inflation at steady state. The selection of a kurtosis value of 2.8 for the PRISMA data stems from the acknowledgment that the value of 4.1 does not account for unobserved heterogeneity. In a comparable investigation employing French CPI data, Alvarez et al. (2021) control for heterogeneity through an appropriate filter. As a result, their analysis yields a reduced kurtosis estimate that is 32% lower, providing a basis for deriving the value of 2.8 mentioned above.

Chart 6

The Size of Price Adjustments, $q(\Delta x)$, in France – Food and Beverages sector

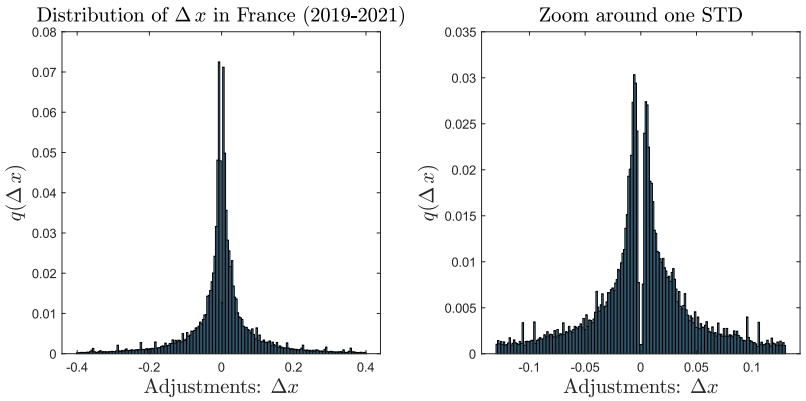

Notes: The right panel shows the distribution for regular price changes with absolute value less of than one standard deviation which in this case is 13%.

Next, we use a GMM estimator to calibrate the parameters of the effort cost function κ, γ to match the standard deviation and kurtosis of price changes. With the described estimation method, the model is exactly identified, and there is a one-to-one mapping between the mentioned moments and the parameters of the cost function. The calibrated parameters are shown in Table 2. Recall that the kurtosis of a Calvo model is equal to 6, while the kurtosis of a canonical menu cost model is 1. The data suggest a somewhat intermediate situation. Next, we present a few key predictions of the model about the frequency and the size distribution of the price changes.

Table 2

Calibration for Price-Setting Behaviour in the Euro Area before 2022

Parameters	CPI data (PRISMA data) Gautier et al. 2022	Supermarket data (IRi data) Karadi et al. 2023	Food and Beverages data (PriceStats data)
$\sigma = \sqrt{N} \cdot STD(\Delta x)$	0.10	0.14	0.23
γ	3.30	4.19	2.61
κ	0.21	0.28	0.09
Matched moments			
STD(Δx)	0.10	0.13	0.15
Kurt(Δx)	$2.8^{(a)}$	3.2	2.4
Frequency (N)	1	1.1	2.4

Notes: (a): The matched value of kurtosis for the PRISMA data is corrected for heterogeneity using a multiple from Alvarez et al. (2021) who perform the correction for French CPI data. Parameters γ, κ are calibrated using a GMM estimator to match the standard deviation and kurtosis of price changes in data (Table 1). The drift of price gaps is μ = 2%. The additional parameters are set to standard values: η = 6, ρ = 0.05.

Frequency of price changes. The cross-sectional distribution of firms' price gaps $f(x)$ and the generalized hazard function $\Lambda(x)$ can be used to compute several objects that are observable in data. The steady-state frequency of price adjustments N is given by

$$N = \int_{-\infty}^{\infty} f(x)\Lambda(x)dx \qquad (4)$$

The equation has a simple interpretation: it counts the total number of firms with a given price gap, f , and their probability of adjustment in a time period (say a year), Λ. These price adjustments originate from the firm's effort ℓ to control the price gaps.

Distribution of the size of price changes. Recall that upon any price change, the firm resets its gap from x to the optimally chosen x^*, i.e. the size of the adjustment is $\Delta x = x^* - x$. This occurs with probability $\Lambda(x)$ per unit of time. Then the distribution of the size of price changes has the following density $q(\Delta x)$:

$$q(\Delta x) \equiv \frac{\Lambda(x) f(x)}{N} \qquad (5)$$

The left panel of Chart 7 shows that the calibrated model can capture some key features of the data in Chart 6: the distribution of price changes is bimodal, with a dip at zero. The latter is a major difference compared to the prediction of the Calvo model, where the constant hazard implies a mode at zero, i.e., that the most frequently observed price change has a tiny size. This prediction is counterfactual and it is a telltale of the fact that price-setting behavior displays state dependence: prices are adjusted only when necessary.

Chart 7

Distribution of the size of price adjustments, $q(\Delta x)$, in two models

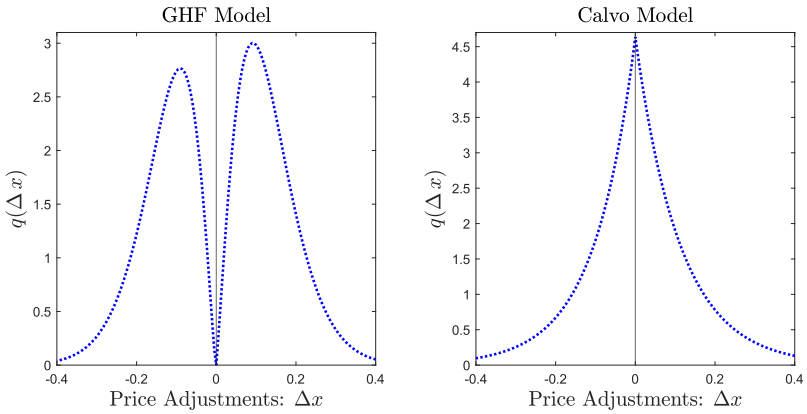

Notes: The model uses the calibration for the food and beverages sector of euro area countries in Table 2. The calibration matches a frequency of N = 2.4 price changes a year, the kurtosis and standard deviation of price changes of 2.4 and 15%, respectively.

