Review of Economic Studies (2025) **92**, 2165–2196 https://doi.org/10.1093/restud/rdae082

Advance access publication 9 August 2024

Empirical Investigation of a Sufficient Statistic for Monetary Shocks

Fernando Alvarez
University of Chicago and NBER, USA
Andrea Ferrara
Northwestern University, USA
Erwan Gautier
Banque de France and Paris-Dauphine University, France
Hervé Le Bihan
Banque de France, France and Curtin University, Australia
and
Francesco Lippi
LUISS University and EIEF, Italy

First version received October 2021*; Editorial decision May* 2024*; Accepted August* 2024 *(Eds.)*

In a broad class of sticky-price models, the non-neutrality of nominal shocks is captured by a simple sufficient statistic: the ratio of the kurtosis of the price change distribution over the frequency of price changes. We test the sufficient statistic proposition using data for a large sample of products representative of the French economy. We first extend the theory to allow for empirically relevant monetary shocks with a transitory predictable component. We then use the microdata to measure kurtosis and frequency for about 120 producer price indices industries and 220 consumer price indices categories. We use a Factor-Augmented Vector Autoregressive (FAVAR) model to measure the industries' response to monetary shocks, under alternative identification schemes. The estimated degree of non-neutrality correlates with the kurtosis and the frequency consistently with the predictions of the theory. Several robustness checks are discussed.

Key words: Impulse response functions, Monetary shocks, Generalized hazard function, Sticky prices, Sufficient statistic

JEL codes: E3, E5

The editor in charge of this paper was Kurt Mitman.

1. INTRODUCTION

A central question in macroeconomics concerns the speed at which prices adjust to fundamental shocks. When such adjustments are slow monetary policy is non-neutral, *i.e.* nominal shocks have long-lasting output effects. In general, several features of a model contribute to determine the "degree of non-neutrality" of the economy. A recent theoretical result identifies a sufficient statistic for monetary shocks in a broad class of new Keynesian models: the non-neutrality is proportional to the ratio of the *kurtosis* of the (non-zero) price change distribution over the *frequency* of price changes. The result is useful because it suggests what features are important to determine the degree of non-neutrality. The result was established by Alvarez, Le Bihan *et al.* (2016) for the sticky price model of Nakamura and Steinsson (2010), nesting two workhorse models of macroeconomics: Calvo (1983) and Golosov and Lucas Jr (2007). The result was extended by Alvarez *et al.* (2022) to a broader class of state-dependent models using the generalized hazard function setup of Caballero and Engel (1993, 1999). Alvarez, Lippi *et al.* (2016) showed the same sufficient statistic to hold in models where firms follow time-dependent rules, as in Reis (2006). Given the multitude of theoretical setups that produce this prediction, Leahy (2016) considered the empirical test of the sufficient-statistic proposition a priority for this research programme.[1] This paper takes up that challenge and presents such a test.

We begin by extending the theoretical framework, developed for once-and-for-all permanent shocks, to accommodate shocks with a predictable transitory component. Such an extension is important to map the model to the data, where nominal interest rate shocks are typically mean reverting.[2] The extension shows that the sufficient statistic proposition remains informative about monetary non-neutrality even in the presence of mean-reverting shocks.

We then test the sufficient-statistic predictions using microdata for a large number of firms, representative of the French economy, underlying the producer price indices (PPI) and the consumer price indices (CPI). The test is made of three steps. We first estimate the sectoral responses to a monetary shock for about 120 PPI industries and 220 CPI categories, using a Factor-Augmented Vector Autoregression (FAVAR) in the vein of Bernanke *et al.* (2005) and Boivin *et al.* (2009). We summarize the extent of the non-neutrality using the cumulative impulse response of the sectoral prices (CIR^P). As the sufficient-statistic proposition concerns the cumulated response of *output*, we use the theory to derive the testable implications for the cumulated response of *prices*. This allows us to increase the number of cross-sectoral observations since output data are scarce relative to pricing data, and to map the theoretical prediction into a metric that is more robust.[3] The second step consists of using the microdata underlying the sectoral data to measure the frequency, kurtosis, and other moments of the price change distribution in the different sectors. In the third step, we inspect the relationship between the CIR^P and the sectoral moments under the restrictions implied by the theory.

The empirical findings support the claim that the ratio of kurtosis over frequency captures the degree of monetary non-neutrality across sectors. For the case of PPI data, both the frequency and the kurtosis consistently appear as statistically significant factors in accounting for the cross-sectional heterogeneity of the estimated CIR^P. The magnitudes of the regression coefficients are consistent with the predictions of the theory. Moreover, "placebo" tests show that moments *not* suggested by the theory, such as the average size, standard deviation, and skewness

1. He wrote: "I would not expect this equation to fit the data perfectly, but it would be very nice to know if these statistics are at all informative" (page 462 of Leahy 2016).

2. We handle this problem using the mean-field-game setup developed by Alvarez *et al.* (2023). The solution method revolves around a linearization, along the lines explored numerically by Boppart *et al.* (2018).

3. The output response depends on sector-specific elasticities that require additional information for testing the sufficient-statistic proposition.

of price changes, are not correlated with the CIR^P when controlling for the candidate sufficient statistic. These results hold across a variety of tests, specifications, and robustness exercises.

In the case of CPI data, the support for the theoretical predictions is weaker. This outcome likely reflects the fact that the model underlying the sufficient statistic result assumes no seasonal sales, an important pattern of consumer prices dynamics. We find that when removing products with frequent sales and product substitutions (in particular, food, clothing, and furniture), the CPI results align more closely with the theory. It is also possible that some pricing patterns for CPI items are driven by mechanisms that are absent from our theoretical set-up, such as consumer search.

Overall we find it noticeable that, across a broad range of specifications, both kurtosis and frequency are related to non-neutrality in a statistically significant manner that is aligned with the theory. We also find that kurtosis and frequency explain a small fraction of the cross-sectional differences in non-neutrality. This finding is consistent with the simplicity of the model and with the large error-in-variables detected in our estimates for the "degree of non-neutrality" (the dependent variable of our statistical test).

Our paper relates to a voluminous applied literature that analyses the implications of price-setting patterns, in particular cross-sectoral heterogeneity, for the propagation of shocks.[4] The novelty of our paper is to test the sufficient-statistic proposition for monetary shocks. The theory guides our empirical analysis: it identifies the variables of interest, how they enter the test, and shows how to interpret the sign and magnitude of the estimated coefficients, as well as the regressions' fit. Previous studies highlighted the importance of the frequency of price changes as a factor behind the cross-sectoral response to an aggregate shock, *e.g.* Nakamura and Steinsson (2010), Gopinath and Itskhoki (2010), Gorodnichenko and Weber (2016), and La'O and Tahbaz-Salehi (2022). The sufficient statistic proposition supplements the predictions for the role of frequency with the prediction for the role of kurtosis.

A related recent analysis by Hong *et al.* (2023), based on the US data, inspects the correlation between the response of sectoral producer price indices and several cross-sectional moments of the price change distribution. The authors show that frequency affects the degree of non-neutrality, but find less clear support for the role of kurtosis and other moments, especially when judged by their statistical "explanatory power". While their evidence is interesting, the lack of an explicit link between the empirics and a structural model prevents the analysis from conveying precise information on the sufficient statistic proposition. In particular, their analysis disregards the effect of measurement error in the constructed variables, a main explanation for the variables' limited explanatory power. Moreover, their regression specifications rarely conform to the specification of the theory we aim to test.

Our paper is also related to Gautier *et al.* (2023) who empirically test the sufficient-statistic prediction using granular data on costs and prices measured at gas stations, following a procedure similar to the one of our paper. Gasoline prices offer a clean laboratory for the test since they allow the authors to quantify, with little measurement error, the price response to a marginal cost shock, as well as the frequency and the kurtosis of price changes. The test strongly supports the sufficient statistic prediction: both kurtosis and frequency, but none of the other moments, affect the industry's non-neutrality with magnitudes that align closely with the theory. One limitation of their results is that they only pertain to a specific product. Our paper performs an investigation for a much broader set of products, covering a large share of the economy at both production and retail levels. This implies that measurement error issues need to be addressed. In

4. See, *e.g.* Bils and Klenow (2004), Burstein *et al.* (2005), Carvalho (2006), Bouakez *et al.* (2009), Imbs *et al.* (2011), Bonomo *et al.* (2023), Carvalho *et al.* (2021), Dedola *et al.* (2021), and Auer *et al.* (2021).

addition, since we look at several sectors, we cannot rely on straightforward sectoral-level measures of the marginal cost (as in the case of gasoline prices). Our approach, therefore, involves the identification of a monetary policy shock, common to all sectors.

The paper is organized as follows. Section 2 recalls the sufficient-statistic result and extends it to mean-reverting monetary shocks. Section 3 derives the theoretical restrictions to be tested on the data. Section 4 uses micro and sectoral data to measure the key ingredients needed to test the theory: (i) the sectoral response of prices (and output) to monetary shocks, and (ii) several cross-sectoral micro moments of the price change distribution. Section 5 presents the baseline results of the test using cross-sectional data. Section 6 investigates the robustness of our findings. Section 7 concludes.

2. A SUFFICIENT STATISTIC FOR MONETARY SHOCKS

We present a new Keynesian model, based on Nakamura and Steinsson (2010), that nests several sticky price models differing in the overall degree of price stickiness. We describe the environment, the firm's price-setting decision, and study the propagation of an aggregate shock. We review the sufficient-statistic result for monetary shocks by Alvarez, Le Bihan *et al*. (2016). The main novelty of our analysis is to extend the baseline case, focused on a once-and-for-all nominal shock, to a shock featuring a transitory component. This allows us to relate the theory to the transitory nominal interest rate shocks often used in empirical work.

2.1. *The model economy*

This section describes the foundations of the model.

2.1.1. Households, wages, money, and interest rates. We follow Golosov and Lucas Jr (2007), augmented to have n industries, and assume household preferences

$$\int_0^\infty e^{-\rho t}\left[\sum_{j=1}^{n}\frac{c_j(t)^{1-\epsilon_j}}{1-\epsilon_j}-\alpha L(t)+\log\left(\frac{M(t)}{P(t)}\right)\right]\mathrm{d}t \text{ and } c_j(t)\equiv\left[\int_0^1\left(A_{ij}c_{ij}(t)\right)^{\frac{\eta_j-1}{\eta_j}}\mathrm{d}i\right]^{\frac{\eta_j}{\eta_j-1}} \tag{1}$$

where ρ is the discount rate, $c_j(t)$ is a CES aggregate across the varieties i sold in industry $j=1,\dots,n$, $L(t)$ is labour, $\alpha>0$ a labour disutility parameter, $M(t)/P(t)$ is real money holdings, A_{ij} are preference shocks, and the parameters $\{\epsilon_j,\eta_j\}$ denote the substitution elasticity between and within industries.[5]

The household maximizes utility subject to the budget constraint

$$M(0)+\int_0^\infty Q(t)\left(\mathcal{T}(t)+W(t)L(t)-R(t)M(t)-\int_0^1\sum_{j=1}^{n}P_{ij}c_{ij}\mathrm{d}i\right)\mathrm{d}t\geq 0$$

where $\mathcal{T}(t)$ is a lump sum transfer, $R(t)$ the nominal interest rate, $Q(t)\equiv e^{-\int_0^t R(s)\mathrm{d}s}$ is the price of the time t nominal bond, $W(t)$ is the nominal wage, and $P_{ij}(t)$ is the price of good i in the industry j. Letting λ denote the Lagrange multiplier of the consumer's budget constraint,

5. See Appendix B in Alvarez and Lippi (2014) for a detailed analysis of this model.

we have the following first-order conditions: $e^{-\rho t}\alpha = \lambda Q(t)W(t)$, $e^{-\rho t}\frac{1}{M(t)} = \lambda Q(t)R(t)$, and $e^{-\rho t}c_j(t)^{-\epsilon_j} = \lambda Q(t)P_j(t)$, where $P_j(t)$ is the price index of industry j.

The first-order conditions imply that the steady-state interest rate is $\bar{R} = \rho + \frac{\dot{M}}{M}$, so that $\bar{R} = \rho$ if the money stock is constant. Moreover, we have $\alpha M(t)R(t) = W(t)$ so that shocks to the money supply or the interest rate map immediately into nominal wages. Using the definition of $Q(t)$, we have $W(t) = W(0)\, e^{\int_0^t (R(s)-\bar{R})\mathrm{d}s}$. Denoting the (after shock) steady-state wage by $\bar{W} = \lim_{t\to\infty} W(t)$, we write

$$\mathcal{W}(t) \equiv \log \frac{W(t)}{\bar{W}} = -\int_t^\infty \left(R(s) - \bar{R}\right) \mathrm{d}s \tag{2}$$

Equation (2) shows that a transitory deviation of the interest rate from the steady-state maps into a path of nominal wages. The often studied once-and-for-all shock to the money supply amounts to a shock that immediately triggers a new steady state level of the nominal wage, $W(t) = \bar{W}$ for all t, with no effects on the path of the interest rate so that $R(t) = \bar{R}$ and $\mathcal{W}(t) = 0$ for all t. In general, a monetary shock is made of two independent components: the permanent effect on nominal wages ($\bar{W}$), and the transitory wage deviations due to the interest rate dynamics (the path of $\mathcal{W}$).[6]

2.1.2. The firm's price setting problem. The firm maximizes profits subject to a random fixed cost for changing prices described in detail below. The production function is $c_{ij}(t) = \frac{L_{ij}(t)}{Z_{ij}(t)}$, where $1/Z_{ij}(t)$ is the firm's labour productivity. We assume that $A_{ij}(t) = Z_{ij}(t)$ so that the marginal cost and the preference shock are perfectly correlated.[7] The (log) markup for firm i in industry j is defined as the price over the unit labour cost: $\mu_{ij}(t) \equiv \log \frac{P_{ij}(t)}{W(t)Z_{ij}(t)}$.

Since the firm faces a demand with constant elasticity, we let μ_j be the time-invariant optimal markup. Define the "markup-gap" for firm i in industry j as

$$g_{ij}(t) \equiv \mu_{ij}(t) - \mu_j = x_{ij}(t) - \mathcal{W}(t) \quad \text{where } x_{ij}(t) \equiv \log \frac{P_{ij}(t)}{\bar{W} Z_{ij}(t)} - \mu_j \tag{3}$$

and $\mathcal{W}(t)$ is given in equation (2). Note that we define the firm's markup gap (g_{ij}) in terms of two variables, namely, x_{ij} and $\mathcal{W}$. This will be useful to highlight the novelty compared to the canonical case with a once and for all shock, where $\mathcal{W} = 0$.

