Who Pays Attention? Firm Size, Information, and Monetary Non-Neutrality*

Sanhitha Jugulum

Columbia University

November 2026

Click here for the most recent version†

Abstract

A large literature models information frictions at the firm level as a microfoundation for price rigidity and hence monetary non-neutrality. The literature mostly abstracts from heterogeneity in firms' information by firm characteristics. In this paper, I use an LLM to construct a measure of firms' attention to aggregate nominal conditions from 10-K filings and document a strong size gradient: larger firms within narrowly defined industries are systematically more attentive. I develop a model of heterogeneous firms and endogenous information acquisition to rationalize this gradient. Firm size in the model is a first-order determinant of the profit loss from pricing mistakes and therefore of the incentive to acquire information to avoid these mistakes. I then use the model to establish the conditions under which the heterogeneity in information across firms matters for aggregate price adjustment. When the larger, more attentive firms also have lower partial equilibrium cost pass-through, the resulting negative covariance between attention and pass-through increases aggregate price rigidity and monetary non-neutrality. I take the model to data using a sufficient-statistics approach to quantify the importance of this new channel. The baseline model-implied Phillips curve is about 40% flatter than under a symmetric firm, homogeneous-attention benchmark. About a fifth of this flattening is attributed to the 'negative covariance between attention and pass-through' in the data. I then calibrate the structural model to the U.S. economy in the mid-2010s and obtain a flattening of similar magnitude. Finally, I conduct a counterfactual analysis using the structural model, in which I vary the productivity distribution to alter the degree of market concentration in the economy, while keeping the demand and cost parameters fixed. I show that changes in market structure can have substantially different effects on monetary non-neutrality once the allocation of attention is allowed to respond.

*I thank my advisor, Professor Jennifer La'O, for her guidance throughout the project. I am also very grateful for the extremely insightful feedback and discussions with Professors Hassan Afrouzi and Laura Veldkamp. I am also thankful for the feedback I received from Professor Jesse Schreger on the use of LLMs. I thank all the members of the faculty and participants at the Economic Fluctuations Colloquium, the Econ Lunch Seminar at Columbia Business School, the Money-Macro workshop at Columbia and the junior macro workshop at the NY Fed.

†Frequently updated. Please click here for the most recent version.

A substantial body of literature is built around the idea that information frictions provide a leading micro-foundation for the real effects of nominal shocks. Imperfectly informed firms do not fully incorporate changes in aggregate conditions into their pricing decisions, causing prices to respond sluggishly to nominal shocks. This idea can be traced back to Lucas (1972) and was further developed in Woodford (2001), Mankiw and Reis (2002), and Mackowiak and Wiederholt (2009). Canonical models typically feature ex ante symmetric firms and are therefore largely silent on whether firms' attention to aggregate conditions varies systematically with their characteristics. This paper asks two questions. First, are there systematic differences across firms in how closely they track aggregate nominal conditions? Second, if such heterogeneity does exist, how does it matter for overall monetary non-neutrality?

I document that firms do in fact differ systematically in how well they track aggregate nominal conditions. Using firms' 10-K filings, I construct a granular firm-year measure to proxy for firms' attention to inflation and monetary conditions. I uncover a strong size gradient: larger firms within narrowly defined industries are systematically more attentive. Then I ask why larger firms optimally acquire more information about aggregate nominal conditions. I answer this using a model of heterogeneous monopolistic firms with variable-elasticity demand à la Kimball (1995) and endogenous information acquisition. Firms differ in productivity, generating heterogeneity in their relative size within sectors. Firms' relative size is a first-order determinant of the cost of mispricing: a percentage pricing error generates greater profit losses for a larger firm, giving larger firms a stronger incentive to acquire more precise information about nominal shocks. In the model, the stakes from mispricing scale with size, while the costs of acquiring information do not. I use this model to then show that when larger, more attentive firms have lower partial-equilibrium cost pass-through—the empirically relevant case (Amiti, Itskhoki and Konings, 2019; Burya and Mishra, 2022)—better information is concentrated among firms whose desired prices are less responsive to changes in marginal cost. This dampens the aggregate price response and increases monetary non-neutrality. Quantitatively, relative to a fully symmetric case with homogeneous attention, the model implies a 18 percent higher output response.

The first part of the paper establishes the key empirical fact in more detail. I use the universe of annual 10-K filings by publicly traded U.S. firms between 2005 and 2020 to construct a granular text-based measure of firm attention to inflation and aggregate nominal conditions. Because firms are required to disclose their risks in their 10-K filings, these reports are particularly suited for studying how firms discuss the aggregate nominal environment. Each filing is processed through

a sequence of LLM queries. The procedure first identifies passages that substantively discuss inflation, monetary policy, or broader aggregate nominal conditions. A second step evaluates the confidence expressed in the language of those passages: GPT-5-mini assigns each passage a score that captures how confidently the firm characterizes the nominal environment. I combine these two dimensions into a *confidence-adjusted attention* measure, defined as the share of passages devoted to aggregate nominal conditions weighted by the 'confidence score' of the language used. A firm therefore receives a high score if it both discusses aggregate nominal conditions extensively and does so with relatively high confidence. This measure is quite intuitive. A firm that pays greater attention to the nominal environment should not only discuss it more frequently but also characterize it with greater precision. I also show that the measure, at the aggregate level, broadly tracks movements in inflation and is also related to cross-industry differences in price rigidity.

Using this measure as a proxy for firm attention to aggregate nominal conditions, I document a systematic size gradient. Within narrowly defined industries and years, firms with larger revenue shares exhibit greater attention to inflation and aggregate nominal conditions. The relationship is robust to alternative industry definitions and to alternative attention proxies constructed from the same filings. Thus, the central empirical fact is that there is meaningful heterogeneity in how closely firms track aggregate nominal conditions, and that this heterogeneity is systematically related to firms' relative size.

An advantage of the text-based measure is the granularity at which attention can be measured. Direct surveys of U.S. firms' inflation expectations have historically provided limited information on firm characteristics. More importantly, surveys that track firms' inflation expectations are very few. The very recent Survey of Firms' Inflation Expectations (SoFIE) run by the Federal Reserve Bank of Cleveland provides an important source of firm-level expectations data. However, many firm characteristics are reported in relatively coarse categories. Firm-size bins, in particular, are much coarser than the continuous measures available in accounting data. By contrast, 10-K filings can be linked directly to Compustat, allowing attention to be related to detailed measures of firm size and other characteristics. I nevertheless use various survey evidence as an external check and show that the size gradient in attention is qualitatively consistent with the patterns observed in surveys.

The next section develops a static general equilibrium model of heterogeneous price-setting firms. The economy consists of a continuum of sectors, and within each sector there are a continuum of firms producing differentiated goods. Firms differ permanently in productivity, generat-

ing heterogeneity in their relative size and relative prices within sectors. Within each sector, firms face variable-elasticity demand à la Kimball (1995). Firms face strategic complementarity in pricing. A firm's desired price depends in part on the prices set by other firms in its sector. In this monopolistic economy, all firms compete against a single sectoral price index. As a result, firms with different prices relative to the index face different demand elasticities and partial-equilibrium cost pass-throughs. Importantly, here the partial-equilibrium cost pass-through also summarizes the strength of strategic complementarity in pricing: lower-pass-through firms place greater weight on changes in the sectoral price index when setting their own price. As a result, partial-equilibrium pass-through also orders firms' desired price responses to an aggregate nominal shock in general equilibrium.

Firms do not directly observe the aggregate nominal shock and are rationally inattentive. In addition to setting prices conditional on their information, they optimally choose how much information to acquire about the shock. In this environment, it translates into choosing the precision of their private signal. The information problem has a simple interpretation. Firms acquire information to reduce the expected losses from setting a price different from the one they would choose under full information. These losses depend on firms' steady-state characteristics, and firm size enters directly. A given percentage pricing error is more costly for a larger firm, whereas the cost of acquiring information does not rise in proportion to size. Larger firms, therefore, have stronger incentives to resolve uncertainty and to optimally choose more informative signals. In the Gaussian environment of the model, this translates into a higher Kalman gain, the weight the firm places on new information when updating its beliefs. I characterize the conditions under which this direct size effect dominates other determinants of information acquisition, thereby generating the positive relationship between firm size and attention observed in the data.