3.1 The propagation of an aggregate cost shock

In Chart 2, we reported that inflation and the frequency of price adjustments rose quickly after a large energy shock. In this subsection, we provide a thought experiment that rationalizes these facts. Namely, it takes an economy at steady state and hits it with a marginal cost shock, as will be made precise below. We take calibrations of the price-setting model presented in Table 2 to study the propagation of large and small shocks. We will illustrate that, under a GHF model, large and small shocks have different implications for passthrough and the frequency of price adjustments.

In this study, we will focus on the implications of a large shock characterized by a 20% rise in marginal costs. The rationale behind this specific figure stems from the observed escalation in energy prices, which surged by 200% between January 2021 and April 2022, as depicted in Chart 1. By considering the proportion of energy expenses, estimated to constitute 10% of total input costs, we determine that this value of 20% accurately captures the magnitude of the shock under investigation.

Take an economy characterized by a steady state cross-sectional distribution of price gaps f and a policy rule Λ. The economy is then hit by an unexpected once-and-for-all shock to marginal cost that displaces the distribution of price gaps δ percentage points to the left, as in Chart 8. Firms would like to increase their prices to close their gaps. This incentive shapes the dynamic response of the price level and the frequency of price changes after the shock. We will describe the transition of these variables back to steady state for small and large shocks.

Chart 8

Displacement of price gaps: Large vs Small Cost shock

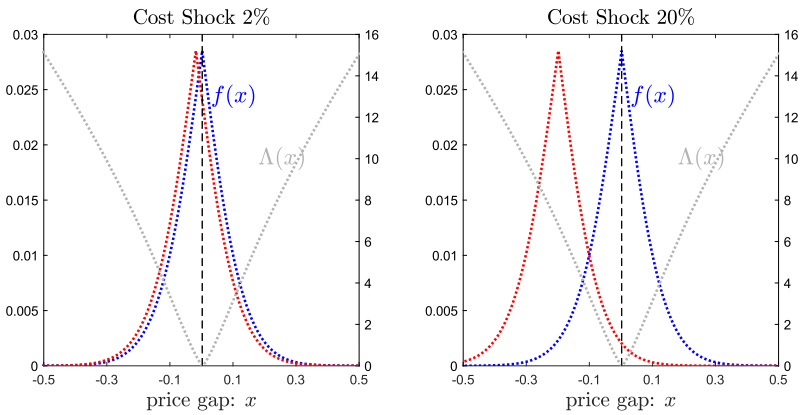

Notes: The model uses the calibration for the food and beverages sector of euro area countries in Table 2. The calibration matches a frequency of N = 2.4 price changes a year, the kurtosis and standard deviation of price changes of 2.4 and 15% respectively. The GHF is plotted relative to the right vertical axis.

To understand the mechanism behind the propagation of the shock, notice that in Chart 8, the distribution after the shock places most firms' prices in a region far from their desired prices (at a price gap of around −20%). In this region, the probability of adjustment $\Lambda(x)$, plotted in dotted black, is higher. This means that a large shock triggers an increase in the number of price adjustments, whereas for a small shock this effect is much weaker.

Chart 9
Size-dependent propagation of shocks

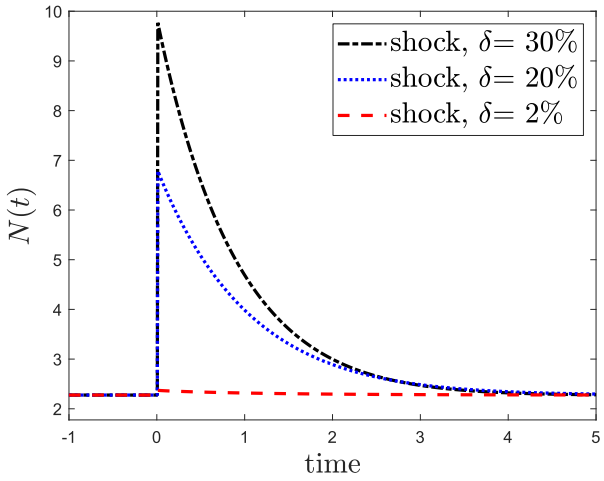

Notes: The model uses the calibration for the food and beverages sector of euro area countries in Table 2. The calibration matches a frequency of N = 2.4 price changes a year, the kurtosis and standard deviation of price changes of 2.4 and 15% respectively.

Chart 9 plots the response of the frequency and the price level after shocks of different sizes. The model predicts that the frequency of price changes increases sharply after a large shock. 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. Notice that this effect cannot be generated by the constant hazard of the Calvo model.

Chart 10 Dynamics of inflation

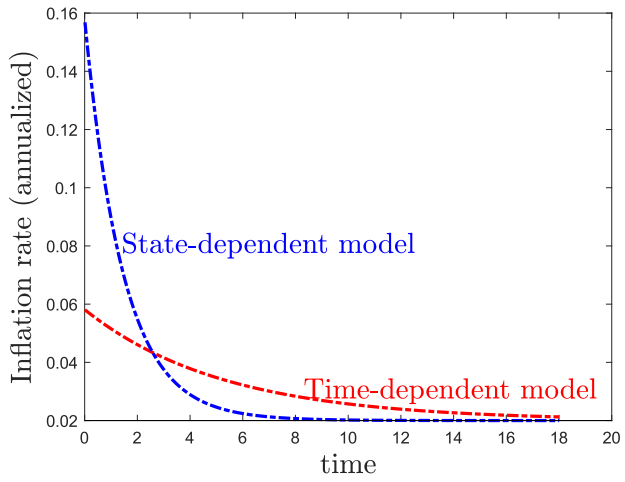

Notes: The model displays the dynamics of inflation of a cost shock with δ = 20%. The state-dependent model uses the calibration for the food and beverages sector of euro area countries in Table 2. The calibration matches a frequency of N = 2.4 price changes a year, the kurtosis and standard deviation of price changes of 2.4 and 15% respectively. The time-dependent model is a Calvo model with N = 2.4.