Assume that $\log Z_{ij}$ follows a driftless diffusion, so that each firm is hit by idiosyncratic shocks $\mathrm{d}x_{ij} = \sigma_j d\mathcal{B}_{ij}$, where each $\mathcal{B}_{ij}$ is a standardized Brownian motion, independent across i and j. The presence of sticky prices, and the stochastic productivity shocks, imply that the firm's markup will not be equal to μ_j at every moment. We assume the initial conditions are such that $\int \log Z_{ij}(t)\mathrm{d}i = 0$ for all industries j.

To adjust its price and control the markup the firm must pay the fixed cost ψ_j. Alternatively, with a rate ζ_j per unit of time, the firm can adjust the price at no cost (a free adjustment

6. The nominal rate changes correspond to a path of money growth, implied by the above equations.

7. This assumption, also used in Woodford (2009) and 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.

opportunity à la Calvo). In what follows, we concentrate on an arbitrary firm i in industry j. The firm solves the following stopping-time problem

$$\min_{\{\tau_k, x_k^*\}_{k=1,2,\ldots}} \mathbb{E}\left[\int_0^\infty e^{-\rho t} B_j \left(x_{ij}(t) - \mathcal{W}(t)\right)^2 \mathrm{d}t\right] + \sum_{k=1}^{\infty} e^{-\rho \tau_k} \mathcal{I}(\tau_k)\ \psi_j \tag{4}$$

where τ_k is a price-reset time, and the indicator function $\mathcal{I}(\tau_k) = 0$ if the stopping time is due to a free-adjustment opportunity. The quadratic flow cost is derived from a second-order approximation of the firm's profit function around the optimal price, where $B_j \equiv \frac{\eta_j(\eta_j - 1)}{2} > 0$ is determined by the curvature of the demand in industry j. A feature of this model is that the firm's decisions are not affected by first-order deviations of the nominal interest rate or aggregate consumption.[8]

Intuitively, the firm's problem is to control $x(t)$ in order to track $\mathcal{W}(t)$. The time invariant parameters B_j, ζ_j, ψ_j and σ_j^2 depend on the industry j. Absent aggregate shocks, *i.e.* $\mathcal{W} = 0$, each firm's gap is only affected by the idiosyncratic productivity shocks. The firm's steady-state policy in industry j consists of a region where control is not exercised if $x \in [\underline{x}_j, \bar{x}_j]$; outside of this region control is exercised and the state is reset to x_j^*. Since the state is driftless the symmetry of the problem implies that $\underline{x}_j = -\bar{x}_j$ and $x_j^* = 0$, *i.e.* it is optimal for firms to "close the price gap" upon adjustment. In general, the firm's decisions are given by three (non-stochastic) time paths: $\underline{x}_j(t), \bar{x}_j(t)$ and $x_j^*(t)$.

We note that the Calvo-plus setup, akin to the model of Nakamura and Steinsson (2010), nests a large class of sticky price models, including the canonical menu-cost model and the Calvo model.[9] These models are indexed by a single parameter, the "Calvoness index" $\ell_j \equiv \sqrt{\frac{2\zeta_j}{\sigma_j^2/\bar{x}_j^2}}$, namely, the ratio between the number of free adjustments (ζ_j), and the number of adjustments that occur in a canonical menu cost model ($\sigma_j^2/\bar{x}_j^2$). If $\ell_j \to 0$ then the model corresponds to the canonical menu cost problem, while $\ell_j \to \infty$ gives the Calvo model. Both the frequency of price adjustment and the kurtosis depend on the Calvoness index. In particular, as shown in Alvarez, Le Bihan *et al.* (2016), the kurtosis of the size distribution of price changes is an increasing function of ℓ_j only, ranging between $Kurt = 1$ for $\ell_j = 0$ to $Kurt \to 6$ as $\ell_j \to \infty$.

2.1.3. Modelling aggregate shocks. The class of models we consider posits that monetary shocks affect the firm's marginal costs. In these models nominal wages and the money supply are proportional to each other, so that a positive monetary shock $\delta > 0$ increases the wages of all firms. Each firm aims to track such costs with its own price, to keep the markup close to the profit maximizing level. As shown in equation (2), the log of nominal wages follows $\log W(t) = \log \bar{W} + \mathcal{W}(t)$. If the shock is permanent it increases $\bar{W}$, reducing the markup gaps x of all firms by $-\delta$, see equation (3). At this new wage level firms charge prices/markups that are too low, thus the output level increases (this is the impact effect). Over time prices will permanently catch up with nominal wages and output will return to the steady-state level.

2.1.4. Once and for all aggregate shock. In the traditional analysis focusing on permanent monetary shocks, $\mathcal{W}(t) = 0$ for all $t > 0$ since the firm's nominal cost jumps up at time zero and remains constant afterwards, as it happens after a once and for all increase of the money supply. In this case, the firm's decision rules are unaffected by the aggregate shock (see Proposition 7

8. A key assumption for this result is that the demand system has a constant elasticity. See Proposition 7 in Alvarez and Lippi (2014) for the proof.

9. This setup can be further enriched using a generalized hazard function as in Caballero and Engel (1999, 2007); see Alvarez *et al.* (2022) for an analysis of this case.

in Alvarez and Lippi 2014) and upon adjustment firms "close the gap" (*i.e.* $x_j^*(t) = 0$ for all t). It is immediate from equation (2) that the nominal interest rate remains flat at the steady state level $\bar{R}$, a fact at odds with the monetary shocks considered in the empirical literature, where the nominal rates are not constant.

2.1.5. Aggregate shock with a predictable component. A novel element of our analysis is to consider a shock that involves a whole path $\mathcal{W}(t)$. These shocks amount to a perturbation of the entire *path* of the aggregate nominal costs. As shown in equation (2), this occurs when the monetary shock consists of a transitory shock to the interest rate, as often considered in monetary analysis. In this case, the firm's decisions are given by the time paths: $\underline{x}_j(t)$, $\bar{x}_j(t)$ and $x_j^*(t)$ for each industry j, and the optimal pricing policy at each t is given by the interval $(\underline{x}_j(t), \bar{x}_j(t))$ so that if $x(t)$ is in this interval the firm does *not* exercise control, *i.e.* inaction is optimal. Instead, if $x(t) \notin (\underline{x}_j(t), \bar{x}_j(t))$, the firm immediately changes its price from $x(t^-)$ to $x(t^+) = x_j^*(t)$.[10] The optimal policy $x_j^*(t)$ depends on the future path of $\mathcal{W}(s)$, for $s > t$, and hence the optimal policy upon adjustment is in general different from "closing the gap" ($x_j^*(t) \neq 0$).

2.2. *The sufficient statistic result*

The output of industry j, in deviation from steady state, $Y_j(t)$, is proportional to the cross-section mean of the price gaps. Using equation (3), we have

$$Y_j(t) \equiv -\frac{1}{\epsilon_j}\int g_{ij}(t)\mathrm{d}i = \frac{1}{\epsilon_j}\left(\mathcal{W}(t) - \int x_{ij}(t)\mathrm{d}i\right) \tag{5}$$

where ϵ_j is the industry-specific income elasticity which depends on the "demand side" of the economy, *i.e.* it is independent of the firm's price setting decisions. We define the cumulative impulse response of output, CIR^{Y_j}, as:

$$CIR^{Y_j} \equiv \int_0^\infty Y_j(t)\,\mathrm{d}t \tag{6}$$

where $Y_j(t)$ is the aggregate output t periods after the shock, measured in deviation from the steady-state output. The variable CIR^{Y_j} is a convenient statistic that summarizes with a single number the overall impact of the monetary shock on output in industry j.

2.2.1. The CIR with a once-and-for-all shock. Consider an invariant distribution of price gaps x that is hit by an unexpected permanent monetary shock of size $\delta > 0$. The shock immediately raises the nominal wage $\bar{W}$, while $\mathcal{W}(t) = 0$ for all t since the nominal interest rate does not change, so that all firms' markups fall by δ (log points). The shock triggers a dynamic response of output, following equation (5).

Letting $CIR_0^{Y_j}$ denote the output CIR after a once-and-for-all shock (denoted by the 0 subscript), the sufficient-statistic result of Alvarez, Le Bihan *et al.* (2016) establishes that after a *small* nominal shock δ we have

$$CIR_0^{Y_j}(\delta) = \frac{\delta}{\epsilon_j}\frac{Kurt_j}{6\,Freq_j} + o(\delta^2). \tag{7}$$

10. The superscripts + and − denote, respectively, the right and left limit of the variable.

The result states that the cumulated output response to a monetary shock is accurately approximated by the ratio of the kurtosis of the (non-zero) price change distribution ($Kurt_j$) to the frequency of price changes ($Freq_j$). The approximation is accurate up to second-order terms. The result in equation (7) is striking. It holds in a large class of inherently different models, from time-dependent models à la Calvo, to canonical menu-cost models à la Golosov-Lucas, intermediate cases such as the Calvo-Plus by Nakamura and Steinsson (2010) or inherently random-menu cost models such as those of Caballero and Engel (1993, 1999), as established by Alvarez *et al*. (2022).

The effect of the frequency is well understood: a higher frequency of price changes implies that adjustment is faster and hence the economy is more flexible (a smaller output effect). The effect of kurtosis is more subtle: it indicates that two industries with the same frequency can have substantially different flexibility. Kurtosis captures the fact that in an economy with heterogeneous agents, the response to an aggregate shock depends on the shape of the cross-sectional distribution of the agents's price gaps, a fact emphasized in several papers by Caballero and Engel. Consider, for instance, an economy where price setting is staggered every T periods, à la Taylor, and one where price setting follows a Calvo rule with an average duration equal to T. These economies have the same frequency of price changes but the Calvo economy features a thicker tail of "late adjusters", firms whose prices remain fixed long after T periods. Such a feature is captured by the kurtosis of the size of price changes, even in models where a time-dependent rule is followed as in Carvalho and Schwartzman (2015) and Alvarez, Lippi *et al*. (2016). Intuitively, kurtosis summarizes the degree of cross-sectional heterogeneity in the timing and size of price-setting behaviour. Equation (7) proves that this feature is important for the propagation of monetary shocks.

2.2.2. The CIR with a transitory shock.

Let CIR^{Y_j} denote the cumulative output defined in equation (6) for the case allowing for both the permanent and the transitory shock. Consider a shock with a permanent component $\log \bar{W} = \delta$ and a transitory component $\mathcal{W}(t) = \delta\omega_0 e^{-\gamma t}$, where γ parametrizes the half life of the shock and ω_0 its impact effect on the interest rate. For a small shock of size δ, we have the following result (see Appendix A for full details and the proof):

Proposition 1. *Consider $\rho \to 0$ and a small aggregate shock to nominal costs,* $\log W(t)$*, with both a permanent and a transitory component, namely,* $\log W(t) = \log \bar{W} + \mathcal{W}(t)$ *where* $\log \bar{W} = \delta$ *and* $\mathcal{W}(t) = \delta\omega_0 e^{-\gamma t}$. *Let* $\tilde{\sigma}_j \equiv \sigma_j^2/2$ *and* $\ell_j \equiv \sqrt{\frac{2\zeta_j}{\sigma_j^2/\bar{x}_j^2}}$. *The Cumulative Impulse Response of output is given by*

$$CIR^{Y_j} = CIR_0^{Y_j} + \frac{\delta}{\epsilon_j}\frac{\omega_0}{\gamma}\left\{\frac{\gamma/\tilde{\sigma}_j}{(\ell_j^2+\gamma/\tilde{\sigma}_j)} + \frac{\ell_j^2}{\left(1-e^{\ell_j}\right)^2}\left[(e^{2\ell_j}+1)\frac{csch\left(\sqrt{\ell_j^2+\gamma/\tilde{\sigma}_j}\right)}{\sqrt{\ell_j^2+\gamma/\tilde{\sigma}_j}} - (2e^{\ell_j})\frac{\coth\left(\sqrt{\ell_j^2+\gamma/\tilde{\sigma}_j}\right)}{\sqrt{\ell_j^2+\gamma/\tilde{\sigma}_j}}\right]\right\} \tag{8}$$

The proposition allows us to explore the robustness of the sufficient statistic result given by equation (7). The proposition highlights that the differences between CIR^{Y_j} and $CIR_0^{Y_j}$ relate to the degree of Calvoness of the model (ℓ), the persistence of the shock ($\gamma/\tilde{\sigma}$), and the size of

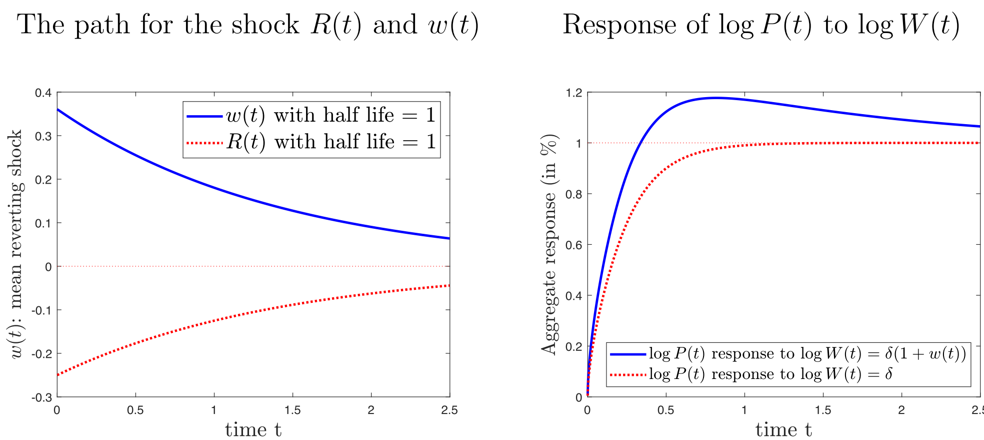

FIGURE 1

Mean reverting interest rate shock

Notes: The nominal shock equals $\log W(t) = \delta(1+\omega(t))$. The left panel plots the transitory component $\omega(t) \equiv w_0 e^{-\gamma t}$, and the associated nominal interest rate shock $R(t) - \bar{R} = \delta \hat{R}_0 e^{-\gamma t}$. We set $\delta = 0.01$ and consider a -25 basis points shock to the interest rate with a half-life of 1 year, which corresponds to $\hat{R}_0 = -1/4$ and $\gamma = 0.69$. By equation (2) then $\omega_0 = 0.36$. The right panel shows the aggregate price response to the $W(t)$ shock (thick solid line) and to the once and for all shock ($\omega(t) = 0$, dotted line).

the shock on impact ω_0. Once these parameters are quantified, the proposition will be used to evaluate the accuracy of the sufficient statistic proposition.