Aggregating firms' pricing decisions delivers a key theoretical result. With ex ante symmetric firms, all firms share a common price pass-through, so the identity of the firms receiving more or less precise information is irrelevant for aggregate price adjustment. What matters is average responsiveness to information. Once firms differ in pass-through, however, the allocation of attention becomes consequential. When larger, more attentive firms also have lower-cost pass-through —as is empirically relevant (Amiti, Itskhoki and Konings, 2019; Burya and Mishra, 2022) — attention is concentrated among firms whose desired prices are less responsive to marginal-cost shocks. This negative covariance dampens aggregate price adjustment and increases monetary non-neutrality. More generally, monetary non-neutrality in this class of models is governed by

the joint distribution of firms' steady-state size, demand elasticity, pass-through, and information/Kalman gain. This is a sufficient statistics characterization of aggregate non-neutrality. The next section takes this model to the data using a 'sufficient statistics' approach.

I then take the aggregation result directly to the data. The theory identifies a small set of firm-level steady-state objects that determine aggregate price adjustment: firms' relative size, partial-equilibrium cost pass-through, demand elasticity, and the Kalman gain. I construct empirical counterparts of each of these objects and use their joint distribution to quantify the degree of monetary non-neutrality.

The main measurement challenge is the Kalman gain. The 10-K exercise provides a granular proxy for firm attention, but the text measure is often contaminated by measurement error and lacks an interpretable economic unit. By itself, it does not map into a firm-level Kalman gain. I therefore estimate a state-space model to recover firm-level Kalman gains from the textual attention measure. I assume that the textual attention maps into a latent attention measure that determines the precision of firms' signals on monetary shocks. I use inflation to summarize the dynamics of monetary shocks to discipline the mapping [1]. In particular, the time-series and cross-sectional moments of the 10-K measure, and the aggregate moments from the Survey of Firms' Inflation Expectations are used to discipline the mapping via SMM. Importantly, the Kalman gain is estimated independently, without imposing the structural relationship between firm characteristics and the Kalman gain implied by the model.

Steady-state relative size is measured in Compustat. I compute firms' sales shares within their primary NAICS 6-digit industry for each period. I then average their share over short windows to filter out transitory fluctuations. Then, I need estimates of firm-level elasticity and pass-through. Using Nielsen scanner data, Burya and Mishra (2022) estimate how demand elasticities and partial-equilibrium cost pass-through vary with firms' market shares. I evaluate these estimated schedules at firms' measured steady-state relative sizes to impute elasticities and pass-throughs for firms in my sample. [2] The resulting empirical joint distribution reveals a negative relationship between firm pass-throughs and Kalman gains. In particular, larger firms within an industry have higher Kalman gains but lower cost pass-throughs. Consequently, the covariance

[1] This assumption remains perfectly valid under the static GE model I write down. Learning about inflation and monetary shocks are equivalent. More detailed argument and caveats are needed when I move to a dynamic environment.

[2] The imputation relies on the assumption that firms' market shares in their NAICS-6 industries proxy their market shares at the primary product market level, the level at which Burya and Mishra (2022) analysis is conducted. Many external validations supporting this assumption are presented in the paper in section 3.

between the Kalman gain and pass-through is negative within the average sector. Because the Kalman gains are recovered independently of the pricing characteristics, this negative relationship is not imposed by the structural model. It therefore serves as an independent validation of the model predictions.

I quantify the importance of this systematic heterogeneity in attention using three nested environments. The first is a CES benchmark with symmetric information. All firms have the same unit partial-equilibrium cost pass-through, and the same Kalman gains equal to the average Kalman gain in the data. In this case, aggregate price adjustment depends only on the average Kalman gain. I then introduce the empirical distributions of pass-through and demand elasticity and assign the same Kalman gain to every firm within a sector. This isolates the additional rigidity generated by incomplete and heterogeneous pass-through while shutting down attention allocation across firms. Finally, I restore the empirical distribution of firm-level Kalman gains, allowing attention to be allocated across pricing types as observed in the data.

Relative to the CES benchmark, introducing heterogeneous elasticities and pass-throughs while holding attention homogeneous flattens the Phillips curve by about 33 percent. Restoring the observed allocation of attention yields an additional 7 percentage-point flattening, for a total of about 40 percent in the baseline economy. This translates to an additional 18 percent larger output response for 1% shocks to money supply than under the CES benchmark. The systematic allocation of attention toward lower-pass-through firms provides an economically meaningful additional source of monetary non-neutrality.

The direct quantification takes the observed joint distribution of firm characteristics as given. I next turn to the structural model to study how the distribution changes as the underlying market structure changes. This distinction is important because a change in the productivity distribution does not simply change the distribution of market shares. It simultaneously changes the distributions of firms' elasticities and pass-throughs and, consequently, through endogenous information acquisition, the distribution of information. The structural model allows all of these margins to adjust jointly.

I calibrate a representative-sector version of the model with a parametric Kimball demand system, a flexible productivity distribution, and endogenous information acquisition. The demand and productivity parameters are estimated jointly using the simulated method of moments, targeting within-industry moments of concentration, markups, and cost pass-through observed in the US in the mid 2010s. Conditional on these parameters, I discipline the scale and curvature of

information costs using moments of the empirical Kalman-gain schedule. The model fits the data reasonably well. A key reason for the fit is a new parametric Kimball demand functional form adopted in this paper. This form combines the desirable features of popular functional forms à la Klenow and Willis (2016) and Dotsey and King (2005) to generate an empirically plausible joint distribution of sales share, elasticities and pass-throughs. The extant parametric Kimball demand systems are unable to deliver all three simultaneously.

The calibrated model produces the same quantitative importance of the 'negative covariance' mechanism found in the direct measurement exercise. Repeating the same nested comparison, the calibrated economy generates a Phillips curve that is about 36 percent flatter than under CES, with roughly one-fifth of the flattening attributable to the negative attention–pass-through alignment channel. Thus, the structural model generates quantitatively similar contribution of attention allocation to aggregate price rigidity.

Finally, I use the model to study how changes in market concentration affect monetary non-neutrality. Holding the demand system and information-cost technology fixed, I vary the productivity distribution to construct less concentrated economies. The implied concentration changes are comparable in scale to documented U.S. increases between the 1980s and 2010s. Moving from these less concentrated economies to the baseline reveals three offsetting forces. Higher concentration raises sales-weighted average attention, which, by itself, would make aggregate prices more responsive to nominal shocks. At the same time, the activity shifts toward lower-pass-through firms, strengthening real rigidity. This dampens the increase in average attention effect. Additionally, the dispersion in firm pass-through and size rises, making the attention-pass-through alignment more negative via endogenous information acquisition. This acts as an additional dampener. In the calibrated model, these two negative forces more than offset the increase in average attention, strengthening monetary non-neutrality. The quantitative effect is nevertheless moderate: the Phillips-curve slope falls by approximately 2–6% despite a substantial rise in concentration. This counterfactual highlights the broader implication of the paper. Changes in market structure alter monetary non-neutrality not only through the familiar real-rigidity channel but also by changing how information is acquired and allocated among firms. Accounting for this can change the inference about monetary non-neutrality materially.

Related literature This paper relates to four strands of literature. First, the paper contributes to the literature on *information frictions and nominal rigidity*. A long tradition studies how incomplete information can generate sluggish price adjustment and real effects of nominal shocks; prominent

examples include Lucas (1972), Mankiw and Reis (2002), Woodford (2001), Nimark (2008), Mackowiak and Wiederholt (2009), Angeletos and La'O (2009), Angeletos, Iovino and La'O (2016), and more recently Hellwig and Venkateswaran (2025) and Gagliardone and Tielens (2026). Much of this literature studies the case of symmetric firms. Instead, I study the case of asymmetric firms with endogenous information acquisition. This introduces systematic cross-sectional heterogeneity in firms' information about nominal conditions, and aggregation reveals that *allocation* of information across firms has an additional effect on aggregate price adjustment. The key result is that, once firms differ in their desired price responses, monetary non-neutrality depends not only on average attention but also on how attention covaries with firms' pricing characteristics. This provides a distinct source of aggregate rigidity even holding average information fixed.