We now study how a large cost shock propagates to the price level in an economy characterized by the Euro Area food and beverages sector data. The left panel of Chart 10 presents the dynamics of inflation after a 20% shock to marginal costs in

such an economy. Namely, for the GHF model calibrated with the PriceStats data in Table 2. The figure further depicts the inflation dynamics implied by a Calvo model with an identical steady-state frequency of price adjustments (dotted red line). The figure illustrates the economy's swifter response to the large shock in comparison to the corresponding reaction predicted by the Calvo model. Put another way, the figure shows the forecast error associated with the use of a Calvo model: failing to account for the large increase in the frequency of price revisions, the Calvo model leads to a substantial initial underestimation of inflation followed by a subsequent overestimation.

Chart 11

The distribution of price changes after a large shock

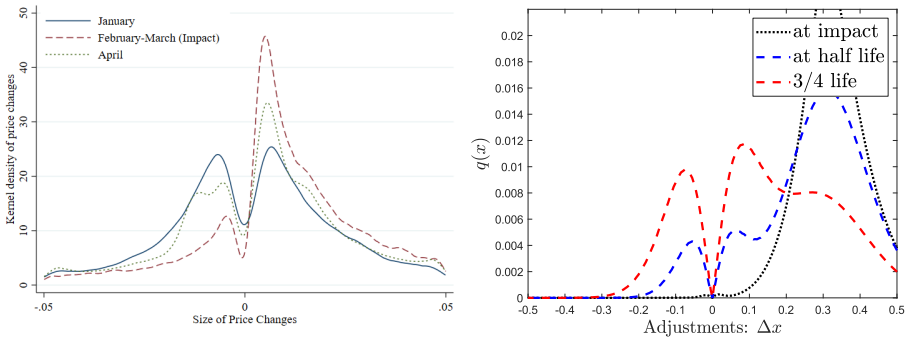

Notes: The model uses the calibration for the food and beverages sector of euro area countries in Table 2. The calibration matches a frequency of N = 2.4 price changes a year, the kurtosis and standard deviation of price changes of 2.4 and 15% respectively.

Chart 11 presents an additional validation of the model, showcasing the dynamics of the price change distribution in response to a large shock. The right panel illustrates the distribution of price changes following the shock, while the left panel displays the distribution of price changes for France before and after the energy price shock. Notably, the model successfully reproduces the qualitative characteristics observed in the French data, namely an asymmetric distribution with a greater mass of positive price adjustments.

In summary, the state-dependent model, supported by the granular empirical evidence, features a dynamic response to a large cost shock that resembles the patterns for inflation and the frequency of price adjustments following the recent energy cost shocks in Europe. Furthermore, the forecast of the frequency of price changes and the path of inflation can be markedly different depending on the model the analyst is using. When shocks are large, the Calvo model, or more generally any time-dependent model, will fail to capture the change in the frequency of price adjustments and hence yield a biased forecast of inflation, as illustrated in Chart 10.

4 Quantifying the welfare cost of inflation

In this section, we analyze the welfare costs in NK models in two steps. First, we measure both welfare costs, χ and φ, in a steady state. We assess these costs using data from the low inflation period that prevailed before 2022. Second, we study the

dynamics of these welfare costs following a large shock to the firms' marginal costs. Specifically, we explore the change in welfare costs that follow the large energy price increases in March 2022.

We focus on two inefficiencies that arise in the new Keynesian framework due to the sticky-price friction. The first is due to price gap dispersion: sticky prices introduce a wedge between the marginal rate of substitution and the marginal rate of transformation for consumers and workers. These wedges give rise to an inefficient allocation of resources. We call this the welfare cost of "misallocation" and denote it by χ. The second one arises when firms waste resources to keep prices close to their optimal levels. This corresponds to a deadweight loss, akin to the shoe-leather cost of inflation in textbook models of money demand. We call this the welfare cost of "price management" and denote it by φ. Since the welfare costs depend on the distribution of price gaps, which is unobservable, we derive a mapping between the welfare costs (χ and φ) and a set of observable moments from the distribution of the size of price changes and the frequency of price changes.

We use three different data sets to inform our analysis. First, we use granular data for the food and beverages sector of several European countries taken from PriceStats (Cavallo, 2018). Although these data cover only a fraction of the CPI, the high quality of the measurement, based on daily observations, is important to identify key features of price-setting behavior. Second, we use some descriptive statistics for the Euro Area taken from the Price-setting Microdata Analysis (PRISMA) network (Gautier et al., 2022) that uses data underlying the CPI in several Euro area countries. Thirdly, we use some descriptive statistics from Karadi et al. (2023), who use supermarket data for some Euro Area countries.

4.1 Measuring the welfare costs in NK models in steady state

Misallocation. As explained above, the first component of the welfare costs is due to wedges introduced by price gaps. Namely, it is related to price gap dispersion. These welfare costs also scale with the parameter of elasticity of demand. Intuitively, this is due to a larger resource-shifting reaction caused by a more sensitive demand. It can be shown that the welfare cost of misallocation denoted by χ is given by

$$\chi(t) \equiv \frac{U-\bar{U}}{U_c \cdot \bar{C}} \approx \frac{\eta}{2} Var_t(x), \qquad (6)$$

where the index t emphasizes that this expression can be used in steady state as well as along a transition after a shock. The cost χ is expressed in terms of GDP percentage loss relative to the efficient GDP level. Notice that the expression above, up to first order, is the one used in traditional monetary policy analyses presented by Gali (2008) (pp. 63) and Woodford (2003) (pp. 396).[16]

The expression in equation (6) gives us a direct map from the distribution of price gaps to welfare losses due to misallocation. However, price gaps are unobservable, see equation (2), since they require observing firms' productivity. To address this

[16] Recall that we analyse a steady state problem so second order effects (risk) are not present.

issue, we derive a mapping from data moments to the variance of price gaps. It can be shown that in a low inflation steady state, the variance of price gaps satisfies,

$$Var(x) = \frac{Var(\Delta x) \cdot Kurt(\Delta x)}{6}.$$

This expression, together with equation (6) and an estimate for the elasticity of demand, allows the researcher to gauge the level of misallocation costs due to sticky prices in steady state.