We consider the transitory component $\mathcal{W}(t) = \delta w_0 e^{-\gamma t}$ to parametrize the initial size of the interest rate shock (through ω_0 and equation (2)) and its persistence $1/\gamma$. Note that $\mathcal{W}(0) = \delta\omega_0 = -\int_0^\infty (R(s) - \bar{R}) \mathrm{d}s$. If $R(t) \equiv \bar{R} + \delta \hat{R}_0 e^{-\gamma t}$, then a given $\hat{R}_0$ implies that $\omega_0 = -\hat{R}_0/\gamma$. For instance, a 1% increase in the long-run nominal wage corresponds to $\delta = 0.01$. To supplement this shock with a 25 basis points reduction of the interest rate with a half life of 1 year, we set $\hat{R}_0 = -1/4$ and $\gamma = 0.69$, which implies $\omega_0 = 0.36$. The left panel of Figure 1 shows an example of a mean reverting interest rate shock, starting with an "expansionary" reduction of the interest rate (equal to 25 basis points) and an exponential decay with a half life of 1 year. The panel also shows the corresponding sequence for $w(t) \equiv w_0 e^{-\gamma t}$. The right panel shows the corresponding response of aggregate (log) prices.

We use Proposition 1 and the parametrization described above to quantify the deviation from the benchmark result with respect to the half life of the shock. The left panel of Figure 2 considers a "small shock" similar to the one described in Figure 1 where the interest rate decreases by 25bp on impact. The vertical axis reports the "normalized" ratio $CIR^{Y_j}/CIR_0^{Y_j}$, namely, the ratio of the CIR with the transitory shock relative to the CIR without it. This ratio is 1 if the shock's half life is zero, since in this case there is only the permanent component and there is no deviation between CIR^{Y_j} and $CIR_0^{Y_j}$. The ratio also converges to 1 as the shock becomes infinitely persistent (*i.e.* the transitory component vanishes). The biggest deviations occur for shocks with a half-life of about 1 year (about half the frequency of the price changes, which is set equal to 2 per year in the figure). Even so the maximal deviation is rather limited, below 10% of the prediction of the CIR for the case of the permanent shock. The different curves in the figure refer to different degrees of Calvoness. It appears that the largest deviations occur for the pure menu cost model (kurtosis equal 1), and that the deviations are smaller as the model gets closer to the Calvo model (high kurtosis).

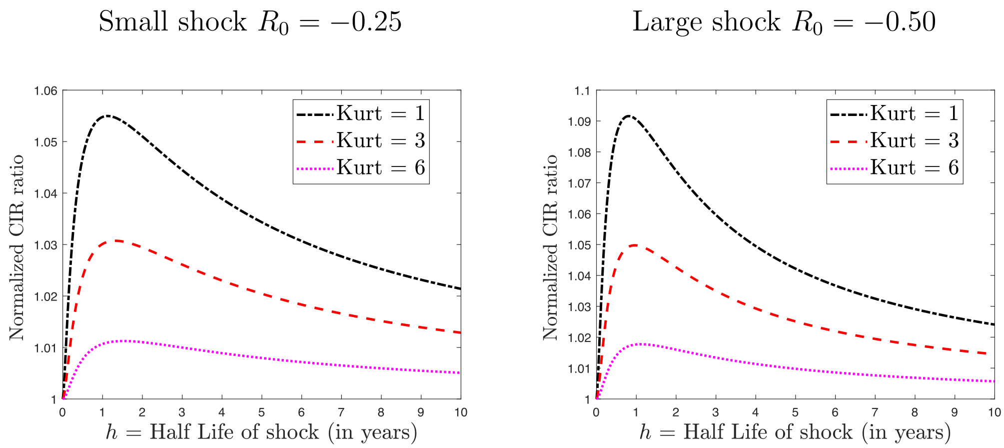

FIGURE 2

Normalized CIR^Y as a function of the shock's duration

Notes: The figure plots the cumulative impulse response of the model with the transitory shock relative to the one of the model without the transitory shock. Each panel displays three models, indexed by the degree of kurtosis (a function of ℓ). The computation uses equation (8).

Overall, these deviations are small, in view of the fact that equation (7) predicts deviations of the effects of monetary shocks that are in the range of 600% (as kurtosis varies from 1 to 6). In particular, for intermediate values of kurtosis as measured in the data (around 3, see Section 4.2 and Table 1 in this paper), the maximum value of the deviations is about 3%. Overall, we find this result reassuring about the informativeness of the sufficient statistic result, even in the presence of transitory shocks.

2.2.3. Key assumptions and limitations of the sufficient statistic result. Three assumptions are needed for equation (7) to hold in the benchmark case of a once-and-for-all shock. The first one is that the model has no inflation, so that several model objects are symmetric. While the assumption of zero inflation might seem restrictive, we argue that it provides a good approximation to models where inflation is low.[11]

The second assumption is that upon adjustment the firm resets x to zero. This assumption is violated in models with transitory shocks as discussed above, or in economies with high inflation, or in models with "price plans" or "sales", such as in Eichenbaum *et al*. (2011). It was shown in the discussion of Proposition 1, and in particular in Figure 2, that the violations due to transitory shocks are not a first-order concern. However, the presence of sales is potentially troublesome. If "sales" are not used to respond to aggregate shocks, then the sufficient statistic result can be recovered provided the number of price changes is measured net of the temporary ones, as in Kehoe and Midrigan (2015). But if temporary price changes are used to respond to shocks then the sufficient statistic result fails, as shown in Alvarez and Lippi (2020). In such cases, equation (7) is not a good summary of the impulse response.

A third assumption is that x follows a Brownian motion. This allows us to use stochastic calculus to analytically characterize the firm's optimal policy and the associated cross-sectional distribution of desired adjustments. In a model with leptokurtic shocks, such as Midrigan (2011), one cannot prove that kurtosis and frequency are enough to summarize the CIR^Y. However,

11. See proposition 7 in Alvarez, Le Bihan *et al*. (2016) for a rigorous argument, and Alvarez *et al*. (2019) for evidence supporting this claim.

TABLE 1
Micro moments of price adjustments: descriptive statistics

	Nb products	Mean	Q1	Q2	Q3	SD
Panel A: Frequency of price changes						
CPI	223	0.105	0.038	0.085	0.142	0.104
PPI	118	0.190	0.086	0.123	0.185	0.208
Panel B: Kurtosis of non-zero price changes—with robustness						
CPI—baseline	223	5.048	3.332	4.422	5.748	2.962
PPI—baseline	118	5.068	3.927	4.615	5.857	1.851
CPI—outlier $\|\Delta p\| < 0.5\%$	222	4.792	3.152	4.215	5.492	2.844
PPI—outlier $\|\Delta p\| < 0.5\%$	118	4.616	3.559	4.281	5.166	1.738
CPI—outlier $\|\Delta p\| > 35\%$	223	6.292	3.860	5.477	7.386	4.339
PPI—outlier $\|\Delta p\| > 25\%$	118	7.805	5.532	6.956	9.042	3.952
CPI—heterogeneity	223	3.413	2.188	3.181	3.837	2.019
PPI—heterogeneity	118	3.917	2.638	3.435	4.497	2.036
Panel C: Mean of non-zero price changes (%)						
CPI	223	1.263	0.296	0.988	2.127	2.120
PPI	118	0.793	0.204	0.722	1.405	0.906
Panel D: Standard deviation of non-zero price changes (%)						
CPI	223	7.590	6.024	7.294	9.301	2.331
PPI	118	4.149	3.606	4.134	4.674	0.872
Panel E: Skewness of non-zero price changes						
CPI	223	−0.260	−0.425	−0.247	−0.102	0.362
PPI	118	−0.274	−0.559	−0.275	0.028	0.444
Panel F: Average inflation (in %, absolute values)						
CPI	223	1.883	0.663	1.531	2.368	2.123
PPI	118	1.556	0.903	1.327	1.984	1.111

Notes: Calculations on CPI microdata are made over the period 1994–2019 (30 million of monthly price quotes). Prices of rents, cars, fresh food products, electricity, and clothing goods are non-available or excluded. Price changes due to sales and promotions are excluded (using the INSEE flag). VAT change and euro–cash changeover periods are excluded as well. Calculation on PPI data are made over the period 1994–2005. We report some descriptive statistics of the distribution of product-specific moments of price rigidity for PPI and CPI products (statistics are unweighted). "Frequency" reports the ratio between the number of price changes and the total number of prices. "Mean", "Standard deviation", "Skewness", and "Kurtosis" are calculated on the distribution of non-zero log price changes, expressed in percentages. In our baseline calculations, we have excluded all price changes below than 0.1% in absolute values and larger than 25% in absolute values for CPI price changes and 15% for PPI price changes. Panel F provides statistics on the average product-specific inflation in absolute values over the period 2005–2019. "Heterogeneity" refers to the measure of kurtosis taking into account for possible product heterogeneity following the methodology in Alvarez *et al.* (2022) (using five lags of price changes).

for moderate deviations from the Brownian benchmark, consistent with the data on the distribution of firms' nominal shocks, the formula continues to provide a useful benchmark (see Section 5 in Alvarez, Le Bihan *et al.* 2016 and the numerical results in Gautier and Le Bihan 2022).

We computed the impulse response by perturbing the stationary state of the model. Hence, what we compute is the "expected impulse response", *i.e.* averaging over the initial conditions, each of which is an entire distribution. Our setup could be used to study impulse responses where the initial condition is not the stationary distribution. This has been done in other models such as the empirical analysis by Caballero *et al.* (1997) in the context of employment, and the theoretical characterization for a price-setting model by Caplin and Leahy (1997). In our context, such analyses would require more extensive data, to fit the CIR at different aggregate states of the economy, as well as new theoretical results to characterize the state-dependent CIR.

3. AN EMPIRICAL TEST FOR THE SUFFICIENT STATISTIC RESULT

This section uses the predictions developed above to derive an empirical test of the theory. We will consider an economy made of several sectors, indexed by j, assuming that firms within a sector are similar, *i.e.* that they have the same response to a common monetary shock. The thought experiment is to hit this economy with an aggregate monetary shock and to use the variation in the responses observed across the sectors to test the theory.

The multi-sector set-up outlined above allows us to consider sectors that differ in the variability of the idiosyncratic shocks (σ_j), as well as in the pricing frictions (ψ_j and ζ_j). Equation (7) suggests testing the theory using a linear empirical relation between the product-level CIR of output over a long horizon, and the observed product-level ratios of kurtosis to the frequency of price changes. However, highly disaggregated sectoral output or real consumption series (at a monthly frequency) that match exactly the level of disaggregation and high frequency of observations typical of price data, are usually not available. In particular, in the case of France, there are no available monthly consumption volume data available at the same level of disaggregation as the CPI (we conjecture the same holds for other countries). In our empirical set-up, we thus rely on the cumulated impulse response of *prices* rather than *output*. One advantage of this strategy is also that both the micro and sectoral sets of variables derive from the same source of micro prices, ensuring consistency.

To obtain this alternative test, let us derive the relation between the cumulated response of output in sector j at horizon T, $CIR_T^{Y_j}$, and the one of the price level at the horizon T, $CIR_T^{P_j} \equiv \int_0^T P_j(t)\mathrm{d}t$, following a monetary shock of size δ. To lighten up notations, we assume a permanent shock so that $\log \bar{W} = \delta$ and $\mathcal{W}(t) = 0$ for all t. Using equation (6) we have

$$CIR_T^{Y_j} \equiv \int_0^T Y_j(t)\mathrm{d}t = \frac{1}{\epsilon_j}\int_0^T \left(\delta - P_j(t)\right)\mathrm{d}t = \frac{1}{\epsilon_j}\left(\delta T - CIR_T^{P_j}\right) \tag{9}$$

where δT is the cumulated change of the price level after a permanent increase of the money supply. Using equation (7), we can thus write $\lim_{T\to\infty}(\delta T - CIR_T^{P_j}) = \frac{\delta}{6}\frac{Kurt_j}{Freq_j} + o(\delta^2)$, or

$$CIR_T^{P_j} = \delta T - \frac{\delta}{6}\frac{Kurt_j}{Freq_j} \tag{10}$$

where the equation is an approximation since we omit $o(\delta^2)$ and we measure T for a finite horizon.[12] One feature of this specification (using CIR of prices instead CIR of output) is that the predictions for prices are independent of the sectoral elasticity ϵ_j, which simplifies how the regression coefficient should be interpreted. Another feature is that the equation has a straightforward interpretation: δT is the cumulated price response in the case of flexible prices, and hence the term $-\frac{\delta}{6}\frac{Kurt_j}{Freq_j}$ is the correction due to sticky prices.

For a feasible test of equation (10), we replace $CIR_T^{P_j}$ and $\frac{Kurt_j}{Freq_j}$ by the estimates $\widehat{CIR}_T^{P_j} = CIR_T^{P_j} + \nu_{cir,j}$ and $\widehat{\frac{Kurt_j}{Freq_j}} = \frac{Kurt_j}{Freq_j} + \nu_{kf,j}$ where $\nu_{cir,j}$ and $\nu_{kf,j}$ denote the corresponding measurement error. The variables denoted by "hats" are estimated from the microdata and the time series as discussed in the next section.

12. The error due to the finite horizon approximation is bound above by the Calvo model, where $\lim_{T\to\infty}\delta T - CIR_T^{P_j} = \delta/Freq$. The error with a finite T, relative to the exact value $\delta/Freq$, is $e^{-Freq\ T}$, or about $7\cdot 10^{-4}$ when evaluated at the baseline values $T = 36$ months and $Freq \approx 0.2$.