Second, the paper relates to the growing empirical literature on firms' macroeconomic expectations. Survey-based work has documented substantial departures of firm expectations from full-information rational expectations and considerable heterogeneity across firms; see, among others, Coibion, Gorodnichenko and Kamdar (2018), Coibion, Gorodnichenko and Kumar (2018), Tanaka et al. (2020), and Candia, Coibion and Gorodnichenko (2024). More recently, textual measures have made it possible to study firm attention over longer panels and for broader samples of firms. Song and Stern (2025) construct measures of attention to macroeconomic news from firms' regulatory filings, while Flynn and Sastry (2024) use regulatory filings and conference calls to measure attention to aggregate conditions. My measurement exercise is complementary to this work. I use an LLM-based classification of firms' 10-K filings to isolate discussion specifically related to inflation, monetary policy, and aggregate nominal cost conditions. The LLM-based measure classifies a broader set of documents as discussing aggregate nominal conditions. The LLMs also allow me to construct a measure that accounts for nuances such as the expressed uncertainty in the disclosure language. The empirical focus is the cross-section: I document a strong gradient in nominal attention with firms' relative size within narrowly defined industries. I then combine the text measure with moments from firms' inflation expectations to map observed textual attention into firm-level Kalman gains. This last step produces a model-relevant measure of belief updating that can be brought directly into the aggregation exercise.

Third, the paper contributes to the literature on *endogenous information acquisition*. Beginning with Sims (2003), rational-inattention models study how agents optimally choose the precision and content of their information. Applications to firm pricing include Mackowiak and Wiederholt (2009), Maćkowiak, Matějka and Wiederholt (2018), Afrouzi and Yang (2021), and Afrouzi

(2024), among others. Related work studies how incentives to acquire information vary with the economic environment or the stakes associated with decision errors; see, for example, Van Nieuwerburgh and Veldkamp (2006), Kohlhas and Walther (2021), and Flynn and Sastry (2024). The distinctive feature here is that information acquisition varies endogenously with *steady-state firm characteristics*. Because the profit loss from mispricing is proportional to firm size, larger firms have more at stake when their prices deviate from their full-information targets and therefore optimally acquire more information about nominal shocks. This mechanism generates the firm-size gradient in attention observed in the data. It differs from Afrouzi (2024), where variation in firms' information arises from their competitive environment and incentives to learn about competitors, and from work on state-dependent attention, where information acquisition varies primarily over the business cycle. Here, the emphasis is instead on persistent cross-sectional differences in information acquisition across heterogeneous firms. This cross-sectional heterogeneity in turn turns out to have a sizable effect on the aggregate price responsiveness to monetary policy.

Finally, the paper relates to the literature on *market structure and the transmission of monetary policy*. Rotemberg and Woodford (1996) emphasize the role of imperfect competition and strategic complementarities for aggregate fluctuations. More recent work studies how oligopolistic market structure and market power affect monetary non-neutrality, including Mongey (2021), Wang and Werning (2022), Afrouzi (2024) and Alexander et al. (forthcoming). Offering a complementary perspective, Bardóczy, Bornstein and Salgado (2026) studies monetary transmission in the presence of labor market power in an oligopsony model. Closest to the monopolistic environment studied here is the work of Baqaee, Farhi and Sangani (2024). They provide related evidence on heterogeneous firm responses and the reallocation induced by monetary shocks. They focus on the reallocation of output from high-elasticity to low-elasticity firms induced by monetary policy shocks due to systematic differences in pass-throughs. In this paper, the channel is embedded, and an additional unexplored channel is highlighted. In a sense, not only does elasticity systematically differ by firm pass-through, but so does the degree of nominal rigidity due to endogenous information acquisition. Thus, this paper introduces a complementary channel through which market structure matters for monetary policy. Firm size is systematically related both to information acquisition and to firms' partial-equilibrium cost pass-through. As a result, changes in the distribution of firm size alter not only average real rigidity and average attention, but also the *alignment* between the two. Market structure therefore affects monetary non-neutrality partly by determining *which pricing types are informed*, over and above the conventional real-rigidity or

oligopoly channels. In addition to emphasizing the

The paper proceeds as follows. Section 1 describes the construction of the 10-K attention measure and documents the main empirical fact. Section 2 develops the heterogeneous-firm model and derives the main theoretical results. Following this, Section 3 connects the model to the data via a sufficient-statistics approach to quantify the importance of information heterogeneity. Finally, Section 4 calibrates a structural model and presents the counterfactual. Finally, Section 5 concludes the paper.

1 Motivating Evidence

The goal of this section is to document the empirical fact that motivates the rest of the paper: firms differ systematically in how closely they track inflation and other aggregate nominal conditions. Larger firms within narrowly defined NAICS industries track these conditions more closely. Throughout the paper, I refer to this as greater *attention* to the aggregate nominal environment. In most models of information acquisition, greater attention implies a lower posterior uncertainty about the state, the aggregate nominal environment in this case. Empirically, this posterior uncertainty is hard to observe at a granular firm-level. I therefore construct a firm-year text-based proxy for attention using firms' annual 10-K filings. Annual 10-K filings provide a useful window into how firms describe the aggregate environment relevant for their business. In particular, the Business, Risk Factors, Management's Discussion and Analysis, Market Risk, and Financial Statements sections contain firms' discussion of important economic conditions, risks, and their implications. I extract two features of this language that are particularly informative about how closely firms track aggregate nominal conditions: the amount of substantive discussion devoted to inflation and monetary conditions, and, conditional on such discussion, the 'confidence' expressed in the language. The first captures the extent to which nominal conditions feature in the firm's disclosure; the second captures how confidently the firm characterizes those conditions when they are discussed.

The main advantage of the text-based approach relative to recent firm surveys such as SoFIE or the European SAFE is granularity. These surveys are extremely useful for measuring aggregate expectations and uncertainty, but firm characteristics are typically reported in relatively coarse size and industry categories. By contrast, 10-K filings can be linked directly to Compustat, allowing the text measure to be related to firm size and other characteristics. I therefore use the 10-K measure

for the baseline cross-sectional analysis and, where possible, use survey evidence as an external benchmark. The following subsections describe the data sources and the sample and then turn to describing the methodology used to construct the measure of attention at a granular firm-year level for all public companies that file 10-K with the SEC.

Data and Sample The primary input data consists of the universe of 10-K annual reports filed with the SEC between 2005 and 2020 [3], retrieved from EDGAR. Each filing is matched to the corresponding firm-year in *Compustat*, allowing me to obtain a range of firm characteristics, including its primary NAICS classification, revenues, and employment. We restrict to firms with at least one non-missing observation of employment or revenue over the sample. I end up with a sample of roughly 130,000 candidate filings and around 14000 firms. Each filing is then processed through an LLM pipeline to extract a text-based measure of firm attention to inflation.

Measuring Attention with an LLM Each filing is processed in five stages to arrive at the final measure of firm attention to inflation. Table 1 summarizes the stages. The methodology first in-

Stage	Task
A	Extract sections 1, 1A, 7, 7A, and 8
B	Split each section into 150–300 word paragraphs
C	Nominal-relevance classifier
D	Holistic uncertainty score $u_{ij} \in [0, 10]$
E	Weighted aggregation to the filing level

Table 1: The five-stage pipeline from raw 10-K to filing-level attention

volves extraction of the relevant sections of each 10-K. Sections 1 (Business), 1A (Risk factors), 7 (Management's Discussion and Analysis), 7A (Market Risk Disclosures) and 8 (Financial Statements) are most likely places where firms discuss their exposure and qualitative forecasts of aggregate nominal environments. These sections are then split into coherent paragraphs that contain enough context for the LLM to glean the main theme of discussion. The remaining stages identify passages that substantially discuss the aggregate nominal environment, evaluate the linguistic markers of confidence expressed in disclosures on these topics, and aggregate the passage-level outputs to the filing level.

Aggregate Nominal Topics The classifier reads each passage and returns an indicator for whether it substantively discusses inflation and the price level, monetary policy and central banks,

[3] Note that the legal requirement to declare all the risk factors that firms face in the 10-K took effect in 2005.