We are not the first ones to attempt measuring price gap dispersion. However, we emphasize the limited informativeness of measuring price level dispersion, even at the product level (Sheremirov, 2020), with regard to assessing misallocation. It is crucial to recognize that variations in prices can be attributed to differences in productivity, which do not necessarily indicate inefficiency.[17] Instead, examining the standard deviation of price changes and its kurtosis proves to be an appropriate approach in this context.

Firstly, emphasizing the magnitude of price changes yields valuable insights into inefficiency wedges, given that such changes reveal the pre-adjustment price gap x, a point made by Nakamura et al. (2018). However, relying solely on this statistic may be insufficient, as various models can exhibit identical standard deviations of price changes while displaying different levels of misallocation.

The kurtosis of price changes, which encodes the "selection" effect, is the second key statistic to measure misallocation. A large kurtosis is an indicator of a relatively large mass of late price adjusters which implies larger misallocation (Alvarez et al., 2016). For example, fixing the standard deviation of price changes, the Calvo model, with a kurtosis of 6, does not feature "selection" of price changes since the adjusters are a random sample each period. In contrast, a canonical menu cost model has a kurtosis of 1. The evidence in Table 1 seems to suggest an intermediate point with kurtosis hovering around 2 and 3.5.

In a related study, Blanco et al. (2022) have estimated the extent of misallocation and menu costs leveraging data underlying the CPI in the UK. Their estimations share a similar order of magnitude when compared to our favored calibration for the Euro Area, with a misallocation cost accounting for 2 percentage points of GDP and menu costs amounting to 2.4 percentage points of firms' revenues for the UK.[18]

Nakamura et al. (2018) investigate the effect of inflation on misallocation. They argue that misallocation comoves with the size of price changes. They aim to shed light on the effects of inflation using observations on the absolute size of price changes. Intuitively, an increase in the size of the price changes suggests an increase in misallocation. While such an approach provides information on the change of the welfare costs, it is silent about their level, as we argued above.

---

17 Sheremirov (2020) acknowledges this limitation and undertakes a significant endeavour to control for many determinants of productivity across retailers to shed light on inefficient price dispersion.

18 Their model includes some features absent in ours e.g., strategic complementarity in pricing decisions.

Price management. The second source of inefficiency is due to the forgone resources used to perform price-management activities by firms.

Zbaracki et al. (2004) present an interpretation based on empirical findings from managerial reports in the United States. The authors posit that pricing activities require managers to spend resources on processes such as information acquisition, decision-making, and communication costs. Additionally, they claim that the magnitude of these allocated resources increases in a convex manner with the absolute size of the price change. They estimate the costs of price management to be around 1 percentage point of GDP.

We can observe that the price-management technology of the model presented in Section 2 shares the qualitative features described in Zbaracki et al. (2004). Recall that firms face a convex cost $(\kappa\ell)^{\gamma}$ corresponding to the effort ℓ of managing prices. Since the optimal effort rate $\ell^{*} = \Lambda(x)$ increases with the absolute size of the price change $|x^{*} - x|$ then the price-management cost is convex in the absolute size of adjustment.

The expression giving the cost of price management in the GHF model is

$$\phi(t) \equiv \frac{1}{\eta} \cdot \mathbb{E}_t\left[\left(\kappa\Lambda(x)\right)^{\gamma}\right] \qquad (7)$$

This equation has the following interpretation: it counts the total effort cost (resources) used to affect the probability of price adjustment. Since these costs are expressed as a percentage of forgone firm's profits, the expression is then divided by η, the share of profits in GDP, to obtain a measure as a percentage of GDP.

Application to data. We next analyze the data through the lenses of the results just established. Table 3 reports the estimated steady state welfare costs, measured over the low inflation period before 2022, using the datasets described before.

Table 3

Estimates of Welfare Costs for the Steady State before 2022

Observables				Estimates		
Euro area CPI data (PRISMA data, period 2005-19, Gautier et al. 2022)						
	Frequency	**STD (Δx)**	**Kurtosis (Δx)**	$\widehat{\mathrm{Var}}(x)$	$\hat{\chi}$	$\hat{\Phi}$
	1	0.10	2.8[(a)]	0.005	0.015	0.004
Euro area supermarket data (IRi data, period 2013-17, Karadi et al. 2023)						
	Frequency	**STD (Δx)**	**Kurtosis (Δx)**	$\widehat{\mathrm{Var}}(x)$	$\hat{\chi}$	$\hat{\Phi}$
	1.1	0.13	3.2	0.009	0.027	0.005
Euro area Food and Beverages data (PriceStats data, period 2019-21)						
	Frequency	**STD (Δx)**	**Kurtosis (Δx)**	$\widehat{\mathrm{Var}}(x)$	$\hat{\chi}$	$\hat{\Phi}$
France	3.6	0.13	2.3	0.006	0.019	0.007
Germany	1.4	0.13	2.5	0.007	0.020	0.007
Greece	2.0	0.10	2.7	0.005	0.014	0.003
Ireland	1.6	0.32	1.6	0.027	0.080	0.042
Italy	2.1	0.23	2.1	0.019	0.057	0.021
Netherlands	2.8	0.11	2.3	0.004	0.013	0.005
Spain	2.9	0.16	2.5	0.010	0.021	0.009
EA Average	2.4	0.15	2.4	0.009	0.027	0.008
Other countries Food and Beverages data (PriceStats data, period 2019-21)						
	Frequency	**STD (Δx)**	**Kurtosis (Δx)**	$\widehat{\mathrm{Var}}(x)$	$\hat{\chi}$	$\hat{\Phi}$
Poland	2.8	0.21	2.2	0.015	0.046	0.017
UK	0.7	0.29	1.9	0.027	0.080	0.033

Notes: Data sources for observables are discussed in Table 1. (a): We use a value of kurtosis of 2.8 for the PRISMA data in an attempt to correct for heterogeneity, see the discussion in Table 2. The estimates for the welfare costs are obtained from the model calibrations in Table 2 and assume an price elasticity of demand equal to 6 which implies a markup equal to 20%. The calibration matches the frequency of price changes per year, the kurtosis and standard deviation of price changes.