We can then write (10) as

$$\widehat{CIR}_T^{P_j} = \delta T - \frac{\delta}{6}\left(\frac{\widehat{Kurt_j}}{Freq_j}\right) + \nu_{cir,j} + \frac{\delta}{6}\nu_{kf,j}. \tag{11}$$

The measurement errors $\nu_{cir,j}$, $\nu_{kf,j}$ are two instances of classical measurement error in the right- and left-hand side variables. The measurement error in the ratio $\frac{Kurt_j}{Freq_j}$ generates a standard attenuation bias for the coefficient of interest. The size of this bias depends on the variance of $\nu_{kf,j}$ relative to the variance of $\frac{Kurt_j}{Freq_j}$. The measurement error in $CIR_T^{P_j}$ does not bias the estimated coefficients but reduces the regression's R^2: a large variance of $\nu_{cir,j}$, relative to the variance of $\frac{\delta}{6}\frac{Kurt_j}{Freq_j}$, leads to a low value of R^2. In Section 5, we will quantify the magnitude of the measurement errors using bootstrap methods and discuss implications for parameter estimates and overall fit in our regressions.

Our empirical test is thus based on the linear regression:

$$\widehat{CIR}_T^{P_j} = \alpha + \beta\left(\frac{\widehat{Kurt_j}}{Freq_j}\right) + \nu_j \tag{12}$$

where $\alpha = \delta T$ and $\beta = -\delta/6$ are the theory-implied values of the regression coefficients and ν_j is the regression's error term. In our empirical exercises, we normalize the monetary policy shock so that $\delta = -1$, leading, under a strict interpretation of the model, to the prediction that $\beta = 1/6$. The interpretation is that we focus on a contractionary monetary shock that reduces the long-run price level by 1%. We refer to this regression as the baseline regression, or as the "constrained regression", since the specification imposes that kurtosis and frequency enter the regression as a ratio.

We can further decompose equation (12) to investigate the restriction imposed by the theory on how kurtosis and frequency relate to the CIR. For that, we rely on a first-order Taylor expansion around the sample means of frequency and kurtosis, respectively, $\bar{F}$, $\bar{K}$, that gives $\widehat{CIR}_T^{P_j} \approx CIR^{\bar{P}_T} - \frac{\delta}{6}\frac{\bar{K}}{\bar{F}}\frac{\widehat{Kurt_j}}{\bar{K}} + \frac{\delta}{6}\frac{\bar{K}}{\bar{F}}\frac{\widehat{Freq_j}}{\bar{F}}$. From this expression, we derive an "unconstrained" version of the empirical test:

$$\widehat{CIR}_T^{P_j} = \beta_0 + \beta_k\left(\frac{\widehat{Kurt_j}}{\bar{K}}\right) + \beta_f\left(\frac{\widehat{Freq_j}}{\bar{F}}\right) + \nu_j \tag{13}$$

The theory implies that $\beta_k = -\beta_f$, *i.e.* the slope coefficients of the regressors $\frac{\widehat{Kurt_j}}{\bar{K}}$ and $\frac{\widehat{Freq_j}}{\bar{F}}$ are expected to have opposite signs and to be equal in absolute value.

4. MEASURING MONETARY SHOCKS AND SECTORAL MOMENTS

This section discusses the data used in the analysis, and the construction of the empirical statistics needed to test the sufficient statistic result. We use variation across products to test the theory. We rely on the existing cross-product variability in the price adjustment statistics, and on the fact that equation (12) is expected to hold across different sectors.[13] We need to estimate two types of statistics: (i) the cumulative impulse response of prices (CIR^P) computed at the sectoral level, and (ii) the moments of the distribution of price changes for the corresponding products. Sections 4.1 and 4.2, respectively, present our approach and results in computing those statistics.

Before providing more details on the construction of the objects underlying our test, we stress two important features of our empirical approach. First, we make use of a cross section of moments computed from two micro datasets of prices in France: a first one covering consumer prices and the other one producer prices. Both datasets are relevant for our purpose, and each has distinctive advantages. Consumer prices are observed directly and somewhat less prone to measurement issues (since they can be directly observed in outlets), offer a broader coverage of the economy (goods and services versus only goods for PPI products) and consumer inflation is used for the definition of the monetary policy target. Producer price data are conceptually closer to the firms' pricing problem studied in standard macro models and are not affected by sales and temporary promotions.

The second feature is that we identify the monetary shocks by imposing that they have the properties highlighted by the theory (in the spirit of the "sign restriction" approach). In particular, we want a (contractionary) shock to decrease output in the short run, to have a permanent negative effect on the price level, and to have no long-run effect on output. These characteristics are consistent with the theoretical model described above and are thus desirable to perform a test of the sufficient statistic result. Note that in principle any common shock to the marginal cost of firms could be used to test the theory. Oil price shocks would for instance qualify, but empirically the sectoral dynamics following such a shock is strongly heterogeneous making it hardly useable for a test in a finite sample. On the contrary, an aggregate monetary shock has the desirable features that it will eventually move all nominal prices by the same amount, leaving relative prices unaltered. We exploit this homogeneity property in our long-run identification of the monetary shock. Finally, we stress that another feature of our approach is that the construction of the CIR^P variables does not use the microdata nor the sectoral moments, so there is no reason to expect any bias in favour (or against) the sufficient statistic result.

4.1. *Measuring the sectoral response to a monetary shock*

To estimate the CIR^P for a large number of sectors, we employ a FAVAR, a method developed by Bernanke *et al*. (2005) and Boivin *et al*. (2009). We closely follow the approach of Boivin *et al*. (2009) as they focus on the response of sectoral inflation rates to monetary policy shocks. A brief description is as follows: the FAVAR is a model in which the dynamics of a large number of time series is governed by the evolution of a small number of factors, that are typically—but not necessarily—unobserved and follow a VAR process (see Appendix B for a more detailed description of the FAVAR model).

Formally, the vector of a large number N of time series X_t, called informational time series, are related to the factors F_t by the following equation: $X_t = \Lambda F_t + e_t$, where F_t is a vector of dimensions $K+M$ of respectively unobserved and observed factors, and e_t is a vector N x 1 of error terms with zero mean. Following Boivin *et al*. (2009), we allow one factor, the interest rate i_t, to be observed, so $F_t \equiv [\tilde{F}_t \; i_t]'$, where the unobservable factors $\tilde{F}_t$ are to be estimated. Consistently, matrix Λ has dimensions $(N, K+1)$, with $K+1$ the number of factors ($K = 5$

13. In the paper, we use indifferently the terms "sectors" and "products". For PPI, product and sector classifications fully overlap, whereas for CPI, we will use product-specific price indices.

in our application). Notice that the observable factors and the informative time series are two distinct objects that do not have any time series in common. The factors are assumed to follow a VAR process: $F_t = \Phi(L)F_{t-1} + \upsilon_t$ where $\Phi(L)$ is a lag polynomial of finite order and υ_t is an error term with zero mean and covariance matrix Q.

We are interested in estimating the response of the disaggregated time series of prices (PPI and CPI) after a monetary shock. In the first step, factors are computed from a Principal Component Analysis using the informative time series. We include three types of "informative time series" in vector X_t (see Appendix B for details): (i) macroeconomic data for France, (ii) financial and monetary variables relevant for the euro area, and (iii) disaggregated series of industrial production indices (IPI), producer price indices (PPI) and consumer price indices (CPI), for France (seasonally adjusted and taken in log differences). In addition, our analysis uses the 3-month Euribor as a measure of the monetary policy variable. This variable is treated as an observable factor and is filtered following motivations and a procedure detailed below. Data are monthly and the sample period is Jan. 2005 to Dec. 2019. From this first step, we extract five principal factors (those with the largest contributions to the overall variance) and we then estimate a VAR model with 12 lags for the 5 factors and the interest rate. From this VAR, we can retrieve the impulse response function (IRF) of all sectoral prices to an aggregate shock. The dynamics of inflation in sector j in our FAVAR are governed by:

$$\pi_{jt} = \lambda_j F_t + e_{jt} \tag{14}$$

where λ_j is a vector of loadings, recovered as the relevant row of matrix Λ. From these sectoral IRFs, the $CIR_T^{P_j}$ is calculated as the cumulated response of sectoral price levels over a large number of periods (see next section for a discussion).

4.1.1. Identifying monetary policy shocks and the price responses. In our baseline setup, to identify a contractionary monetary shock in our FAVAR model, we use a Cholesky decomposition of the variance–covariance matrix of the VAR innovations. Following a standard timing restriction, the Euribor is ordered as the last variable in the VAR. Notice that imposing a Cholesky decomposition in this setup does not imply that the IRFs of informative time series cannot respond simultaneously to the monetary shock.

In addition, in our baseline approach, to compute the CIR corresponding to the identified monetary policy shocks, we impose a "long run neutrality" restriction on relative prices, and output. Specifically, it is imposed that (i) output comes back to its original level in the long run after a monetary shock and (ii) all sectoral prices have identical responses—equal to that of the average price across sectors—in the long run.[14] Both of these restrictions are consistent with the money neutrality hypothesis. We follow Boivin *et al.* (2009) to implement the latter restriction in the baseline FAVAR specification.

To complement our baseline, we also consider an alternative set of estimates without the "long run neutrality" restriction on the effects of monetary shocks on relative prices—a case also considered by Boivin *et al.* (2009). It is worthwhile to note that this set of estimates does not differ from the baseline case as regards the identification of the monetary policy shock—only the effect on relative prices is altered, thus affecting the cross-sectional CIRs however. In addition, we explore another alternative procedure to compute the CIRs using a High-Frequency Identification (HFI) of monetary shocks, following Gertler and Karadi (2015). Such an approach,

14. In practice, we impose these restrictions at the horizon of 8 years. Note this horizon is independent of, and substantially longer than, the one over which we compute the cumulated IRFs (36 months in the baseline).

popular in recent applied work, allows us to deal with simultaneity issues without resorting to a timing assumption (as in the Cholesky approach). To implement this approach, we use monetary surprises in the euro area computed by Altavilla *et al*. (2019), relying on market interest rate changes around the times of ECB Governing Council meetings.[15]

In all empirical exercises, we normalize the shock, so that the monetary policy shock produces a 1% long-run decrease in the aggregate price level. This normalization assumption (which has no bearings in terms of inference) departs from the usual approach to normalization imposing that the shock produces an effect on the impact on the nominal interest rate. The normalization allows an easier comparison with our theoretical model (where the size of the shock is proportional to the long-run response of the price level) and facilitates the interpretation of results relating the CIR^P to the sufficient statistic.

4.1.2. Filtering the Euribor. The theory suggests that a (contractionary) monetary policy shock triggers a transient, and negative, impact on inflation and output. The VAR estimates based on unfiltered interest rate data produce IRFs that are not consistent with these predictions, a feature we relate to the marked downward trend in the nominal interest rate over the sample period (see Figure A.1 in the Appendix—a pattern likely related to the decline in the "natural rate" of interest).[16] Furthermore, the theory also suggests that all the sectors should have a negative IRFs of prices after a contractionary monetary shock.

In our FAVAR model, we thus filter the interest rate to ensure that the FAVAR model produces a negative and transient response of output and inflation after a contractionary monetary policy shock. To do so, we use a one-sided HP filter that does not use future data at any point in time, as a hedge against introducing spurious correlations when using the filtered interest rate in the VAR model. Our approach is to use a one-sided HP filter with a parameter λ^{HP} that maximizes the number of PPI and CPI sectors with negative IRFs after 36 months of the shock. Appendix C provides more details on our strategy to select the value for λ^{HP}.[17] Notice that our selection criterion relies on the sign of IRFs after 36 months, not on the *CIR* which will be used in our regressions. Furthermore, our procedure for selecting the filter parameter makes no use of the microeconomic data or the sectoral moments, so it is not biasing towards finding some relevance of the sufficient statistic results. Our FAVAR estimation procedure is designed to produce a shock that has the same negative price effect in all sectors and can be interpreted as a monetary policy shock.

We select a value of $\lambda^{HP} = 500,000$ for which about 70% of PPI and CPI products have a negative IRF after 36 months (see Appendix C for more details).[18] Note that the latter statement refers to the *point values* of the IRFs after 36 months. It is thus not in contradiction with the *cumulative* IRFs (CIR^P) being always negative, a more relevant prediction of the theoretical model. In fact, as Table A.1 in the Online Appendix indicates, depending on the horizon considered and the identification scheme, the estimates of CIR^P turn out to be negative for 95% to 100% of the products.

15. In another robustness exercise, we also report results using a longer term interest rate—the 2-year German Bond rate—as the policy rate, and using the same HFI approach, to account for non-conventional monetary policy shocks.

16. Identifying well-behaved monetary policy shocks for the euro area is particularly challenging over the sample period, in particular due to the proximity of the effective lower bound on interest rates—see Andrade and Ferroni (2021) and Jarocinski and Karadi (2020) for investigations in the context of information shocks.

17. Note that the literature does not agree on a specific value for the one-sided HP filter, unlike with the standard two-sided HP filter.

18. As a robustness check we also ran the whole empirical exercise including FAVAR and OLS product-level regressions with $\lambda^{HP} = 1$ million. The results are qualitatively and quantitatively the same.

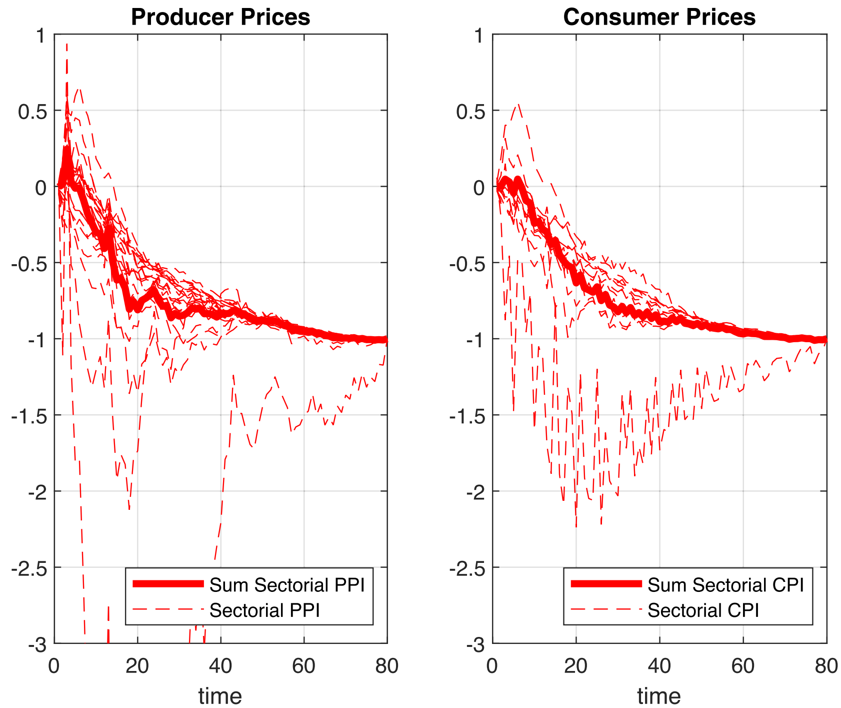

FIGURE 3

Sectoral responses of PPI and CPI to a contractionary monetary shock

Notes: Sectoral IRFs of PPI products (left panel) and sectoral IRFs of CPI products (right panel). All product-level IRFs are computed at a disaggregate product level; for CPI, the level of disaggregation is 5 digit-level of the ECOICOP classification (*i.e.* "01.1.1.1"), whereas for PPI, the product level is the four-digit level of the NACE rev2 classification of sectors. The dashed lines correspond to sectoral IRFs computed at a two-digit product level (the simple average over the most disaggregated product level IRFs used then in our OLS regressions), the thick solid line plots the average IRF computed over all disaggregated product-level IRFs.