Prompt 1: Stage C: Nominal-relevance classifier (`gpt-4o-mini`)

System: *You are a classifier. Respond only with valid JSON.*

User: Does this passage substantively discuss ANY of the following? *Inflation and price level:* inflation, deflation, stagflation, hyperinflation, price level, rising prices, price pressures, pricing environment, price stability, purchasing power, cost inflation, input cost environment, cost pressures, cost escalation, cost pass-through, PPI, CPI, producer/consumer prices. *Monetary policy and central banks:* monetary policy, Federal Reserve, Fed, FOMC, ECB and other central banks; rate hikes/cuts, tightening/easing cycle, QE, QT, policy rates, fed funds rate, benchmark rates, yield curve, monetary conditions. *Nominal cost environment:* passages where the firm attributes cost pressure to the BROAD MACRO ENVIRONMENT rather than a specific named event or firm-specific factor. *Thematic coverage:* also flag passages that substantively discuss aggregate monetary conditions, economy-wide price levels, or interest rates even if none of the specific terms above appear. *Do not flag* passages where the cost pressure, uncertainty, or risk is attributable to firm-specific idiosyncratic conditions or named discrete events tied to the firm — even if prices or costs are mentioned.

Return: `{"relevant": true|false, "concept": "inflation_price_level"|"monetary_policy"|"nominal_cost_environment"|"mixed"|"none", "confidence": "high"|"low"}`

or broader aggregate nominal cost conditions. It is explicitly instructed to distinguish economy-wide nominal conditions from firm-specific cost shocks or named idiosyncratic events. The classifier also records a broad concept label including 'inflation and price level', 'monetary policy' and 'nominal cost environment'. Prompt 1 reports the exact prompt used to classify passages.

The share of passages that passed the classifier's filter is already a simple and interpretable measure of the prominence of discussion on nominal conditions in the filing. Discussion alone, however, does not distinguish a passage that characterizes the nominal environment precisely from one that explicitly emphasizes uncertainty or inability to characterize. I therefore use a second LLM step to capture this additional dimension of the disclosure.

Linguistic Confidence in Disclosure Each filtered passage that is classified as discussing nominal topics is passed through another model, `gpt-5-mini`, which assigns an uncertainty score $u_{ij} \in [0, 10]$. A score of zero corresponds to language that characterizes the nominal environment with high confidence, while a score of ten corresponds to language indicating that the nominal environment cannot be characterized. The model also records whether the passage is forward-

Prompt 2: Stage D: uncertainty scorer (`gpt-5-mini`)

System: *You are scoring passages from corporate 10-K filings. Respond only with valid JSON.*

User: This passage has been identified as discussing inflation, CPI/PPI, or monetary policy. Score the uncertainty expressed in the language. *uncertainty (0–10):* how uncertain does the firm sound about the inflation/monetary policy environment? 0 = complete confidence and confidence about the nominal outlook; 10 = cannot characterize the nominal environment at all. *Also return:*

- `temporal_orientation`: `forward` / `backward` / `mixed` (forward = about future inflation/policy path; backward = about past cost realizations).
- `boilerplate`: true/false (true = generic language any firm could use unchanged).
- `key_quote`: most informative phrase, max 10 words.
- `n_key_sentences`: count of sentences directly concerning inflation, monetary policy, or aggregate nominal conditions.

- `key_sentences`: those sentences verbatim.

Return: `{"uncertainty": <0-10>, "temporal_orientation": ..., "boilerplate": true|false, "key_quote": "...", "n_key_sentences": <int>, "key_sentences": [...]}`

looking, backward-looking, or mixed, as well as whether the language is generic boilerplate. The score should be interpreted as a feature of the disclosure language. Prompt 2 reports the exact instruction.

Define passage-level linguistic confidence as

$$p_{ij} \equiv 1 - u_{ij}/10 \in [0, 1].$$

The baseline filing-level measure is the *confidence-adjusted attention* index

$$A_{it} = \frac{1}{|\mathcal{B}_{it}|} \sum_{j \in \mathcal{C}_{it}} p_{ij} = \underbrace{\frac{|\mathcal{C}_{it}|}{|\mathcal{B}_{it}|}}_{\phi_{it}} \times \underbrace{\frac{1}{|\mathcal{C}_{it}|} \sum_{j \in \mathcal{C}_{it}} p_{ij}}_{\bar{p}_{it}}, \tag{1}$$

for filings with $|\mathcal{C}_{it}| > 0$, and $A_{it} = 0$ when no passage is classified as nominally relevant. Here, $|\mathcal{B}_{it}|$ is the total number of passages extracted for filing (i, t) and $|\mathcal{C}_{it}|$ is the number of passages classified as discussing the above-mentioned topics. The first term, ϕ_{it}, is the *extensive margin*: the share of all the passages in the filing devoted to aggregate nominal conditions. The second, $\bar{p}_{it}$, is the average linguistic confidence of those passages. Thus, a high value of A_{it} can arise because nominal conditions receive more discussion, because that discussion is more precise, or both. The index is a text-based proxy for attention; in Section 3, I later map its variation to an economically

interpretable Kalman gain using survey moments.

What does A_{it} capture? Two diagnostics of the data help clarify the content of A_{it}. First, Figure 1 shows the composition of passages that are flagged by the classifier as discussing aggregate nominal topics. Direct discussions of monetary policy account for 39% of admitted passages and inflation or the price level for another 17%; broader nominal-cost passages account for 41%. The last category is intentionally broad; passages classified under this category often discuss aggregate interest rates such as LIBOR, foreign exchange, capital and credit markets/financial conditions. Appendix A reports a detailed breakdown of topics for a random subset of hand-classified passages. Appendix A shows that the paper's main conclusions are unchanged when the measure is restricted to the narrower monetary-policy and inflation/price-level categories, explicitly excluding passages that are bucketed under the broad umbrella 'nominal cost environment'. Second,

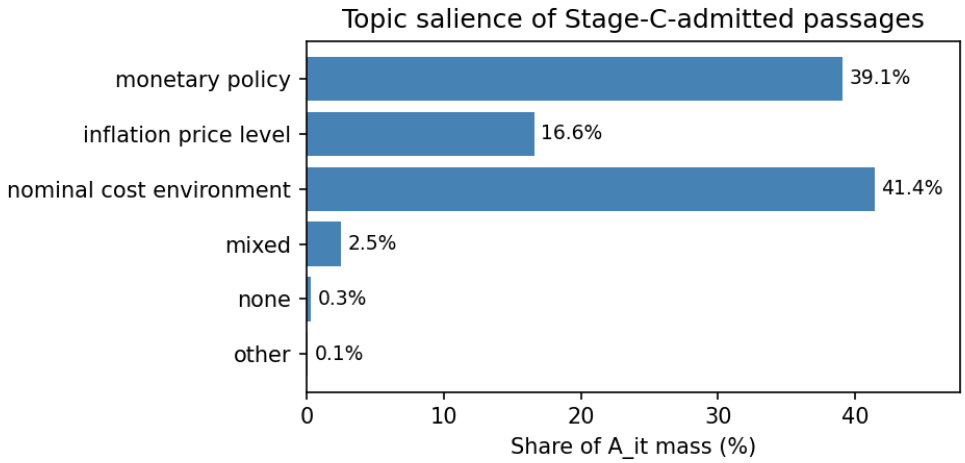


Figure 1: The figure shows the topic-wise composition of passages that passed the Stage C filter. Approximately 40% of the passages directly discuss monetary policy, and 17% discuss the price level and inflation. The rest of the passages are broadly bucketed under nominal cost environment.

the 'confidence' score captures recognizable differences in disclosure language. High-p_{ij} passages tend to contain concrete or quantitatively specific statements; intermediate passages often outline a conditional relationship with qualified assumptions; the low-confidence band is dominated by statements explicitly stating that the effects of nominal conditions or the nominal conditions cannot be predicted. Table 2 provides representative examples. Essentially, the takeaway is that the score is grading disclosure language on a dimension a human 10-K reader can recognize.