For clarity, we used the granular European countries' data on food and beverages to illustrate the application of our main findings. But the results can be applied to measure the welfare costs in various sectors and countries where micro-data are available. The requirements for this assessment are the frequency, standard deviation, and kurtosis of the distribution of the sizes of price changes and an estimate for the elasticity of demand. We attempt a preliminary extension of our exercise using the pricing statistics collected by the Price-setting Microdata Analysis (PRISMA) in Gautier et al. (2022) and supermarket data for four Euro area countries taken from Karadi et al. (2023). As described in Table 1, we take the statistic of frequency of price changes from Table 7 of Gautier et al. (2022), whereas the other observables were kindly provided by the authors upon request. The choice of 2.8 as a value for the kurtosis of price changes is discussed in Table 2. The model calibrated to the Euro Area using the PRISMA data (Gautier et al., 2022) suggests that the costs of misallocation are in the order of 1.5% of GDP (fifth column of the table). The model also allows us to gauge the welfare costs due to the price management, which turn out to be smaller, in the order of 40 basis points of GDP (last column of the table). This magnitude is close to the direct estimates of 1% in the literature (Zbaracki et al., 2004).

The findings derived from the analysis presented in Table 3 are contingent upon an accurate estimation of three key values: the price elasticity of demand, the standard deviation of price changes, and the kurtosis of price changes. Based on our observations, the estimated values for kurtosis generally lie within the range of 2 to 3.5, while standard deviation values typically range from 10% to 20%. Consequently, while it is prudent to approach point estimates with caution, we can reasonably assert that the estimated magnitudes fall within a narrow range depending on the particular price-setting behavior in the economy.

Next, we will briefly discuss the effects of a change in the steady-state inflation rate and the associated welfare costs. However, we argue that this approach is not suitable for studying the welfare implications of the recent inflation surge. Consequently, we will proceed under the assumption that the steady-state inflation rate remains constant. This amounts to assuming that the shock only has a temporary effect on inflation but that it does not alter the long run inflation expectations. Our focus will be on studying the impact of a large cost shock, which we interpret as the recent energy experience in Europe during 2022

Steady state inflation. One viable approach to analyze the costs associated with inflation is to consider alternative steady states with different inflation rates. Interestingly, most models considered before display little variation in the level of welfare costs as long as inflation remains below two digits. Chart 12 illustrates this result: the steady-state welfare costs as a function of inflation have a flat profile around zero inflation. This result occurs because, when inflation is low, the costs of misallocation in equation (6) are closely related to the size of price changes, which aligns with the firms' desire to minimize such gaps. Stated differently, the firms' optimal response to steady-state inflation mitigates the costs of misallocation.

In particular, Alvarez and Lippi (2022) show that all even moments of the distribution of price changes feature a low sensitivity to steady-state inflation. Chart 12 illustrates that both the cost of misallocation and price-management activities increase by 50 basis points of GDP as steady-state inflation rises from 0% to 20%. As a matter of fact, the Calvo model allows us to obtain a closed-form solution of the price gap dispersion that arises under steady-state inflation. The cost χ in the Calvo model is proportional to Var(x) following equation (6), where the latter is given by

$$(Var\ x)(\mu) = \left(\frac{\mu}{N}\right)^2 + \frac{\sigma^2}{N}$$

We observe that the variance of price gaps is flat at $\mu = 0$, meaning that for small changes in steady-state inflation, the costs of misallocation remain essentially unchanged.

The observations above need not be true for high inflation environments in steady state, as in e.g., Argentina or Turkey. Nevertheless, recent results by Baley and Blanco (2021) allow us to obtain an inverse mapping between observables and the variance of price gaps under high inflation.

Chart 12

Steady-state inflation vs welfare cost (in % of GDP)

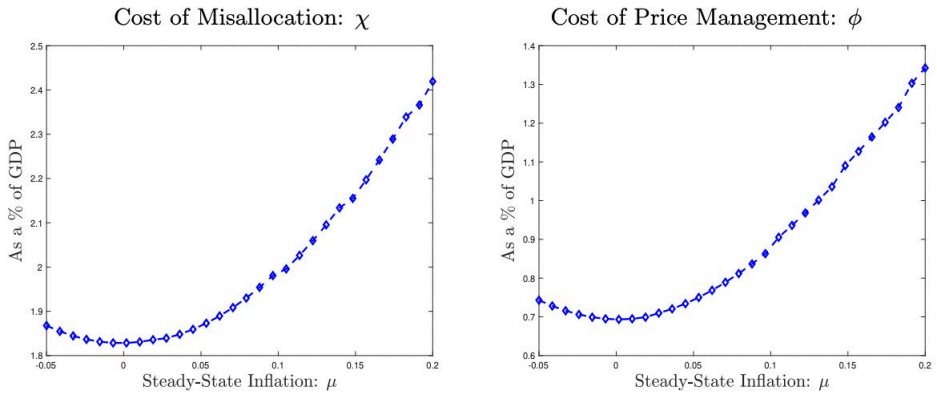

Notes: The model uses the calibration for the food and beverages sector of euro area countries in Table 2. The calibration matches a frequency of N = 2.4 price changes a year, the kurtosis and standard deviation of price changes of 2.4 and 15% respectively.