4.1.3. VAR results: IRFs and CIR^P. Our estimated FAVAR provides theory-consistent results for the responses of aggregate variables to a monetary shock. After a contractionary policy shock the interest rate increases and subsequently decreases, going back to its steady-state level in the long run (IRFs are presented in Appendix Figure A.4). Industrial production immediately shrinks after a contractionary monetary policy shock, then gradually recovers. The production price index and the consumption price index both decline following the shock, converging toward the new steady-state values.

We focus our analysis on the objects used to test the sufficient statistic result, namely, the responses of sectoral producer and consumer prices, as derived from the FAVAR. Figure 3 reports the estimated IRFs of production and consumer price series. In each panel, dashed lines are the IRFs of different sectors partially aggregated, at the two-digit level for PPI, and one-digit level for CPI.[19] The thick solid line is the average of all the dashed lines. In both figures, we impose that the long-run price response is -1% at a long horizon (8 years). The transitory dynamics are however heterogeneous across sectors. Most of them display a trough in prices between 1 and 3 years after the shock.

19. Our PPI/CPI series are available at the four-digit and five-digit levels, and the dashed red lines are constructed as the arithmetic average of estimated IRFs.

Finally, using the estimated IRFs of the PPI and CPI, we construct the CIR^Ps for each sector/product category, as the sum of the respective IRF from time zero up to a time horizon T. We select a baseline value of $T=36$ months to compute CIR^Ps (see Table A.1 in Appendix for descriptive statistics on product-specific CIR^Ps) but we will also provide robustness analysis using two different values of T (24 and 48 months).

4.2. *Measuring micro moments*

4.2.1. CPI microdata. For consumer price microdata, we rely on longitudinal datasets of monthly price quotes collected by the Institut National de la Statistique et des Études Économiques (INSEE) to compute the monthly French CPI (Consumer Price Index). Stacking datasets used in Baudry *et al*. (2007), Berardi *et al*. (2015), and Berardi and Gautier (2016) and extending the dataset to September 2019, we obtain a long sample covering a period of about 25 years between August 1994 and September 2019.

The dataset contains about 30 million of price quotes and covers about 60% of the CPI weights.[20] Price changes are computed as log-differences of prices, and we exclude price changes due to sales. To compute price adjustment moments, we have first dropped data collected around VAT changes (*i.e.* in Aug.–Sept. 1995, Sept.–Oct. 1999, April–May 2000, July–Sept. 2009, Jan.–Feb. 2012, and Jan.–Feb. 2014) and before and after the euro cash changeover (between Aug. 2001 and June 2002). We have also dropped price changes smaller than 0.1% in absolute values, in both datasets, in order to control for possible small price changes due to measurement errors (Eichenbaum *et al*. 2014).

We compute price adjustment statistics excluding sales, as the model is not able to reproduce price changes due to sales. Our identification of sales relies on an INSEE flag variable that identifies whether a price corresponds to a sale price, either in the form of seasonal sales or temporary promotional discounts. We identify products at the five-digit level of the ECOICOP product classification, which is the most disaggregated level for which sectoral price indices are available. For each product, we compute the frequency of price changes as the ratio between the number of price changes and the total number of prices for this product. We also compute the kurtosis of price changes, as well as other moments of the price change distribution (such as the average, the standard deviation and the skewness of price changes) at the product level. Overall, our baseline dataset contains price adjustment moments for 223 different "ECOICOP-5" CPI products.

The measurement of kurtosis is notoriously a challenging issue, as large values of price changes, and outliers, can have an important impact on kurtosis. Very large kurtosis values tend to be obtained when not correcting for measurement errors.[21] In our baseline, we drop from the calculations price changes larger than 25% in absolute values, which corresponds to about 5% of all price changes. As robustness, we provide results with alternative values for the thresholds used to define outliers and address measurement error concerns (for very large or very small price changes in absolute values). Drawing on Alvarez *et al*. (2022), we also provide results using a measure of kurtosis that uses a correction for unobserved heterogeneity. Alternative kurtosis measures are highly correlated across products.

20. Some categories of goods and services are not available in our sample: centrally collected prices, among which car prices and administered prices (*e.g.* tobacco) or public utility prices (*e.g.* electricity), as well as other types of products such as fresh food or rents.

21. Note however that excluding sales by itself does not decrease the degree of kurtosis, see for instance Gautier and Le Bihan (2022).

4.2.2. PPI microdata. We rely on micro price data collected by INSEE to construct the French Producer Price Index. This dataset is the same as the one used in Gautier (2008) where further details are available. Reported prices must be observed at the "factory gate", excluding transport and commercialization costs, or invoiced VAT.[22] Our sample contains more than 1.5 million price reports between January 1994 and June 2005. Overall, more than 90% of the price quotes used to compute the French PPI are available. The PPI covers all products manufactured and sold in France by industrial firms, which includes sections C (Mining and quarrying), D (Manufacturing), and E (Electricity, gas, and water supply) of NACE Rev 2 classification. As for CPI, price changes are computed as log-differences of prices.

For each NACE four-digit sector, we compute both the frequency of price changes and the kurtosis of non-zero price changes, as well as other moments of the price change distribution. Unlike with CPI, large price changes are much less frequent (reflecting that sales or temporary promotions are absent in the business-to-business context of producer prices) and only 2% of all price changes are larger than 22% in absolute value. To measure kurtosis, we drop price changes larger than 15% in absolute values (which correspond to less than 5% of all price changes) and we test the robustness of our results to this definition of price change outliers. We restrict the sample to sectors for which an aggregate sectoral price index is available from the statistical office, so as to match micro moments with time-series macro evidence in our subsequent analysis. This results in a baseline sample containing 118 sectors.

Basic statistics for the microdata underlying both the CPI and the PPI are presented in Table 1. Consumer prices are more rigid than producer prices, with average frequencies of price changes of 10.6% and 19%, respectively. The distribution of price changes has fat-tails for both datasets, with a virtually identical value of the unweighted average kurtosis of 5.0 in both datasets. One main important takeaway is that there is some cross-sectoral dispersion in the frequency and kurtosis of price changes, for both consumer prices and producer prices—as apparent from the interquartile ranges or the standard deviations. The frequency of price changes however seems to show relatively more cross-sectoral variability than the kurtosis of price changes. While alternative corrections for measurement error and unobserved heterogeneity do change the average value of kurtosis, they do not substantially affect the degree of cross-product heterogeneity however.

Cross-sectoral characteristics of both our CPI and PPI datasets are consistent with available international evidence. As regards consumer price data, Berardi *et al.* (2015) using the same data, provide a detailed comparison of CPI data moments in France with those in the United States, based on detailed moments reported by Nakamura and Steinsson (2008). They conclude that patterns are quite similar, whenever sales-related price changes are disregarded (as the pattern of sales is however much more prevalent in the United States).[23] Regarding producer prices, Vermeulen *et al.* (2012) provide a comparison of the patterns of price setting in the United States and 6 euro-area countries, including France—relying for that particular country on the same dataset as we use. They conclude that patterns of producer price rigidities are very similar—albeit the size of price changes is typically larger in the United States than in Europe. The above-mentioned international evidence mainly focuses on the frequency of price changes, as well as on the first two moments of the distribution of price changes. Evidence is scarcer on kurtosis. For US PPI data, Hong *et al.* (2023) report an average kurtosis of 4.9. With consumer price data, Cavallo (2018) reports a median kurtosis of 4.8 in a large sample of countries based

22. Contrary to CPI prices, there is no flag for temporary promotions or sales. We assume, consistent with Nakamura and Steinsson (2008), that there are no sales in producer prices.

23. See also Gautier *et al.* (2024) for a detailed comparison of euro area evidence.

on “scraped” data. These values, all obtained after correcting for measurement errors in the same spirit as we do, are thus much in line with our baseline values.

5. TESTING THE THEORY: RESULTS

This section presents the results of the empirical tests developed in Section 3 using as inputs the product-level CIR^P (as measured in Section 4.1), and product-level moments of price adjustments (as measured in Section 4.2).

5.1. *Estimates of the baseline empirical specification*

This section presents our baseline estimation results. As detailed in Section 3, the theory predicts that, in case of a contractionary monetary policy shock, the coefficient associated with $Kurt/Freq$ ratio should be positive in the regression for CIR^P. Smaller values of the $Kurt/Freq$ ratio may stem from more frequent price adjustments, more price selection (smaller kurtosis), or both, implying more flexibility of the aggregate price index (the *lower* new steady-state level of prices is reached faster).

Table 2 reports the estimation results for equation (12), the baseline “constrained” regressions for horizon $T = 36$ months. The regressions for the CIR^P of PPI products are presented in Panel A while those for the CIR^P of CPI products are presented in Panel B. In each panel, we report the results for three specifications: (i) identification using Cholesky and imposing long-run restriction on relative prices (Columns 1 and 2) which is our baseline specification; (ii) Cholesky identification without imposing any restriction on the long-run effect on relative prices (Columns 3 and 4); and (iii) identification based on HFI and imposing long-run restriction on relative prices (Columns 5 and 6). For each of the three specifications, we run the regressions as in equation (12), without including any product “fixed-effects” (Columns 1, 3, and 5). Columns 2, 4, and 6 report results of equation (12) including dummy variables for two-digit level sectors for both CPI and PPI products. There are 38 such broad two-digit sectors in our sample for the CPI, and 24 in the case of the PPI.[24] We label this specification as the “fixed-effects” one. We are agnostic on whether the “fixed effect” case, or the no “fixed effect” one, is the most relevant specification. If the kurtosis-to-frequency ratio mainly varies across broad sectors, then introducing fixed effects will remove useful information and may obfuscate the results in a finite sample of data. By contrast, beside providing a hedge against spurious correlation, including fixed effects is relevant if the relation between CIR^P and the pricing-moments mostly holds within broad sectors. Both specifications inform us on the sources of product variability that help to identify the relation between CIR^P and the cross-sectional moments: across-sector differences versus within-sector variability.

For producer prices (Panel A), the estimated slope coefficient associated with the $Kurt/Freq$ ratio turns out to be positive and statistically significant in most cases. Adding “fixed effects” sectoral dummy variables weakens the significance of the estimated parameters, but the results are qualitatively and—for many specifications—quantitatively the same as in our baseline regressions. These results are consistent with the theoretical framework. Coefficients obtained in the case without the long-run restriction (Columns 3 and 4) are significant and with

24. As robustness, Table A.12 in the Appendix reports results using dummy variables for more aggregate sectors (6 for PPI and 12 for CPI).

TABLE 2
Baseline OLS regression results: "constrained" specification—36-month horizon

Identification	Cholesky		Cholesky		High-Freq. IV	
Long-run Restriction	Yes		No		Yes	
Product FE	No	Yes	No	Yes	No	Yes
PANEL A: PRODUCER PRICES						
Kurt/Freq	0.206***	0.116**	0.543**	0.231	0.189***	0.135**
	(0.0744)	(0.0544)	(0.247)	(0.175)	(0.0646)	(0.0555)
Constant	−27.93***	−17.85***	−42.64***	−34.19***	−34.08***	−27.29***
	(4.563)	(3.132)	(15.57)	(6.979)	(3.834)	(6.984)
Observations	118	118	118	118	118	118
R^2	0.095	0.534	0.058	0.468	0.110	0.452
P-val $\beta = 1/6$	0.598	0.358	0.129	0.713	0.736	0.568
PANEL B: CONSUMER PRICES						
Kurt/Freq	−0.000424	0.0426***	0.0290	0.0944**	0.0133	0.0333***
	(0.0145)	(0.0148)	(0.0387)	(0.0430)	(0.0104)	(0.0116)
Constant	−16.58***	−11.77***	−19.31***	−20.05***	−23.89***	−21.63***
	(1.958)	(1.185)	(5.776)	(4.434)	(1.584)	(0.764)
Observations	223	223	223	223	223	223
R^2	0.000	0.440	0.003	0.335	0.007	0.650
P-val $\beta = 1/6$	0.000	0.000	0.001	0.095	0.000	0.000

Notes: The table reports results of OLS regressions (equation (12)) where the dependent variable is the product-specific $CIR_T^{P_j}$ (calculated for the horizon $T = 36$ months and using $\delta = -1$) and the right-hand-side variable is the ratio $Kurt/Freq$. Each observation corresponds to a disaggregate CPI or PPI product. For CPI, the level of disaggregation is 5 digit-level of the ECOICOP classification (*i.e.* "01.1.1.1"), whereas for PPI, the product level is the four-digit level of the NACE rev2 classification of sectors. PPI covers the manufacturing sectors, whereas CPI covers about 60% of the whole French CPI (main products excluded are rents, cars, utilities like electricity). Product-fixed effects are defined at the two-digit level for both CPI and PPI products (*i.e.* 38 product-fixed effects for the CPI, and 24 in the case of the PPI). Robust standard errors are reported in parentheses. Stars (*) indicate the significance level of the p-values: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

expected signs but are larger than in the baseline case (Columns 1 and 2), presumably reflecting a larger degree of variability of the sectoral CIR^P (see Table A.1 in Appendix).[25]

For consumer prices, the results (Panel B of Table 2) are mixed. When sectoral fixed effects are not included, the coefficient associated with the $Kurt/Freq$ ratio is small and not significant. When sectoral fixed effects are included, the coefficients are positive and significant in the three specifications. These results suggest that the relationship between CIR^P and the pricing moments is driven mainly by within-sector variability rather than broad sector differences. For CPI, the incidence of sales could be of particular importance to explain why the relationship does not hold when looking at broad differences across sectors. The extent of sales could indeed affect price adjustment moments even if we have excluded price changes due to sales in the calculation of these moments. In particular, if a large majority of price adjustments are due to sales or promotions in one sector, the pricing moments excluding these changes might be not very representative of the typical price changes. To explore this, we develop alternative estimation exercises removing all food, clothing/footwear and furnishings goods, as within these broad sectors, most products are largely affected by seasonal sales and replacements.[26] In another set of additional estimates, we exclude CPI products for which more than 10% of all price changes are

25. The estimated values relying on the FAVAR with long-run restrictions are consistent with our theoretical setup where monetary shocks are neutral in the long run. The FAVAR without these restrictions puts less constraint on the data but is less connected to the theory.