Band	Representative complete sentences
High $p \in [0.7, 1.0]$ ($\approx$ 22% of passages)	*"The decrease in interest expense is due to lower LIBOR rates and lower average borrowings during the year."* (industrial, 2010)
	"A hypothetical 100 basis point decrease in interest rates would reduce the estimated fair value of our mortgage loans by approximately $4.1 million." (financial, 2007)
Medium $p \in [0.5, 0.6]$ ($\approx$ 10% of passages)	*"If capital market conditions result in rising interest rates, we may be unable to refinance this indebtedness on acceptable terms."* (utility, 2013)
	"A low interest rate environment existed for approximately seven years, which compressed our net interest spreads." (community bank, 2016)
Low $p \in [0.0, 0.4]$ ($\approx$ 68% of passages)	*"The ultimate impact of future interest rate increases could result in unanticipated reductions in our net operating income."* (REIT, 2023)
	"The nature and impact of future changes in fiscal and monetary policies cannot be predicted." (holding company, 2010)

Table 2: Representative complete sentences across three confidence bands, where $p_{ij} = 1 - u_{ij}/10$. High-confidence passages are numeric and factual; medium-confidence passages are structured conditionals; low-confidence passages are disclaimer-heavy. Full five-band curated gallery in Appendix A

Independent checks on quality with a sentence-level classifier: As an independent check on this interpretation, I run a separate sentence-level classifier on the same filings. This classifier is not used in constructing A_{it}; instead, it records simple features of sentences. It classifies a sentence as discussing idiosyncratic or macro topics and further buckets them into broad categories, including inflation, monetary policy, exchange rates, interest rates, labor, legal compliance, environment, demand, and so on. It also flags if these sentences contain numerical content, an explicit time reference, directional language, ambiguity, or boilerplate. Higher-A_{it} firm-years contain systematically more numeric and time-specific sentences discussing 'macro-nominal topics'. Appendix A reports the exercise in detail. The key conclusion is that high A_{it} firms systematically reveal higher-quality quantitative statements and also qualitative forecasts about specifically aggregate nominal topics. As a placebo test, I check whether high A_{it} also correlates with more quantitative/time-specific sentences about topics other than 'macro-nominal'. The relationship is essentially absent for technology operations and geopolitics, and substantially weaker for labor. Appendix A reports these results. Highly attentive firms under the measure A_{it} disclose quality information specifically on macro-nominal topics as opposed to being firms with generally better disclosure quality across all topics.

Finally, Figure 2 shows which margin accounts for the variation in A_{it} across different quintiles of A_{it}. Approximately 9% of filings did not contain any passage discussing the topics in question. A_{it} is zero for these firms, and they are all in the lowest quintile of the distribution of A_{it}. Both extensive (ϕ_{it}) and intensive ($\bar{p}_{it}$) rise across quintiles of the index, but the extensive margin moves much more sharply: ϕ_{it} rises from 0.019 in the bottom quintile to 0.205 in the top quintile, while

$\bar{p}_{it}$ rises from 0.28 to 0.51. Thus, high-A_{it} firms are more often firms that devote more of their disclosures to aggregate nominal conditions, with linguistic confidence being a secondary source of variation.

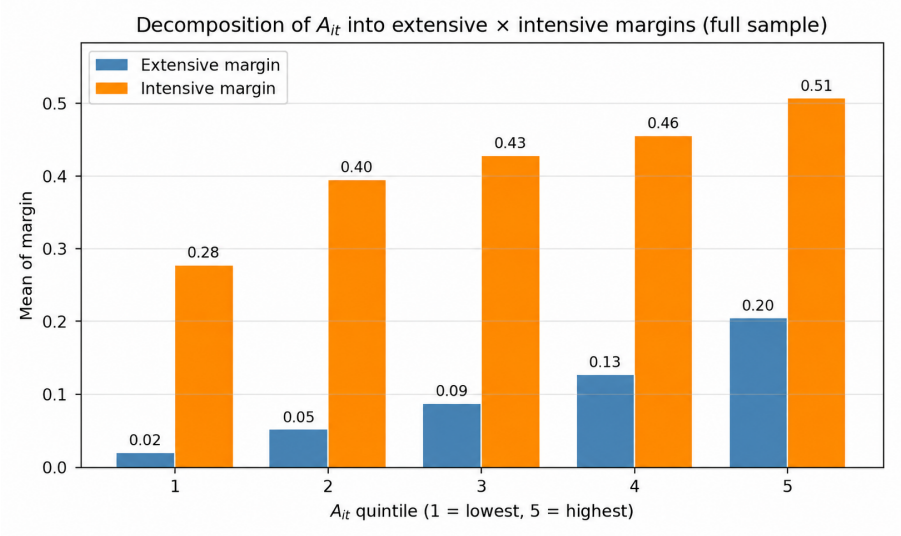


Figure 2: Decomposition of A_{it} into extensive $\times$ intensive margins by A_{it} quintile. The extensive margin ϕ (share of macro-relevant passages) rises eleven-fold across the distribution; the intensive margin π (mean confidence on relevant passages) rises less than two-fold. Dotted lines show grand means across firm-years ($\bar{\phi} = 0.098$, $\bar{\pi} = 0.424$).

External sensibility checks: The text measure also exhibits economically sensible variation. Its aggregate time series rises when inflation and monetary policy are more salient or volatile. At the industry level, average measured attention is higher in industries where costs are more exposed to inflationary pressures and lower in industries where prices are more rigid; Appendix B presents the exercises in detail.

Firm-size and inflation attention The key empirical fact is a strong positive relationship between firm size and the text 'attention' within narrowly defined industries. Figure 3 plots A_{it} against a firm's percentile rank in the within-NAICS-6-by-year revenue-share distribution, after residualizing A_{it} for NAICS-6-by-year fixed effects to keep scales comparable across industries. The pattern is striking: Larger firms have systematically higher measured attention. Moving from the bottom to the top of the within-industry size distribution raises A_{it} by roughly one-half of a standard deviation. The gradient is robust to alternative industry definitions and to alternative measures constructed from the same text, including the extensive margin alone and narrower measures based only on inflation/price-level and monetary-policy passages; Appendix A reports

these exercises. While this cross-sectional pattern does not seem surprising, there is very little empirical evidence in the literature that documents cross-sectional heterogeneity in firm attention. The result is related to Song and Stern (2025), who document that larger firms exhibit greater attention to macroeconomic conditions in general. Using Song and Stern (2025) data, one can also construct topic-specific textual measures of attention by restricting dictionaries to topic-specific words, including FOMC, inflation, and foreign exchange. There are two main differences. First, the LLM-based classification captures a substantially broader set of filings that discuss aggregate nominal conditions. Restricting the Song-Stern measure to nominal topics-inflation, FOMC, and foreign exchange, together with their associated dictionaries-about 60% of filings in the sample contain at least one such reference. By contrast, the LLM classifier identifies substantive discussion of aggregate nominal conditions in roughly 91% of filings. This difference likely reflects, at least in part, the ability of the LLM to identify nominal discussion from context rather than requiring the presence of a predefined keyword or phrase. Second, the empirical emphasis is different. I focus on the cross-sectional allocation of information about nominal conditions across firms and, in particular, how it varies systematically with firm characteristics. Song and Stern instead emphasize time-series variation in broader macroeconomic attention and its implications for the transmission and efficacy of monetary policy.

This raises two important questions: First, why do larger firms acquire more precise information about the aggregate nominal factors such as inflation and monetary policy? Next and more importantly, does this systematic relationship affect the degree of monetary non-neutrality? The following section addresses this question through the lens of a relatively simple heterogeneous firm general equilibrium model with aggregate nominal shocks.