4.2 The welfare cost of a large cost shock

This section analyzes the welfare costs using a calibrated economy's reaction to a large cost shock. We argue that a large portion of the recent inflation surge can be attributed to the large shock in energy prices that hit Europe in 2022. Additionally, we posit that this shock has not substantively affected expectations regarding the steady state level of inflation, which we maintain to be approximately 2% across our exercises.

As documented above, a large shock triggers a sizeable increase in price-setting activity. A proper assessment of the welfare costs must account for the fact that since firms are more active (and more price changes are observed), both the degree of misallocation χ and the resources for price-setting activity φ are likely to change. The cumulative welfare costs of misallocation and price-management activities are given by

$$CIR[\chi(t)] = \int_0^t (\chi(s) - s)ds, \quad , \quad CIR[\phi(t)] = \int_0^t (\phi(s) - \phi)ds, \qquad (8)$$

where $\chi(t)$ and $\phi(t)$, are defined in equation (6) and equation (7), respectively, and their steady-state values χ and ϕ are computed using the steady-state cross-sectional distribution of price gaps f . The cumulative response $CIR[\chi(t)]$ in equation (8), measures the welfare costs of misallocation above and beyond the steady state cost. Likewise, the cumulative response $CIR[\phi(t)]$ gives the total excess welfare cost of price management after time t has elapsed from the large shock.

The left panel Chart 13 plots the cost of misallocation for the three shock sizes studied using the calibration to the Euro Area with data from the PRISMA network (Gautier et al., 2022). The cumulative costs of misallocation after a 20% shock are around 1.5 percentage points of GDP, as reported in Table 4. At the time of the shock, the variance of price gaps is the same as in steady state. However, the mass of firms that are displaced very far from their desired price are likely to adjust to x^*

immediately after the shock, and this increases the variance of price gaps. More precisely, we can see in Chart 8 that x^* is far from the mean price gap so the distribution displays more variance than at steady state.

The cost of price management has no counterpart in our data, but we can use the model to infer these aggregate losses. The cumulative costs of price management after a 20% shock are around 1.2 percentage points of GDP, reported in Table 4. As mentioned before, the magnitude of the shock renders different levels and dynamics for the responses. For a large shock, the price-management costs are initially higher. First, this is because on average firms' price gaps x are further away from x^* Which renders a high incentive to engage in costly price-management efforts. Second, the cost of price-management increases in a convex manner with the absolute size of the change.

Table 4 reports estimates of the welfare costs of an inflationary 20% shock using the model. The estimated welfare costs for our preferred Euro Area calibration are 2.7 percentage points of GDP. Moreover, the estimated costs implied by the supermarket data by Karadi et al. (2023) are 3 percentage points of GDP, whereas the costs implied by the food and beverages sector are 90 basis points of GDP. The stark difference between these welfare costs is attributed mainly to differences in the frequency and the kurtosis of price changes. Namely, the food and beverages data have a higher degree of price flexibility showcased by more frequent adjustments and a relatively lower mass of late adjusters.

Table 4

Estimates of welfare costs for the Euro Area triggered by a 20% cost shock

Model calibration δ = 20%	CIR$[\hat{\chi}]$	CIR$[\hat{\phi}]$
CPI data (PRISMA data, Gautier et al. 2022)	0.015	0.012
Supermarket data (Karadi et al. 2023)	0.019	0.011
Food and Beverages data	0.004	0.005

Notes: The estimates for the cumulative welfare costs are obtained from model calibrations described in Table 2.

Although unreported, we measured a lower standard deviation of price changes for 2022-2023. One might be tempted to think that this standard deviation of price changes translates into lower misallocation, but this intuition is misleading. In fact, a lower standard deviation of price changes after a large shock is consistent with the dynamics predicted by the model and is associated with a higher cost of misallocation.

In the right panel of Chart 13 we also show the evolution of the mean absolute size of price changes after a shock. This statistic has been used to gauge changes in the costs of misallocation (Nakamura et al., 2018). Indeed, the dynamics of this variable do behave qualitatively as the time profile of the costs of misallocation in this model.

Chart 13

Impulse responses of moments of the distribution of price changes

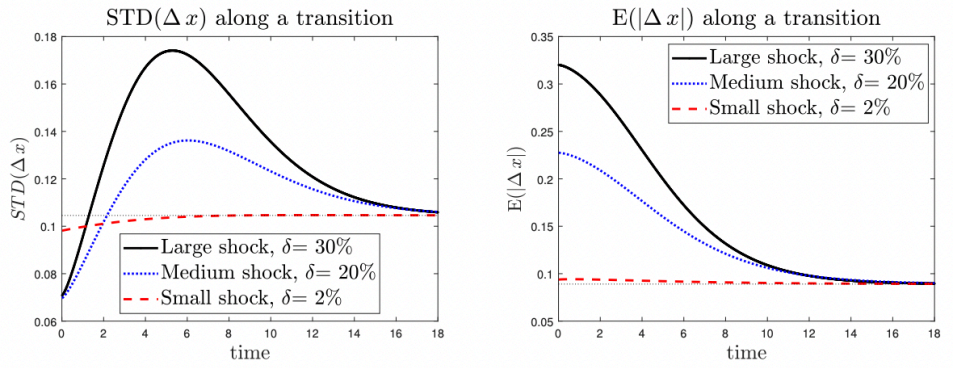

Notes: The model uses the calibration for the Euro area with PRISMA data which covers many sectors in several euro area countries (Gautier et al., 2022). For the calibration see Table 2. The calibration matches a frequency of N = 1 price change a year, the kurtosis and standard deviation of price changes of 2.8 and 10% respectively.

Economies with different frequencies of price changes. We presented a welfare analysis for a model calibrated to a frequency of 1 price change per year for the Euro Area. However, it is important to understand how these welfare costs vary across economies with different degrees of price flexibility as measured by their frequency N. It can be proved that the cumulative welfare costs from two economies with the same distribution of price changes are inversely proportional to their frequency of price changes.