26. These products correspond to COICOP 01.1, 03, and 05 in the product classifications.

TABLE 3
Regression results: role of sales—consumer prices

	Case 1: Excluding food, clothing/footwear, furnishings			Case 2: Products with % of sales prices below the median		
Identification	Cholesky	Cholesky	High-Freq. IV	Cholesky	Cholesky	High-Freq. IV
Long-run Restriction	Yes	No	Yes	Yes	No	Yes
PANEL A: Constrained model						
Kurt/Freq	0.0481***	0.0787	0.0426***	0.0434**	0.174***	0.0489***
	(0.0182)	(0.0569)	(0.0153)	(0.0210)	(0.0583)	(0.0172)
Constant	−22.03***	−21.45**	−27.20***	−20.47***	−41.03***	−28.48***
	(3.464)	(10.64)	(2.939)	(4.120)	(10.83)	(3.396)
R^2	0.067	0.019	0.068	0.050	0.103	0.083
P-val $\beta = 1/6$	0.000	0.124	0.000	0.000	0.893	0.000
PANEL B: Unconstrained model						
$Freq/\bar{F}$	−10.49***	−30.07***	−6.838***	−12.26***	−37.23***	−7.846***
	(1.576)	(4.917)	(0.899)	(1.949)	(4.662)	(0.996)
$Kurt/\bar{K}$	3.827***	0.907	3.581***	2.220*	5.626	4.222**
	(1.305)	(4.330)	(1.346)	(1.268)	(3.669)	(2.023)
Constant	−9.773***	16.87**	−18.99***	−5.295**	11.19**	−19.08***
	(2.021)	(7.010)	(1.942)	(2.465)	(5.600)	(2.871)
R^2	0.646	0.532	0.374	0.644	0.757	0.384
P-val $\beta_f = -\beta_k$	0.001	0.000	0.017	0.000	0.000	0.083
Observations	134	134	134	111	111	111

Notes: The table reports OLS results of the constrained model (equation (12)) for CPI products relating product-specific $CIR_T^{P_j}$ (calculated for the horizon $T = 36$ months and using $\delta = -1$) to the ratio $Kurt/freq$ and OLS results of the unconstrained model (equation (13)) relating product-specific $CIR_T^{P_j}$ to the ratio of the product-level frequency over its average $Freq/\bar{F}$ and the ratio of the product-level kurtosis over its average $Kurt/\bar{K}$. In Case 1, we have removed goods of three broad sectors where sales concentrate (COICOP01.1 Food, COICOP03 Clothing/Footwear, and COICOP05 Furnishing goods). In Case 2, we have removed products for which the share of sales and promotions represent more than 11% of all price changes (this threshold corresponds to the median of this ratio over all CPI products). Product-fixed effects are not included. Robust standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

due to sales (this fraction corresponds to the median value among all CPI products). Results are reported in Table 3. In all specifications, the coefficient associated with the $Kurt/Freq$ ratio is close to the one obtained in our baseline case including sectoral fixed effects. In all except one specification, we find that this coefficient is significant. Overall, these results show that the sufficient statistic predictions emerge more clearly for the CPI products less affected by sales.

The model also offers predictions for the magnitude of the coefficients. In the constrained version of the model, β is predicted to be $-\frac{\delta}{6}$ which is $1/6 \approx 0.167$ since we normalize the shock to $\delta = -1$. The order of magnitude of the estimates for PPI in Table 2 is broadly in line with the theory for the specifications with the long-run restrictions. The last row of each panel in Table 2 shows that we cannot reject the hypothesis that $\beta = 1/6$.[27] For the CPI, instead, the hypothesis that $\beta = 1/6$ is rejected for most specifications. When we consider the products less affected by sales (Table 3), we cannot reject that the coefficient is consistent with the theoretical value in the specifications without long-run restriction.

27. Table A.2 in Appendix D reports more results on formal tests showing that for PPI we cannot reject that the constant is also equal to −36 as predicted by the theory ($\alpha/(-\delta) = -T = -36$).

Finally, we comment on the effect of measurement error on the regression's left and right hand side variables. We quantify the errors in variables using a bootstrap method, following ideas developed by Deaton (1985), see Appendix E. We find that the measurement error in the $Kurt/Freq$ ratio, denoted by ν_{kf} in Section 3, has a much smaller variance than the variance of the regressor itself ($\frac{Kurt}{Freq}$). This implies that the attenuation bias is very small.[28] Instead, we find that the estimated variance of the measurement error in the CIR^P is much bigger than the variance of $\frac{\delta}{6}\widehat{\frac{Kurt_j}{Freq_j}}$ for both PPI and CPI. Such a large measurement error in the dependent variable implies that we should expect R^2 smaller than 0.10 in the OLS regressions.[29] This highlights the pitfall of using the magnitude of the regression's R^2, a procedure followed by Hong *et al*. (2023), as a test of the theory. It is the sign, size, and significance of the coefficients that allow one to test the sufficient statistic proposition and to assess the "information content" of different variables.

5.2. *Estimates of the "unconstrained" empirical specification*

To further investigate the relevance of both the kurtosis and the frequency of price adjustments in explaining the propagation of monetary shocks, Table 4 reports estimation results of the "unconstrained" version of the regression (equation (13)) allowing for a potential different effect of frequency and kurtosis.

For PPI products (Panel A), the estimates are consistent with the theoretical predictions in all specifications. First, after a contractionary shock, if prices are more flexible in a given sector (*i.e.* larger frequency), prices will decline faster and the product-level CIR^P will be more negative. This will induce a negative relationship between the frequency and CIR^P. Second, a smaller kurtosis in a given sector (*i.e.* a larger selection effect) is associated with a more negative reaction of prices after a contractionary shock, resulting in a positive coefficient in the cross-section regression between CIR^P and kurtosis. When not including sectoral fixed effects, coefficients associated with frequency but also kurtosis are all significant at 5% or 10% levels. In the specification including sectoral fixed effects, the sign and the size of the coefficients remain quite similar but they are not significant any more for both frequency and kurtosis. Our interpretation is that the addition of sectoral fixed effects substantially reduces the source of cross-sectional variation and so, lowers the precision of the estimates.

For CPI products (Panel B), we also find—in all cases—a negative and significant relationship across sectors between frequency and the CIR^P, and that the slope coefficient associated with kurtosis, when statistically different from zero, has the expected positive sign. The coefficient associated with frequency is statistically significant in all specifications. When considering CPI products which are less affected by sales (Table 3—Panel B), in all of the specifications the frequency is significantly and negatively correlated with CIR^P. In half of the specifications the kurtosis is significantly and positively correlated with CIR^P.

28. See Table A.16 in Appendix E. For CPI, we find that $Var(\widehat{\frac{Kurt}{Freq}}) \approx 6,900$ and $Var(\nu_{kf}) \approx 390$, implying an attenuation bias for the OLS coefficient in the order of 5% of the coefficient. Similarly for the PPI we find $Var(\widehat{\frac{Kurt}{Freq}}) \approx 800$ and $Var(\nu_{kf}) \approx 40$.

29. In the CPI case, the variance of $\nu_{cir,j}$ is 2,600 and the variance of $\frac{\delta}{6}\widehat{\frac{Kurt_j}{Freq_j}}$ is equal to 200 (*i.e.* $\approx (\frac{1}{6})^2 \times 6,900$). For the PPI case, the variance of $\nu_{cir,j}$ is 1,900 and the variance of $\frac{\delta}{6}\widehat{\frac{Kurt_j}{Freq_j}}$ is equal to 20.

TABLE 4
Regression results—"unconstrained" specification—36-month horizon

Identification	Cholesky		Cholesky		High-Freq. IV	
Long-run Restriction	Yes		No		Yes	
Product FE	No	Yes	No	Yes	No	Yes
PANEL A: PRODUCER PRICES						
$Freq/\bar{F}$	-7.366^{**}	-3.699	-22.35^{*}	-8.245	-6.241^{**}	-3.610
	(3.136)	(2.432)	(11.50)	(10.15)	(2.827)	(2.575)
$Kurt/\bar{K}$	8.864^{**}	5.788	24.19^{*}	21.33	7.065^{**}	3.930
	(4.257)	(3.739)	(14.53)	(13.60)	(3.300)	(3.305)
Constant	-20.45^{***}	-15.93^{***}	-20.79	-38.60^{***}	-26.68^{***}	-23.04^{***}
	(4.488)	(3.867)	(15.44)	(13.13)	(3.234)	(6.445)
Observations	118	118	118	118	118	118
R^2	0.211	0.553	0.164	0.483	0.205	0.462
P-val $\beta_f = -\beta_k$	0.720	0.621	0.897	0.383	0.790	0.916
PANEL B: CONSUMER PRICES						
$Freq/\bar{F}$	-7.188^{**}	-11.95^{***}	-23.46^{***}	-30.86^{***}	-4.931^{***}	-6.566^{***}
	(2.800)	(1.493)	(7.426)	(5.921)	(1.392)	(0.922)
$Kurt/\bar{K}$	4.805^{***}	3.113^{*}	2.142	-3.531	3.572^{***}	2.760^{**}
	(1.700)	(1.591)	(3.672)	(3.549)	(1.159)	(1.095)
Constant	-14.24^{***}	5.506^{*}	4.696	34.43^{***}	-21.30^{***}	-12.76^{***}
	(3.570)	(3.129)	(8.494)	(10.75)	(2.074)	(1.845)
Observations	223	223	223	223	223	223
R^2	0.217	0.725	0.260	0.579	0.177	0.794
P-val $\beta_f = -\beta_k$	0.497	0.000	0.013	0.000	0.446	0.006

Notes: This table reports results of OLS regressions (equation (13)) where the dependent variable is the product-specific $CIR_T^{P_j}$ (calculated for the horizon $T=36$ months, and using $\delta=-1$) and the right-hand-side variables are the ratio of the product-level frequency over its average $Freq/\bar{F}$ and the ratio of the product-level kurtosis over its average $Kurt/\bar{K}$. Product fixed effects are defined at the two-digit level for both CPI and PPI products (*i.e.* 38 product fixed effects for the CPI, and 24 in the case of the PPI). Robust standard errors are reported in parentheses. $^{***}\ p<0.01$, $^{**}p<0.05$, $^{*}\ p<0.1$.

The last row of Table 4 reports a formal test for the hypothesis $\beta_f = -\beta_k$, predicted by the theory. For PPI, we cannot reject the hypothesis in any regression. For CPI, the hypothesis cannot be rejected in the specifications with long-run restrictions without sectoral fixed effects (Columns 1 and 5).[30]

For completeness, we also run regressions where frequency and kurtosis are introduced alone as regressors (see Table 5), although we notice that such a specification is inconsistent with the theory. The results show that the coefficients of these regressions are statistically significant and have the expected sign suggested by the theory. That frequency and kurtosis turn out as significant in separate single-regressor specifications (arguably mis-specified) reflects the weak correlation between frequency and kurtosis in our sample.

We notice that the measurement error associated with the frequency, $\nu_{freq,j}$, is much smaller than the one associated with kurtosis, $\nu_{kurt,j}$. Quantitatively, $\frac{\text{Var}(\nu_{kurt,j})}{\bar{K}^2}$ is one order of magnitude

30. Table A.2 in the Appendix reports p-values of formal Fisher tests for the estimated parameters. For PPI products we cannot reject the hypothesis that the *size* of the coefficients is consistent with model's predictions. The result is less clear for CPI products.

TABLE 5
Baseline OLS regression results: kurtosis alone—frequency alone

Identification Long-run Restriction	Cholesky Yes		Cholesky No		High-Freq. IV Yes	
Product FE	No	Yes	No	Yes	No	Yes
PANEL A: Producer Prices—Kurtosis alone						
$Kurt/\bar{K}$	9.148**	5.240	25.05	20.11	7.305*	3.396
	(4.435)	(3.717)	(15.44)	(13.07)	(3.761)	(3.374)
Constant	−28.10***	−17.99***	−44.00**	−43.18***	−33.17***	−25.05***
	(5.836)	(3.412)	(20.46)	(12.98)	(4.780)	(5.934)
R^2	0.031	0.521	0.020	0.469	0.027	0.421
PANEL B: Producer Prices—Freq. alone						
$Freq/\bar{F}$	−7.404**	−3.613	−22.45*	−7.926	−6.272**	−3.551
	(3.242)	(2.463)	(11.83)	(10.28)	(2.873)	(2.537)
Constant	−11.55***	−10.79***	3.500	−19.65***	−19.59***	−19.55***
	(2.463)	(2.717)	(9.020)	(7.224)	(2.181)	(6.535)
R^2	0.182	0.545	0.146	0.474	0.180	0.457
Observations	118	118	118	118	118	118
PANEL C: Consumer Prices—Kurtosis alone						
$Kurt/\bar{K}$	5.728***	5.177**	5.155	1.801	4.205***	3.894**
	(1.901)	(2.163)	(3.457)	(3.800)	(1.305)	(1.563)
Constant	−22.35***	−15.06***	−21.78***	−18.70***	−26.86***	−24.07***
	(2.513)	(2.318)	(5.596)	(5.598)	(1.939)	(1.615)
R^2	0.039	0.441	0.004	0.322	0.036	0.649
PANEL D: Consumer Prices—Freq. alone						
$Freq/\bar{F}$	−7.402***	−12.16***	−23.56***	−30.63***	−5.090***	−6.748***
	(2.759)	(1.422)	(7.384)	(5.963)	(1.379)	(0.900)
Constant	−9.222***	8.611***	6.933	30.90***	−17.56***	−10.01***
	(2.655)	(2.391)	(6.809)	(9.554)	(1.367)	(1.462)
R^2	0.189	0.718	0.260	0.577	0.151	0.784
Observations	223	223	223	223	223	223

Notes: This table reports OLS results of a model (equation (12)) relating product-specific $CIR_T^{P_j}$ to the ratio of the product-level kurtosis over its average $Kurt/\bar{K}$ and OLS results of a model (equation (13)) relating product-specific $CIR_T^{P_j}$ to the ratio of the product-level frequency over its average $Freq/\bar{F}$. Product fixed effects are defined at the two-digit level for both CPI and PPI products (*i.e.* 38 product fixed effects for the CPI, and 24 in the case of the PPI). Robust standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

larger than $\frac{\mathrm{Var}(\nu_{freq,j})}{\bar{F}^2}$ for both PPI and CPI.[31] This implies that the coefficient for frequency is estimated more precisely than the one for kurtosis (see Appendix E).