2 Model

In this section, I develop a general equilibrium model of heterogeneous firms that set prices under incomplete information about an aggregate nominal shock. The household block is standard, thereby allowing a clear exposition of the novel channels I intend to emphasize. The production side features a continuum of sectors (indexed by j) and within each sector, there is a continuum of firms (indexed by i). Each firm i in sector j produces a differentiated intermediate good, and a competitive sectoral goods producer aggregates the output of all firms in a sector, using a generalized Kimball aggregator à la Kimball (1995), into a composite sectoral good. A key driver of

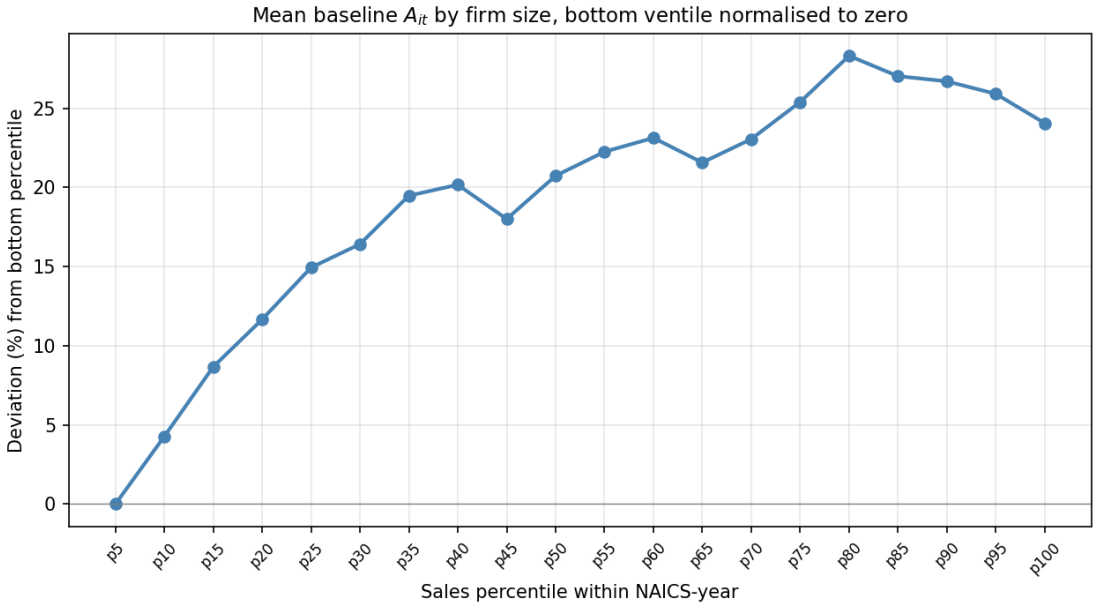


Figure 3: Confidence-adjusted attention A_{it} against firms' within-industry revenue-share percentile. Industry is defined at the NAICS 6-digit level, and the percentile is computed year-by-year; Industries with fewer than ten Compustat firms are excluded. NAICS 6-digit $\times$ year fixed effects are residualized out to emphasize the cross-sectional pattern. Points are means over twenty equal-population percentile bins with the bottom percentile as the base.

heterogeneity in the model is the permanent differences in productivity between firms. Given the homothetic Kimball aggregator, this generates heterogeneity in firm size, markups (and elasticity) as well as cost-shock pass-throughs. Firms do not observe underlying aggregate nominal shocks but can acquire costly information about them. Firms, in addition to making optimal pricing decisions, optimally acquire information about the underlying shock. Finally, sectoral composite goods are aggregated by a final goods producer using a CES production function to produce the final good. In this sector, I abstract from dynamics in the underlying nominal process and study a setting that is essentially a repeated static game. Section XX will allow for dynamics and will essentially reinforce the conclusions from this simple, repeated static setting.

Households

A representative household chooses consumption, labor, and nominal bond holdings. I adopt the Golosov-Lucas (Golosov and Lucas Jr. (2007)) style of preferences to describe the household side

owing to the various conveniences it affords.

$$\max_{\{C_t, L_t, B_t\}_{t\geq 0}} E_0 \sum_{t=0}^{\infty} \beta^t \left[\log C_t - \chi L_t\right] \tag{2}$$

subject to

$$P_t C_t + B_t \leq (1 + i_{t-1}) B_{t-1} + W_t L_t + \Pi_t. \tag{3}$$

To simplify the environment, I assume that the representative household has full information on all the aggregate states. Let nominal expenditure be $Q_t \equiv P_t C_t$. The household first-order conditions imply

$$Q_t = \frac{W_t}{\chi}, \qquad 1 = \beta(1 + i_t) E_t \left[\frac{Q_t}{Q_{t+1}}\right]. \tag{4}$$

I normalize $\chi = 1$, so nominal expenditure and the nominal wage move one-for-one. The monetary authority is thought of as directly controlling the nominal expenditure $P_t C_t \equiv Q_t$ and hence the nominal wages. The only stochastic fundamental in the baseline model is therefore an aggregate nominal wage shock. For the time being, I will assume that shocks to nominal expenditure, and hence to nominal wages, are i.i.d.

$$\widehat{W}_t \sim N(0, \sigma_W^2). \tag{5}$$

where $\widehat{W}_t$ refers to log deviation of W_t from a steady state W^*.

Also, it is easy to see that the intertemporal stochastic discount factor $M_{t,t+1} = \beta \frac{Q_t}{Q_{t+1}}$

Final-good aggregation across sectors

A competitive final-good producer combines sectoral composite goods $Y_{j,t}$ using a CES aggregator,

$$Y_t = \left[\int_0^1 \vartheta_j^{1/\sigma} Y_{j,t}^{(\sigma-1)/\sigma} dj\right]^{\sigma/(\sigma-1)}, \qquad \sigma > 1, \tag{6}$$

where $\vartheta_j > 0$ and $\int_0^1 \vartheta_j dj = 1$ can be thought of as sectoral demand shifters. Expenditure minimization gives sectoral demand

$$Y_{j,t} = \vartheta_j \left(\frac{P_{j,t}}{P_t}\right)^{-\sigma} Y_t \tag{7}$$

and the aggregate price index

$$P_t = \left[\int_0^1 \vartheta_j P_{j,t}^{1-\sigma} dj\right]^{1/(1-\sigma)}. \tag{8}$$

Sectoral aggregation

Within sector j, a competitive sectoral producer aggregates differentiated firm outputs using a Kimball aggregator. The nested CES-Kimball structure I adopt in this setting is similar to that adopted in other papers that study the implications of market structure for monetary policy, including Wang and Werning (2022). Sectoral output is implicitly defined by

$$\int_0^1 G\left(\frac{y_{ij,t}}{Y_{j,t}}\right) di = 1, \tag{9}$$

where G is increasing and concave. Let $\lambda_{j,t}$ denote the (normalized by sectoral expenditure $P_{j,t}Yj,t$) multiplier associated with (9) and define $F \equiv (G')^{-1}$. The aggregator's cost minimization yields conditional demand for firm-level output

$$y_{ij,t} = Y_{j,t} F\left(\frac{P_{ij,t}}{\lambda_{j,t} P_{j,t}}\right). \tag{10}$$

Combining (7) and (10), the demand faced by firm ij is

$$y_{ij,t} = \vartheta_j \left(\frac{P_{j,t}}{P_t}\right)^{-\sigma} Y_t F\left(\frac{P_{ij,t}}{X_{j,t}}\right). \tag{11}$$

Note that the firm's price relative to the index $\lambda_{j,t}P_{j,t}$ determines its demand. In essence, each firm in a given sector competes against a common sectoral price index $\lambda_{j,t}P_{j,t}$. The same relative price also determines a firm's own price elasticity of demand. A change in the price index for a given price set by the firm affects the elasticity of demand and, in turn, the firm's markup. This induces a change in the optimal price, creating strategic complementarity in pricing. The zero profits condition of the aggregator gives the sectoral price index

$$P_{j,t} = \int_0^1 P_{ij,t} F\left(\frac{P_{ij,t}}{\lambda_{j,t} P_{j,t}}\right) di \tag{12}$$

The aggregator constraint becomes

$$1 = \int_0^1 G\left[F\left(\frac{P_{ij,t}}{\lambda_{j,t} P_{j,t}}\right)\right] di. \tag{13}$$

The Kimball class nests CES demand. In particular, if $G(q) = q^{(\varepsilon-1)/\varepsilon}$, (9) becomes a CES aggregator with elasticity ε.

Firms

Technology Firm i in sector j has permanent productivity Z_{ij} and produces using a linear production technology,

$$y_{ij,t} = Z_{ij} L_{ij,t}, \qquad MC_{ij,t} = \frac{W_t}{Z_{ij}}. \tag{14}$$

where $L_{ij,t}$ is the production labor employed by the firm. Productivity is fixed over time and is known to the firm. Importantly, permanent productivity heterogeneity, combined with the Kimball aggregator, generates within-sector cross-sectional variation in firm size, markups, and cost pass-throughs in the steady state that I shall elaborate on below.