For example, although not perfect, in Table 4, we can see that the cumulative welfare costs of the inflation episode for the food and beverages sector is roughly a third of the welfare costs for the aggregate PRISMA data, correspondingly the frequency of adjustments in the food and beverages sector is 2.4 times larger.

5 Concluding remarks

The New Keynesian paradigm, which greatly influences modern monetary economics, assumes that firms' prices are somewhat rigid and unresponsive to fundamental shocks, at least temporarily. We concentrate on two inefficiencies that emerge from this framework. First, the assumption of sticky prices implies distorted prices, impeding efficient resource allocation—this is known as the welfare cost of misallocation. Second, firms waste resources to maintain prices near optimal levels, resulting in a deadweight loss—we term this as the welfare cost of price management.

We propose a methodology to calculate both welfare costs for the period preceding 2022 and during the subsequent high inflation period triggered by substantial energy shocks in Europe. This task is complex because the welfare costs are not directly observable, as is often the case in welfare economics. To measure these costs, we apply a tractable sticky-price model, mapping it to detailed data from PriceStats and the ECB's PRISMA project.

Our findings reveal significant welfare costs for the Euro Area. In the low-inflation environment that prevailed before 2022, the efficiency cost amounts to roughly 2% of GDP. About 3/4 of these steady-state costs result from misallocation, while the remaining 1/4 is attributed to the costly price management activity.

Moreover, we estimate how the recent inflationary shock has affected these welfare costs. We found that the energy shocks led to an above-average surge in costly repricing activity, cumulatively adding to a cost of around 1.2% of GDP. A comparable increase in the welfare cost was caused by a temporary increase of the economy's misallocation. The total cumulated welfare costs of the energy shock thus range at about 2.7% of GDP.

We see our contribution as providing a first step in quantifying the welfare costs of misallocation in NK models. Indeed, we hope that our analysis will serve as a pathway to future research. We emphasize that our findings are initial estimations subject to refinement and development as additional data and improved models become available.

References

Adam, K., A. Alexandrov, and H. Weber (2023, April). “Inflation distorts relative prices: theory and evidence”, *CEPR Discussion Papers*, No 18088.

Aiyagari, S. R., R. A. Braun, and Z. Eckstein (1998). “Transaction services, inflation, and welfare”, *Journal of Political Economy,* Vol. 106, No. 6, pp. 1274–1301.

Alvarez, F., M. Beraja, M. Gonzalez-Rozada, and P. A. Neumeyer (2019a). “From hyperinflation to stable prices: Argentinas evidence on menu cost modles”, *The Quarterly Journal of Economics*, Vol. 143, No. 1, pp. 451–505.

Alvarez, F. and F. Lippi (2022). “The analytic theory of a monetary shock”. *Econometrica*, Vol. 90, No. 4, pp. 1655–1680.

Alvarez, F. and A. Neumeyer (2019, October). “The Passthrough of Large Cost Shocks in an Inflationary Economy", *Working Papers Central Bank of Chile*, No. 844.

Alvarez, F. E., M. Beraja, M. Gonzalez-Rozada, and P. A. Neumeyer (2019b). “From hyperinflation to stable prices: Argentina’s evidence on menu cost models”, *The Quarterly Journal of Economics*, Vol. 134, No. 1, pp. 451–505.

Alvarez, F. E., A. Ferrara, E. Gautier, H. LeBihan, and F. Lippi (2021, November). “Empirical investigation of a sufficient statistic for monetary shocks”, *NBER Working Paper*, No. 29490.

Alvarez, F. E., H. Le Bihan, and F. Lippi (2016). “The real effects of monetary shocks in sticky price models: a sufficient statistic approach”, *The American Economic Review*, Vol. 106, No. 10, pp. 2817–2851.

Alvarez, F. E. and F. Lippi (2014). “Price-setting with menu costs for multi product firms”, *Econometrica*, Vol. 82, No. 1, pp. 89–135.

Alvarez, F. E., F. Lippi, and A. Oskolkov (2022, May). “The macroeconomics of sticky prices with generalized hazard functions”, *The Quarterly Journal of Economics*, Vol. 137, No. 2, pp. 989–1038.

Alvarez, F. E., F. Lippi, and L. Paciello (2016). “Monetary Shocks in Models with Inattentive Producers”, *Review of Economic Studies*, Vol. 83, pp. 421–459.

Argente, D. and M. Lee (2020, 06). “Cost of Living Inequality During the Great Recession”, *Journal of the European Economic Association*, Vol. 19, No. 2, pp. 913–952.

Bach, G. L. and J. B. Stephenson (1974). “Inflation and the redistribution of wealth”, *The Review of Economics and Statistics*, Vol. 56, No. 1, pp. 1–13.

Bailey, M. J. (1956). “The welfare cost of inflationary finance”, *Journal of Political Economy,* Vol. 64, No. 2, pp. 93–110.

Baley, I. and A. Blanco (2021). “Aggregate dynamics in lumpy economies”, *Econometrica*, Vol. 89, No. 3, pp. 1235–1264.

Barro, R. J. (1972, January). “A theory of monopolistic price adjustment”, *Review of Economic Studies*, Vol. 39, No. 1, pp. 17–26.

Berger, D. and J. Vavra (2018). “Dynamics of the u.s. price distribution”, *European Economic Review*, Vol. 103, pp. 60-82.

Blanco, A., C. Boar, C. Jones, and V. Midrigan (2022). “Nonlinear Inflation Dynamics in Menu Cost Economies”. *Mimeo, NYU*.

Bonadio, B., A. M. Fischer, and P. Sauré (2019, 03). “The Speed of Exchange Rate Passthrough”, *Journal of the European Economic Association*, Vol. 18, No. 1, pp. 506–538.

Caballero, R. J. and E. M. Engel (1993a). “Heterogeneity and output fluctuations in a dynamic menu-cost economy”, *The Review of Economic Studies*, Vol. 60, No. 1,pp. 95.