5.3. *"Placebo" tests*

While the above results are consistent with the "sufficient statistic" property, a sufficient statistic property predicts something broader: it implies that the effect of a monetary shock should be related to the ratio "kurtosis over frequency" but also that other moments of the price distribution should not matter in this relationship. To test this prediction, we estimate equation (12) adding three additional moments of the price change distribution computed at the product level:

31. Respectively 0.0123 versus 0.0035 for CPI and 0.0125 versus 0.0016 for PPI (see Table A.16 in the Appendix).

TABLE 6

Regression results—placebo specification—36-month horizon

Identification	Cholesky		Cholesky		High-Freq. IV	
Long-run Restriction	Yes		No		Yes	
Product FE	No	Yes	No	Yes	No	Yes
PANEL A: PRODUCER PRICES						
Kurt/Freq	0.201**	0.110	0.513	0.222	0.188**	0.0987
	(0.0960)	(0.0782)	(0.312)	(0.275)	(0.0768)	(0.0666)
Mean	−1.430	−1.350	−10.90	−4.844	−1.262	1.273
	(1.771)	(1.694)	(6.644)	(7.412)	(1.522)	(1.677)
Skewness	−1.869	−3.824	−16.72	−6.397	−2.207	−3.327
	(4.613)	(4.452)	(17.15)	(20.29)	(2.718)	(3.625)
SD	−1.355	0.445	−6.436	−4.213	−0.182	2.197
	(2.509)	(2.335)	(8.512)	(8.616)	(2.011)	(1.728)
Constant	−21.45**	−19.54	−10.56	−11.56	−32.89***	−38.56***
	(9.551)	(12.41)	(30.11)	(44.33)	(7.834)	(11.33)
Observations	118	118	118	118	118	118
R^2	0.100	0.537	0.075	0.471	0.113	0.465
PANEL B: CONSUMER PRICES						
Kurt/Freq	−0.0196	0.0227	0.0897	0.161***	0.0126	0.0308**
	(0.0221)	(0.0217)	(0.0567)	(0.0617)	(0.0140)	(0.0148)
Mean	1.434**	1.818***	−0.901	0.398	0.136	0.794*
	(0.575)	(0.688)	(1.906)	(1.993)	(0.452)	(0.453)
Skewness	6.473**	4.252	19.08**	5.377	4.488**	4.153**
	(2.829)	(3.088)	(8.986)	(9.945)	(1.782)	(1.929)
SD	−0.840	−0.0583	3.104	5.089**	−0.524	0.260
	(0.712)	(0.864)	(1.881)	(2.180)	(0.524)	(0.508)
Constant	−8.567	−11.07*	−42.41**	−59.90***	−18.85***	−23.23***
	(7.039)	(6.666)	(18.80)	(18.17)	(4.344)	(3.922)
Observations	223	223	223	223	223	223
R^2	0.042	0.458	0.051	0.368	0.030	0.660

Notes: This table reports results of OLS regressions (equation (12)) where the dependent variable is the product-specific $CIR_T^{P_j}$ (calculated for the horizon $T = 36$ months and using $\delta = -1$) and the right-hand-side variables include the product-specific ratio $Kurt/freq$ and three other moments of the product-specific price change distribution: the average price change *Mean*, the skewness of price changes *Skewness*, and the standard deviation of price changes *SD*. Product fixed effects are defined at the two-digit level for both CPI and PPI products (*i.e.* 38 product fixed effects for the CPI, and 24 in the case of the PPI). Robust standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

the average size of (non-zero) price changes, the standard deviation, and the skewness of price adjustments. This exercise can be considered as a “placebo” test of our baseline regressions, testing that our main result is not driven by correlations between frequency or kurtosis and other moments of the price change distribution.

Table 6 provides results for this specification. For PPI products (Panel A), the coefficients associated with the ratio of kurtosis over frequency are highly similar to the ones obtained in the baseline case (Table 2). They are much less precisely estimated however, and they remain significant at the 5% level only in two specifications without sectoral fixed effects. Importantly, neither the average size of price changes, nor the standard deviation of prices changes, nor the skewness of price changes, do have statistically significant effects in any of the six specification. These results are consistent with the theoretical prediction. We have in addition estimated an unconstrained version of the “placebo” regression (results are in Table A.3 in the Appendix). Results for PPI products are broadly robust, although the degree of significance decreases, presumably owing to multi-collinearity.

For CPI products (Panel B), results are more mixed. The coefficient of $Kurt/Freq$ is positive and significant in only two cases and several coefficients associated with the "placebo" moments are significant (in 8 cases out of 18). Results are also quite mixed when looking at the unconstrained version of the "placebo" regressions (see Table A.3 in the Appendix). We have also run the same placebo regressions for CPI products less affected by sales (Table A.4 in the Appendix). The coefficient associated with the $Kurt/Freq$ ratio remains positive and significant in all specifications and only three coefficients (over 18) associated with the "placebo" moments are significant at the 10% level.

As alternative "placebo" tests, we have also considered introducing other covariates that may be confounding factors. Candidates are average inflation, production volatility (available for PPI only), or the degree of "upstreamness" in the production chain (as captured by dummies for broad sectors). Results are reported in Table A.12 and Table A.13 in the Appendix. Overall, results are unaffected: the $Kurt/Freq$ ratio remains significant. One qualification, though, is that the variable average inflation turns out to be significant in some cases.

6. ROBUSTNESS ANALYSIS

This section explores the robustness of our findings with respect to several dimensions: (i) the time horizon of the CIR; (ii) the measurement of kurtosis; (iii) the exclusion of products with a large drift in prices; (iv) the use of the interest rate on a long-term bond as a policy indicator (related to the effective lower bound on interest rates and unconventional policies); (v) using moments of price durations as an alternative sufficient statistic; and (vi) using the CIR of output as a dependent variable.

6.1. *The time horizon of the CIR*

In our baseline results, the CIR^P is computed by cumulating the price deviation for 36 months after the shock. We have carried out various estimations using the alternative time horizon of $T = 24$ and $T = 48$ months. The results are reported in the Appendix in Table A.5, Table A.6, Table A.7, and Table A.8. For the 48-month time horizon, the slope coefficients associated with the ratio $Kurt/Freq$ are almost identical to the ones obtained for the 36-month horizon. When we use CIR^P calculated over a 24-month horizon they are lower but still close. As expected, estimates of the intercepts vary with the time horizon. In all regressions, results are quantitatively and qualitatively close to the baseline results.

6.2. *The measurement of kurtosis*

The measurement of kurtosis is known to be severely affected by unobserved heterogeneity. We run robustness regressions using a measure of kurtosis, based on Alvarez *et al*. (2022), that takes into account product-level unobserved heterogeneity. Results reported in the Appendix Table A.9 are very much in line with the ones in our baseline regressions. For PPI, the coefficient associated with the $Kurt/Freq$ ratio is positive, and significant in all specifications, whereas for CPI the estimated coefficients are not statistically different from 0. In the unconstrained regression, results are also qualitatively and quantitatively similar to the ones obtained in the baseline regressions.

We also investigated the role of very large or very small price changes (in absolute values) for the measurement of kurtosis. In the baseline regressions, we have used kurtosis measures calculated on the sample of price changes smaller in absolute value than 15% for PPI price changes and than 25% for CPI price changes (*i.e.* 5% of all price changes in both cases) and we have excluded price changes below 0.1% in both cases. We also tested the robustness of our results to modifying the thresholds defining extreme price changes (results available from the authors upon request). In the first exercise, we investigated the role of large price changes and we set the thresholds defining extreme values to 25% for PPI price changes and 35% for CPI price changes (*i.e.* about 2% of all price changes). In a second exercise, we set the threshold for small price changes to 0.5% (which corresponds to about 5% of all price changes).[32] The results overall remain in line with the baseline results. The standard errors of coefficients however are higher, lowering the significance of the estimated coefficients, in particular for large producer price changes.

6.3. *Using a long-term yield as policy indicator*

In this robustness exercise, we consider an alternative to the policy rate used in the FAVAR estimation. We identify the shock using an external instrument approach, as in *e.g.* Jarocinski and Karadi (2020). The main motivation is that over the last part of our sample, the short-run policy rate was arguably constrained by the proximity of the effective lower bound for interest rates, and the ECB engaged in unconventional monetary policies intended to influence long-term interest rates.[33] We use the 2-year German sovereign bond rate, a relevant risk-free long-term interest rate, instead of the 3-month Euribor rate.[34]

Results relating to the sectoral CIR^P obtained from this FAVAR model and the sufficient statistic, for PPI products, are in line with the baseline (see Table A.10 in the Appendix). The coefficient associated with the $Kurt/Freq$ ratio is positive and significantly different from zero (and we cannot reject the coefficient being equal to the predicted value of $1/6$). In the unconstrained specification, the estimated parameters associated with the frequency and kurtosis are very close to the ones obtained in the baseline case. For CPI products, the coefficient associated with the $Kurt/Freq$ ratio is positive and significant, but much smaller than in the baseline case (Table 2). In the unconstrained version of the model, the coefficient of frequency is negative and significant, as in the baseline, while the coefficient of kurtosis is not significant.

6.4. *Removing products with sizeable drifts in price levels*

The theoretical predictions of the model are derived under the assumption of low inflation. While this assumption is clearly fulfilled for the aggregate inflation rate in France over our sample period, a concern is that for some specific sectors, it may not be the case. Table 1 provides statistics on the average product-specific inflation rates in absolute values. Product-level inflation rates (in absolute value) are typically small as well: average and median inflation rates are about 1.5% per year, whereas the third quartiles of the inflation distribution are around 2%. In this robustness exercise, we remove all products with a "non-small" average inflation rate (in absolute value). In practice, we define small inflation rates as products with an average annual inflation lower than 5% in absolute value.[35] For PPI products, only two products are removed, whereas for CPI, nine

32. Other definitions of small and large price changes lead to similar conclusions.

33. Note however that the policy rate was negative from 2014, and statements by the ECB indicate that the lower bound was not actually reached afterwards.

34. Jarocinski and Karadi (2020) use the 1-year and 2-year German bond as a policy variable in their analysis of ECB monetary policy.

35. See Gagnon (2009), Nakamura *et al*. (2018) or Alvarez *et al*. (2019) for evidence on price rigidity in higher inflation environments. These authors show that when inflation is below 5%, the frequency of price changes does not vary with the inflation rate.

products are removed. For both PPI and CPI, results are reported in Appendix Table A.11 and they are very consistent with the ones obtained in the baseline regressions.[36]

6.5. *Using the CIR of output*

Originally, the theoretical results were developed using CIR^Y, but deriving the predictions for CIR^P is straightforward as shown above. We focused our empirical analysis on CIR^P for two reasons. Firstly, product-level measures of output are only available at an infra-annual frequency for producer goods, and not for consumer goods. Secondly, the sufficient statistic prediction derived for CIR^Y contains a "nuisance parameter", the industry-specific elasticity (ϵ_j), which is not the case for prices where the prediction simply links CIR^P to the kurtosis over frequency ratio. This extra parameter in the prediction for output could blur the quantitative interpretation of the estimated coefficients and might also complicate the estimation of the correlation between CIR^P and the kurtosis-over-frequency ratio.

However, as a robustness exercise, we have performed estimations for output in the case of PPI, using the sectoral Industrial Production Index as product-specific output variable. The shock is normalized the same way as for producer prices.[37] Results of regressions using the CIR^Y as the dependent variable are reported in Table A.14 of the Appendix. Note that the $Kurt/Freq$ ratio is now expected to have a negative sign, opposite to the case of CIR^P. The results turn out to be mixed and generally weaker than using CIR^P. In all cases (whether with "long-run restriction" or not, and with fixed effects or not) the $Kurt/Freq$ has the expected sign. However, it is not significant, reflecting very imprecise estimates (in particular with fixed effects). Our interpretation is that the weaker results are consistent, and in fact to be expected, in the presence of heterogeneity in the industry-specific income elasticity (ϵ_j).

6.6. *Using moments from the distribution of price durations*

Another robustness test consists in using moments of price durations as a substitute for the candidate sufficient statistic $Kurt/Freq$. Indeed, several authors have exploited the mapping between the duration of price spells and the size of price adjustments, showing that the distribution of the durations is informative about monetary non-neutrality, see *e.g.* Carvalho and Schwartzman (2015) and Baley and Blanco (2021). As shown by Proposition 2 in Alvarez, Lippi *et al*. (2016), with Brownian shocks the distribution of durations provides an alternative formula to compute the CIR, involving two moments: the average price-spell duration and the squared coefficient of variation of durations. The former is obviously related to the average frequency of price changes appearing in equation (7), the latter is a stand-in for the kurtosis of the size of price changes.

Regression results involving these spell duration moments are reported in the Appendix (Table A.15). The specification suggested by Alvarez, Lippi *et al*. (2016), where duration and the coefficient of variation (CV) enter the regression multiplicatively and appear significant with the expected sign, both for the PPI as well as for the CPI sample. As mentioned above, this result is consistent with the ones shown above using the moments from the distribution of price changes.[38] The regressions where the frequency and the CV are entered as separate regressors

36. Using a threshold of 4% to define "small" versus "large" inflation rates leads to similar results.

37. The number of products is larger than in the case of prices because more product-level IPIs (than PPIs) are available over a long-time dimension.