Information Firms do not observe aggregate nominal shocks, i.e., shocks to nominal wage W_t. Firms are rationally inattentive and acquire information subject to a finite information-processing capacity that they optimally choose. Information capacity is produced using labor Let L_{ij}^{info} units of labor be devoted to the production of information processing capacity given by

$$\kappa_{ij} = \left(\frac{L_{ij}^{\text{info}}}{\omega} \right)^{1/p}, \qquad \text{or equivalently} \qquad L_{ij}^{\text{info}} = \omega \kappa_{ij}^{p}, \tag{15}$$

with $p \geq 1$. Firms set prices conditional on their information to maximize expected discounted profits. Following Sims (2003), I model limited attention as a constraint on the amount of information that a firm can process about a payoff-relevant state. Formally, before setting its price, a firm chooses an information structure, represented by a conditional distribution

$$\pi_{ij,t} \left(ds \mid W_t, \mathcal{I}_{ij,t-1} \right),$$

An information structure specifies, for every possible realization W_t of the aggregate state, a probability distribution over signals. Thus, the firm chooses ex ante the stochastic mapping from the underlying state W_t into a signal $S_{ij,t} \in \mathcal{S}_{ij,t}$, where the signal space $\mathcal{S}_{ij,t}$ is left unrestricted and may be multidimensional. After observing the realization of $S_{ij,t}$, the firm's information set becomes

$$\mathcal{I}_{ij,t} = \mathcal{I}_{ij,t-1} \vee \sigma \left(S_{ij,t} \right),$$

and its price must be measurable with respect to $\mathcal{I}_{ij,t}$.

The amount of information acquired about the aggregate state is measured by Shannon mutual

information [4]. Firm ij 's information structure must satisfy

$$I\left(W_t; S_{ij,t} \mid \mathcal{I}_{ij,t-1}\right) \leq \kappa_{ij,t}, \tag{16}$$

where $\kappa_{ij,t}$ denotes the firm's information-processing capacity. Mutual information essentially measures the reduction in uncertainty, as measured by the entropy of the uncertain random variable W_t, generated by the firm's signal; Therefore, mutual information $I\left(W_t; S_{ij,t}\right)$ is given by

$$I\left(W_t; S_{ij,t} \mid \mathcal{I}_{ij,t-1}\right) = H\left(W_t \mid \mathcal{I}_{ij,t-1}\right) - H\left(W_t \mid S_{ij,t}, \mathcal{I}_{ij,t-1}\right),$$

where $H(\cdot)$ denotes entropy. Thus, $\kappa_{ij,t}$ restricts the informativeness of the firm's signal. At this stage, the information structure is left unrestricted. Later, after taking a second-order approximation to the firm's problem, I show that the resulting Gaussian-quadratic tracking problem admits an optimal information structure that can be represented by a single scalar Gaussian signal.

Timing I assume that shocks to nominal expenditure (and hence nominal wages controlled by the central bank) follow an i.i.d. process. The i.i.d. nature of the shocks reduces the apparently dynamic problem into essentially a series of repeated static interactions. The within-period timing assumptions make this amply clear. At the beginning of the period, the aggregate nominal wage shock is realized but is not directly observed by firms. Then, each firm chooses its information-processing capacity and information structure, observes its signal, and sets the price and commits to fulfilling the demand it receives at that price. Then demand is realized at the set price. The firm observes the quantity demanded and meets demand by hiring enough production labor. The labor supply adjusts enough to clear the labor market at the wage rate. The model shares an enduring feature of most New Keynesian models. The output is demand-determined.

Firm's Problem Firm ij chooses its information capacity and information structure, and sets the price to maximize its expected discounted profit stream.

The firm's exact problem is then

$$\max_{\left\{\kappa_{ij,t}, \pi_{ij,t}, P_{ij,t}(\cdot)\right\}_{t\geq 0}} E_0 \sum_{t=0}^{\infty} \beta^t \frac{1}{Q_t} \left[\left(P_{ij,t} - \frac{W_t}{Z_{ij}}\right) y_{ij,t} - W_t \omega \kappa_{ij,t}^p\right]$$

[4]For more information about the use of Shannon information in rational-inattention models, see Sims (2003) and Mackowiak and Wiederholt (2009)

subject to demand (10), information flow constraint (16) and

$$P_{ij,t} \quad \text{is } \mathcal{I}_{ij,t}\text{-measurable}, \quad \mathcal{I}_{ij,t} = \mathcal{I}_{ij,t-1} \vee \sigma\left(S_{ij,t}\right).$$

Here $\frac{1}{Q_t}$ represents the household discount factor. Because prices are flexible, production is demand-determined over the period, and the only fundamental disturbance is i.i.d., the firm's intertemporal problem reduces to a sequence of identical static information-and-pricing problems. I retain the dynamic notation here to clarify the timing of information acquisition and to facilitate a faster setup of the dynamic problem in Section 6.

Equilibrium

Given the stochastic process for nominal expenditure $\{Q_t\}$, an equilibrium consists of household allocations; firm, sectoral, and aggregate prices and quantities; and, for each firm, an information capacity, information structure, and signal-contingent pricing rule such that: (i) households optimize; (ii) the competitive final-good and sectoral producers minimize expenditure given the CES and Kimball aggregators; (iii) each firm optimally chooses its information capacity, information structure, and price subject to its information-processing constraint, taking as given equilibrium mappings from the aggregate state W_t into aggregate and sectoral endogenous variables (iv) firms meet realized demand at posted prices; (v) firms' beliefs about the equilibrium mappings are consistent with those generated by equilibrium pricing decisions and aggregation; and (vi) goods, labor, and asset markets clear.

Approximation

I work with a second-order approximation of the firm's problem around the "Full-information asymmetric steady state ".

The Asymmetric Steady State Heterogeneity in Z_{ij} induces asymmetry in the steady state along three main dimensions: firm size, the price elasticity of demand faced by firms, and cost shock pass-through. Let

$$\eta\left(\frac{P_{ij}}{\lambda_j P_j}\right) \equiv -\frac{\partial \log F\left(\frac{P_{ij}}{\lambda_j P_j}\right)}{\partial \log\left(\frac{P_{ij}}{\lambda_j P_j}\right)}, \qquad \mu\left(\frac{P_{ij}}{\lambda_j P_j}\right) \equiv \frac{\eta\left(\frac{P_{ij}}{\lambda_j P_j}\right)}{\eta\left(\frac{P_{ij}}{\lambda_j P_j}\right) - 1}. \tag{17}$$

where $\eta\left(\frac{P_{ij}}{\lambda_j P_j}\right)$ denotes the price elasticity of demand at relative price $\frac{P_{ij}}{\lambda_j P_j}$. In the steady state,

$$P_{ij}^* = \mu\left(\frac{P_{ij}}{\lambda_j P_j}\right)\frac{W^*}{Z_{ij}}, \qquad \mu\left(\frac{P_{ij}}{\lambda_j P_j}\right) \equiv \frac{\eta\left(\frac{P_{ij}}{\lambda_j P_j}\right)}{\eta\left(\frac{P_{ij}}{\lambda_j P_j}\right) - 1}. \tag{18}$$

Define the super-elasticity of demand $\xi\left(\frac{P_{ij}}{\lambda_j P_j}\right)$

$$\xi\left(\frac{P_{ij}}{\lambda_j P_j}\right) \equiv \frac{\partial \log \eta\left(\frac{P_{ij}}{\lambda_j P_j}\right)}{\partial \log\left(\frac{P_{ij}}{\lambda_j P_j}\right)} \tag{19}$$

Holding the sectoral strategic index $\lambda_j P_j$ fixed, the firm's partial-equilibrium pass-through of marginal cost into price is defined as

$$\rho\left(\frac{P_{ij}}{\lambda_j P_j}\right) \equiv \left.\frac{\partial \log P_{ij}}{\partial \log MC_{ij}}\right|_{\lambda_j P_j} = \frac{\eta\left(\frac{P_{ij}}{\lambda_j P_j}\right) - 1}{\eta\left(\frac{P_{ij}}{\lambda_j P_j}\right) - 1 + \xi\left(\frac{P_{ij}}{\lambda_j P_j}\right)}. \tag{20}$$

Under CES, $\xi\left(\frac{P_{ij}}{\lambda_j P_j}\right) = 0$ for all ij and $\rho\left(\frac{P_{ij}}{\lambda_j P_j}\right) = 1$ for all ij. Kimball demand curvature generates incomplete pass-through. Furthermore, heterogeneity in permanent productivity Z_{ij} creates variation in firm size, elasticity, and cost pass-through within the sector.