Caballero, R. J. and E. M. Engel (2007, September). “Price stickiness in ss models: New interpretations of old results”, *Journal of Monetary Economics*, Vol. 54 (Supplement), pp. 100–121.

Caballero, R. J. and E. M. R. A. Engel (1993b). “Microeconomic adjustment hazards and aggregate dynamics”, *The Quarterly Journal of Economics*, Vol. 108, No. 2, pp. 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*, Vol. 67, No. 4, pp. 783–826.

Calvo, G. A. (1983, September). “Staggered prices in a utility-maximizing framework”, *Journal of Monetary Economics*, Vol. 12, No. 3, pp. 383–398.

Carvalho, C. and O. Kryvtsov (2018). “Price Selection”. *Staff Working Papers of the Bank of Canada*, No. 18-44.

Cavallo, A. (2017). "Are online and offline prices similar? Evidence from large multi-channel retailers." (pdf) American Economic Review 107, no. 1 (January 2017): 283–303

Cavallo, A. (2018). “Scraped data and sticky prices”, *Review of Economics and Statistics*, Vol. 100, No. 1, pp. 105–119.

Cavallo, A. and R. Rigobon (2016). “The Billion Prices Project: Using online prices for measurement and research”, *NBER Working paper*, No. 22111.

Costain, J. and A. Nakov (2011a). “Distributional dynamics under smoothly state-dependent pricing”, *Journal of Monetary Economics*, Vol. 58, No. 6, pp. 646–665.

Costain, J. and A. Nakov (2011b, 03). “Price adjustments in a general model of state– dependent pricing”, *Journal of Money, Credit and Banking*, Vol. 43, pp. 385–406.

Dias, D., C. R. Marques, and J. Santos Silva (2007). “Time- or state-dependent price-setting rules? evidence from micro data”, *European Economic Review*, Vol. 51, No. 7, pp. 1589–1613.

Dixit, A. (1991, January). “Analytical approximations in models of hysteresis”, *Review of Economic Studies*, Vol. 58, No. 1, pp. 141–51.

Doepke, M. and M. Schneider (2006). “Inflation and the redistribution of nominal wealth”, *Journal of Political Economy*, Vol. 114, No. 6, pp. 1069–1097.

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*, Vol. 114, No. 2, pp. 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*, Vol. 61, No. 1, pp. 159–188.

Eichenbaum, M., N. Jaimovich, and S. Rebelo (2011, February). “Reference prices, costs, and nominal rigidities”, *American Economic Review*, Vol. 101, No. 1, pp. 234–62.

Fougere, D., H. Le Bihan, and P. Sevestre (2007). “Heterogeneity in consumer price stickiness”, *Journal of Business & Economic Statistics*, Vol. 25, No. 3, pp. 247–264.

Friedman, M. (1968, October). “The Optimum Quantity of Money and Other Essays". *Aldine, Chicago*.

Gagnon, E. (2009). “Price-setting During Low and High Inflation: Evidence from Mexico”, *Quarterly Journal of Economics*, Vol. 124, No. 3, pp. 1221–1263.

Gali, J. (2008). “Monetary policy, inflation, and the business cycle : an introduction to the New Keynesian framework”, *Princeton University Press*.

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”, *European Central Bank Working paper*, 2669.

Gautier, E. and R. Le Saout (2015, September). “The Dynamics of Gasoline Prices: Evidence from Daily French Micro Data”, *Journal of Money, Credit and Banking*, Vol. 47, No. 6, pp. 1063–1089.

Golosov, M. and R. E. J. Lucas (2007). “Menu costs and phillips curves”, *Journal of Political Economy*, Vol. 115, pp. 171–199.

Gopinath, G. and O. Itskhoki (2010, 05). “Frequency of Price Adjustment and Passthrough”, *The Quarterly Journal of Economics*, Vol. 125, No. 2, pp. 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*, Vol. 11, No. 3, pp. 111–46.

Lagos, R. and R. Wright (2005). “A unified framework for monetary theory and policy analysis”, *Journal of Political Economy*, Vol. 113, No. 3, pp. 463–484.

Lucas, R. E. J. (2000). “Inflation and welfare”, *Econometrica*, Vol. 68, No. 2, pp. 247–274.

Nakamura, E. and J. Steinsson (2010, August). “Monetary non-neutrality in a multisector menu cost model”, *The Quarterly Journal of Economics*, Vol. 125, No. 3, pp. 961–1013.

Nakamura, E., J. Steinsson, P. Sun, and D. Villar (2018, 08). “The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation”, *The Quarterly Journal of Economics*, Vol. 133, No. 4, pp. 1933–1980.

Petrella, I., E. Santoro, and L. de la Porte Simonsen (2018). “Time-varying price flexibility and inflation dynamics”, *CEPR Discussion Paper*, No. 13027.

Sheremirov, V. (2020). “Price dispersion and inflation: New facts and theoretical implications”, *Journal of Monetary Economics*, Vol. 114, pp. 59–70.

Sheshinski, E. and Y. Weiss (1977, June). “Inflation and costs of price adjustment”, *Review of Economic Studies*, Vol. 44, No. 2, pp. 287–303.

Taylor, J. B. (1980, February). “Aggregate dynamics and staggered contracts”. *Journal of Political Economy*, Vol. 88, No. 1, pp. 1–23.

Walsh, C. E. (2010, December). “Monetary Theory and Policy, Third Edition, Volume 1”, *The MIT Press*.

Woodford, M. (2009, October). “Information-constrained state-dependent pricing”, *Journal of Monetary Economics*, Vol 56, pp. 100–124.

Zbaracki, M. J., M. Ritson, D. Levy, S. Dutta, and M. Bergen (2004). “Managerial and customer costs of price adjustment: Direct evidence from industrial markets”, *The Review of Economics and Statistics*, Vol. 86, No. 2, pp. 514–533.