38. The correlation coefficient between $Kurt/Freq$ and $E(d)(1+CV(d)^2)$ calculated across products is equal to 0.64 for PPI and 0.79 for CPI which supports the claim that these two candidate sufficient statistics encode similar information.

also show the correct sign but the statistical significance of the CV regressor is weaker. We note that in several models, such as the ones discussed in Section 2, there is a tight link between the distribution of durations and the distribution of price changes. The two distributions encode the same information (see Appendix E in Alvarez *et al.* 2022 for a formal analysis of this equivalence). Under the null hypothesis that the model is the data-generating process, the two tests are equivalent. Differences in the statistical significance of the regressions might be due to differences in the quality of the data (measurement errors in durations versus size of price changes) or reflect deviations from the assumed normal distribution for the firm's idiosyncratic shocks.

7. CONCLUSION

In a broad class of sticky price models, the non-neutrality of nominal shocks is captured by a simple sufficient statistic: the ratio of the kurtosis of the price change distribution over the frequency of price changes, see *e.g.* Alvarez *et al.* (2022). We tested this theoretical prediction using sectoral and microeconomic data for France both for PPI and CPI products. Our test followed three steps. We first measured the effects of monetary shocks using a FAVAR across a number of industries using data from 2005 to 2019. Secondly, we measured the candidate sufficient statistics using microdata for the same industries. Thirdly, we used these estimates to test the sufficient statistical predictions.

We found clear support for the theoretical predictions, particularly in the PPI data. The estimated industry non-neutrality correlates with the kurtosis and the frequency, in a way that is consistent with the theory. Several robustness tests are investigated and the results appear solid. The support for the theoretical predictions is weaker on the CPI data. This might be due to seasonal sales (or price plans). Such features, prevalent in the CPI, violate the assumptions under which the sufficient-statistic result holds. Another possible confounding factor is the presence of learning (price discovery), a feature shown by Baley and Blanco (2019) to weaken the power of the sufficient statistic, and which is likely more prevalent for CPI. When goods with a large prevalence of sales are removed from the CPI sample it becomes harder to reject the predictions of the theory. Avenues for future research involve investigating whether the sufficient statistical proposition holds in other datasets, as recently done by Gautier *et al.* (2023) for the gasoline industry.

Acknowledgments. The views in this paper are those of the authors and do not necessarily reflect those of the Banque de France or the Eurosystem. We thank the Editor and three anonymous referees for helpful suggestions. We thank John Leahy for pushing us to pursue this question. We are indebted with our discussants Luca Dedola and Klaus Adam and are grateful to seminar participants at the ECB, the University of Wisconsin, the Bank of Italy, EIEF, the University of Michigan, the University of Bielefeld, the Imperial College, the University of Edinburgh, the BIS, the IMF, the Banque de France, the Banco de España, the Boston Fed, the 2021 CEBRA conference, the 2021 Minnesota Workshop in Macroeconomic Theory, the 2021 Hydra Conference. Laurent Baudry, Andrej Mijakovic, and Laura Pittalis provided excellent research assistance. We also thank Insee (the French Statistical Office) for providing us with micro price datasets; the access to these datasets has been made possible within a secure environment offered by CASD—Centre d'Accès Sécurisé Distant (Ref. 10.34724/CASD). Lippi acknowledges financial support from the ERC grant 101054421-DCS.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability Statement

The data and code underlying this research are available on Zenodo at https://doi.org/10.5281/zenodo.11482758. The replication package includes product-level moments of the price change distributions but not the micro CPI and PPI datasets used to compute these moments which are proprietary data provided by Insee (French Statistical Office). The replication package documents how we had access to these micro price datasets.

REFERENCES

ALTAVILLA, C., BRUGNOLINI, L., GURKAYNAK, R. S., *et al.* (2019), “Measuring Euro Area Monetary Policy”, *Journal of Monetary Economics*, **108**, 162–179.

ALVAREZ, F., BERAJA, M., GONZALEZ-ROZADA, M., *et al.* (2019), “From Hyperinflation to Stable Prices: Argentinas Evidence on Menu Cost Models”, *The Quarterly Journal of Economics*, **134**, 451–505.

ALVAREZ, F. and LIPPI, F. (2020), “Temporary Price Changes, Inflation Regimes, and the Propagation of Monetary Shocks”, *American Economic Journal: Macroeconomics*, **12**, 104–152.

ALVAREZ, F., LIPPI, F. and SOUGANIDIS, P. (2023), “Price Setting with Strategic Complementarities as a Mean Field Game”, *Econometrica*, **91**, 2005–2039.

ALVAREZ, F. E., LE BIHAN, H. and LIPPI, F. (2016), “The Real Effects of Monetary Shocks in Sticky Price Models: A Sufficient Statistic Approach”, *The American Economic Review*, **106**, 2817–2851.

ALVAREZ, F. E. and LIPPI, F. (2014), “Price Setting with Menu Costs for Multi Product Firms”, *Econometrica*, **82**, 89–135.

ALVAREZ, F. E., LIPPI, F. and OSKOLKOV, A. (2022), “The Macroeconomics of Sticky Prices with Generalized Hazard Functions”, *The Quarterly Journal of Economics*, **137**, 989–1038.

ALVAREZ, F. E., LIPPI, F. and PACIELLO, L. (2016), “Monetary Shocks in Models with Inattentive Producers”, *The Review of Economic Studies*, **83**, 421–459.

ANDRADE, P. and FERRONI, F. (2021), “Delphic and Odyssean Monetary Policy Shocks: Evidence from the Euro Area”, *Journal of Monetary Economics*, **117**, 816–832.

AUER, R., BURSTEIN, A. and LEIN, S. M. (2021), “Exchange Rates and Prices: Evidence from the 2015 Swiss Franc Appreciation”, *American Economic Review*, **111**, 652–686.

BALEY, I. and BLANCO, A. (2019), “Firm Uncertainty Cycles and the Propagation of Nominal Shocks”, *American Economic Journal: Macroeconomics*, **11**, 276–337.

—— —— (2021), “Aggregate Dynamics in Lumpy Economies”, *Econometrica*, **89**, 1235–1264.

BAUDRY, L., LE BIHAN, H., SEVESTRE, P., *et al.* (2007), “What Do Thirteen Million Price Records Have to Say about Consumer Price Rigidity?”, *Oxford Bulletin of Economics and Statistics*, **69**, 139–183.

BERARDI, N. and GAUTIER, E. (2016), “Adjustments in Consumer Prices in France in Periods of Low Inflation”, *BdF Bulletin, Quarterly Selection of Articles*, **41**, 51–65.

BERARDI, N., GAUTIER, E. and LE BIHAN, H. (2015), “More Facts about Prices: France Before and During the Great Recession”, *Journal of Money, Credit and Banking*, **47**, 1465–1502.

BERNANKE, B. S., BOIVIN, J. and ELIASZ, P. (2005), “Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach”, *The Quarterly Journal of Economics*, **120**, 387–422.

BILS, M. and KLENOW, P. J. (2004), “Some Evidence on the Importance of Sticky Prices”, *Journal of Political Economy*, **112**, 947–985.

BOIVIN, J., GIANNONI, M. P. and MIHOV, I. (2009), “Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data”, *American Economic Review*, **99**, 350–384.

BONOMO, M., CARVALHO, C., KRYVTSOV, O., *et al.* (2023), “Multi-Product Pricing: Theory and Evidence from Large Retailers in Israel”, *The Economic Journal*, **133**, 905–927.

BOPPART, T., KRUSELL, P. and MITMAN, K. (2018), “Exploiting MIT Shocks in Heterogeneous-Agent Economies: The Impulse Response as a Numerical Derivative”, *Journal of Economic Dynamics and Control*, **89**, 68–92.

BOUAKEZ, H., CARDIA, E. and RUGE-MURCIA, F. J. (2009), “The Transmission of Monetary Policy in a Multisector Economy”, *International Economic Review*, **50**, 1243–1266.

BURSTEIN, A., EICHENBAUM, M. and REBELO, S. (2005), “Large Devaluations and the Real Exchange Rate”, *Journal of Political Economy*, **113**, 742–784.

CABALLERO, R. J. and ENGEL, E. M. R. A. (1993), “Heterogeneity and Output Fluctuations in a Dynamic Menu-Cost Economy”, *The Review of Economic Studies*, **60**, 95.

—— —— (1999), “Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S, s) Approach”, *Econometrica*, **67**, 783–826.

—— —— (2007), “Price Stickiness in Ss Models: New Interpretations of Old Results”, *Journal of Monetary Economics*, **54**, 100–121.

CABALLERO, R. J., ENGEL, E. M. R. A. and HALTIWANGER, J. (1997), “Aggregate Employment Dynamics: Building from Microeconomic Evidence”, *The American Economic Review*, **87**, 115–137.

CALVO, G. A. (1983), “Staggered Prices in a Utility-Maximizing Framework”, *Journal of Monetary Economics*, **12**, 383–398.

CAPLIN, A. and LEAHY, J. (1997), “Aggregation and Optimization with State-Dependent Pricing”, *Econometrica*, **65**, 601–626.

CARVALHO, C. (2006), “Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks”, *The B.E. Journal of Macroeconomics*, **6**, 1–58.

CARVALHO, C., LEE, J. W. and PARK, W. Y. (2021), “Sectoral Price Facts in a Sticky-Price Model”, *American Economic Journal: Macroeconomics*, **13**, 216–256.

CARVALHO, C. and SCHWARTZMAN, F. (2015), "Selection and Monetary Non-Neutrality in Time-Dependent Pricing Models", *Journal of Monetary Economics*, **76**, 141–156.
CAVALLO, A. (2018), "Scraped Data and Sticky Prices", *The Review of Economics and Statistics*, **100**, 105–119.
DEATON, A. (1985), "Panel Data from Time Series of Cross-Sections", *Journal of Econometrics*, **30**, 109–126.
DEDOLA, L., KRISTØFFERSEN, M. S. and ZÜLLIG, G. (2021), "The Extensive and Intensive Margin of Price Adjustment to Cost Shocks: Evidence from Danish Multiproduct Firms" (Mimeo, European Central Bank).
EICHENBAUM, M., JAIMOVICH, N. and REBELO, S. (2011), "Reference Prices, Costs, and Nominal Rigidities", *American Economic Review*, **101**, 234–262.
EICHENBAUM, M. S., JAIMOVICH, N., REBELO, S., *et al.* (2014), "How Frequent Are Small Price Changes?", *American Economic Journal: Macroeconomics*, **6**, 137–155.
GAGNON, E. (2009), "Price Setting During Low and High Inflation: Evidence from Mexico", *Quarterly Journal of Economics*, **124**, 1221–1263.
GAUTIER, E. (2008), "The Behaviour of Producer Prices: Some Evidence from French PPI Micro Data", *Empirical Economics*, **35**, 301–332.
GAUTIER, E., CONFLITTI, C., FABER, R. P., *et al.* (2024), "New Facts on Consumer Price Rigidity in the Euro Area", *American Economic Journal: Macroeconomics*, **16**, 386–431.
GAUTIER, E. and LE BIHAN, H. (2022), "Shocks vs Menu Costs: Patterns of Price Rigidity in an Estimated Multi-Sector Menu-Cost Model", *The Review of Economics and Statistics*, **104**, 668–685.
GAUTIER, E., MARX, M. and VERTIER, P. (2023), "How Do Gasoline Prices Respond to a Cost Shock?", *Journal of Political Economy: Macroeconomics*, **1**, 707–741.
GERTLER, M. and KARADI, P. (2015), "Monetary Policy Surprises, Credit Costs, and Economic Activity", *American Economic Journal: Macroeconomics*, **7**, 44–76.
GOLOSOV, M. and LUCAS JR, R. E. (2007), "Menu Costs and Phillips Curves", *Journal of Political Economy*, **115**, 171–199.
GOPINATH, G. and ITSKHOKI, O. (2010), "Frequency of Price Adjustment and Pass-Through", *Quarterly Journal of Economics*, **125**, 675–727.
GORODNICHENKO, Y. and WEBER, M. (2016), "Are Sticky Prices Costly? Evidence from the Stock Market", *American Economic Review*, **106**, 165–199.
HONG, G. H., KLEPACZ, M., PASTEN, E., *et al.* (2023), "The Real Effects of Monetary Shocks: Evidence from Micro Pricing Moments", *Journal of Monetary Economics*, **139**, 1–20.
IMBS, J., JONDEAU, E. and PELGRIN, F. (2011), "Sectoral Phillips Curves and the Aggregate Phillips Curve", *Journal of Monetary Economics*, **58**, 328–344.
JAROCINSKI, M. and KARADI, P. (2020), "Deconstructing Monetary Policy Surprises: The Role of Information Shocks", *American Economic Journal: Macroeconomics*, **12**, 1–43.
KEHOE, P. and MIDRIGAN, V. (2015), "Prices Are Sticky after All", *Journal of Monetary Economics*, **75**, 35–53.
LA'O, J. and TAHBAZ-SALEHI, A. (2022), "Optimal Monetary Policy in Production Networks", *Econometrica*, **90**, 1295–1336.
LEAHY, J. (2016), Comment on "Are State and Time Dependent Models Really Different?" (NBER Macroeconomics Annual 2016, Volume 31 Ch. 5. University of Chicago Press) 458–464.
MIDRIGAN, V. (2011), "Menu Costs, Multi-Product Firms, and Aggregate Fluctuations", *Econometrica*, **79**, 1139–1180.
NAKAMURA, E. and STEINSSON, J. (2008), "Five Facts about Prices: A Reevaluation of Menu Cost Models", *Quarterly Journal of Economics*, **123**, 1415–1464.
—— —— (2010), "Monetary Non-Neutrality in a Multisector Menu Cost Model", *The Quarterly Journal of Economics*, **125**, 961–1013.
NAKAMURA, E., STEINSSON, J., SUN, P., *et al.* (2018), "The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation", *The Quarterly Journal of Economics*, **133**, 1933–1980.
REIS, R. (2006), "Inattentive Producers", *The Review of Economic Studies*, **73**, 793–821.
VERMEULEN, P., DIAS, D. A., DOSSCHE, M., *et al.* (2012), "Price Setting in the Euro Area: Some Stylized Facts from Individual Producer Price Data", *Journal of Money, Credit and Banking*, **44**, 1631–1650.
WOODFORD, M. (2009), "Information-Constrained State-Dependent Pricing", *Journal of Monetary Economics*, **56**, s100–s124.