Assumption 1 (Demand curvature and pass-through ordering)**.** *On the relevant support,* $\eta\left(\frac{P_{ij}}{\lambda_j P_j}\right) > 1$, $\xi\left(\frac{P_{ij}}{\lambda_j P_j}\right) > 0$, *and*

$$\frac{d}{d\left(\frac{P_{ij}}{\lambda_j P_j}\right)}\left[\frac{\xi\left(\frac{P_{ij}}{\lambda_j P_j}\right)}{\eta\left(\frac{P_{ij}}{\lambda_j P_j}\right) - 1}\right] < 0. \tag{21}$$

Lemma 1 (Productivity, size, and pricing type)**.** *Under Assumption 1, higher-productivity firms are larger within their sectors, face lower demand elasticities, charge higher markups, and have lower partial-equilibrium cost pass-through.*

The proof is provided in Appendix C.1. The mechanism is straightforward. A more productive firm has lower marginal cost and therefore charges a lower relative price, $\frac{P_{ij}}{\lambda_j P_j}$. Because F is decreasing, this lower relative price increases the firm's quantity sold and, hence, its sales. This size effect is present under CES demand as well. With Kimball demand, however, variation in relative prices also generates heterogeneity in demand elasticities and superelasticities. In particular, when the superelasticity is positive over the relevant support, a lower relative price is associated

with a lower demand elasticity and therefore a higher markup. Condition (21) further implies that the same decline in relative price is associated with lower cost pass-through. Although this ordering is imposed as an assumption here, it is consistent with the empirical evidence. In particular, Edmond, Midrigan and Xu (2023) and Burya and Mishra (2022) provide evidence consistent with positive demand super-elasticity, while a substantial empirical literature documents that larger firms exhibit lower cost pass-through, with the most notable being Amiti, Itskhoki and Konings (2019).

Before moving ahead, the setup is far enough to illustrate another important feature of this environment. The pass-through statistic $\rho\left(\frac{P_{ij}}{\lambda_j P_j}\right)$ also summarizes the firm's local complementarity with the sectoral strategic price index $\lambda_j P_j$. To see this, recall that the firm's optimal markup depends on its relative price $\frac{P_{ij}}{\lambda_j P_j}$. Log-differentiating the full-information pricing condition gives

$$d\log P_{ij} = d\log MC_{ij} - \underbrace{\frac{\xi\left(\frac{P_{ij}}{\lambda_j P_j}\right)}{\eta\left(\frac{P_{ij}}{\lambda_j P_j}\right) - 1}}_{=\frac{d\log\mu\left(\frac{P_{ij}}{\lambda_j P_j}\right)}{d\log\left(\frac{P_{ij}}{\lambda_j P_j}\right)}} \left[d\log P_{ij} - d\log(\lambda_j P_j)\right].$$

Rearranging this equation is very revealing:

$$d\log P_{ij} = \rho\left(\frac{P_{ij}}{\lambda_j P_j}\right) d\log MC_{ij} + \left[1 - \rho\left(\frac{P_{ij}}{\lambda_j P_j}\right)\right] d\log(\lambda_j P_j). \tag{22}$$

Notice that the firm's elasticity and hence its optimal markup depend only on its price relative to the price index $\lambda_j P_j$. A higher own price, keeping $\lambda_j P_j$ fixed, raises the relative price and lowers the markup (under positive super-elasticity). This markup response dampens the price response to any cost shocks. The higher the markup elasticity, the greater the dampening of the price response and the lower the pass-through of the firm's own cost shock. Conversely, an increase in $\lambda_j P_j$ lowers the firm's relative price and raises its desired markup, inducing the firm to raise its price. Thus, the same markup elasticity governs both the response to own-cost shocks and the degree of strategic complementarity locally. This feature is common across demand systems where the demand for a variety is only a function of its price relative to a common index.

For notational simplicity, I will use the shorthands η_{ij}, ξ_{ij} and ρ_{ij} to refer to steady-state elasticity, super-elasticity and pass-through of firm i in sector j.

Optimal Pricing Let X^* denote the deterministic full-information steady state of variable X, and define

$$\widehat{X}_t \equiv \log X_t - \log X^*.$$

$\widehat{Q}_t = \widehat{W}_t$ are then innovations to nominal expenditure

$$W_t = W^* e^{\widehat{W}_t}, \quad \widehat{W}_t \overset{i.i.d.}{\sim} N\left(0, \sigma_W^2\right).$$

Notice also that because W_t maps to $\widehat{W}_t$ one-to-one, therefore

$$I\left(W_t; S_{ij,t} \mid \mathcal{I}_{ij,t-1}\right) = I\left(\widehat{W}_t; S_{ij,t} \mid \mathcal{I}_{ij,t-1}\right).$$

So moving from levels to log deviations changes nothing about the information constraint.

Define the firm's total steady-state sales share by

$$s_{ij}^* \equiv \frac{P_{ij}^* y_{ij}^*}{P^* Y^*}. \tag{23}$$

The second-order approximation of profits around the full-information is especially useful here. The firm's problem boils down to minimizing the expected deviation of its price from the full-information target.

$$\beta^t \min_{\kappa_{ij},\, \pi_{ij}} \left\{ E\left[\min_{\widehat{P}_{ij,t}} E_{ij,t}\left[\frac{1}{2} \underbrace{s_{ij}^* \frac{\eta_{ij}-1}{\rho_{ij}}}_{\text{curvature of the profit function}} \left(\widehat{P}_{ij,t} - \widehat{P}_{ij,t}^f\right)^2 \right] \right] + \omega \kappa_{ij}^p \right\}, \tag{24}$$

subject to

$$I\left(\widehat{W}_t; S_{ij,t} \,\middle|\, \mathcal{I}_{ij,t-1}\right) \leq \kappa_{ij}, \tag{25}$$

where

$$S_{ij,t} \sim \pi_{ij,t}\left(\cdot \mid \widehat{W}_t, \mathcal{I}_{ij,t-1}\right), \qquad \mathcal{I}_{ij,t} = \mathcal{I}_{ij,t-1} \vee \sigma(S_{ij,t}),$$

and

$$E_{ij,t}[\cdot] \equiv E[\cdot \mid \mathcal{I}_{ij,t}].$$

Here, $P_{ij,t}^f$ is the optimal price they would choose under full information. The full-information

target price is given by

$$\widehat{P}^f_{ij,t} = \rho_{ij}\widehat{W}_t + (1-\rho_{ij})\left(\widehat{P}_{j,t} + \widehat{\lambda}_{j,t}\right). \tag{26}$$

The firm acquires information to track payoff-relevant states $\widehat{W}_t$, $\widehat{P}_t$ and $\widehat{\lambda}_t$ that determine the full-information optimal price $P^f_{ij,t}$. Notice that the local curvature of the profit function around the steady state is specific to firm ij and is determined by the firm's steady state size s^*_{ij}, demand elasticity η_{ij}, and pass-through ρ_{ij}. This curvature informs how costly deviations from the optimal full information price would be in utility terms. Thus, a percentage deviation of the price set from the firm's optimal target is costlier for firms with higher curvature. Details of the derivation can be found in Appendix C.2.

Proposition 1 (Optimal pricing)**.** *1. Given any information set $\mathcal{I}_{ij,t}$, the optimal price is*

$$\widehat{P}_{ij,t} = E_{ij,t}\left[P^f_{ij,t}\right] = \rho_{ij}E_{ij,t}[\widehat{W}_t] + (1-\rho_{ij})E_{ij,t}\left[\widehat{P}_{j,t} + \widehat{\lambda}_{j,t}\right]. \tag{27}$$

2. Under the linear equilibrium conjecture,

$$\widehat{P}_{j,t} = \alpha_j\widehat{W}_t, \qquad \widehat{\lambda}_{j,t} = \ell_j\widehat{W}_t. \tag{28}$$

Then the full-information target becomes

$$\widehat{P}^f_{ij,t} = \underbrace{\left[\rho_{ij} + (1-\rho_{ij})(\alpha_j + \ell_j)\right]}_{\equiv\phi_{ij}}\widehat{W}_t, \tag{29}$$

and the optimal price is

$$\widehat{P}_{ij,t} = \left[\rho_{ij} + (1-\rho_{ij})(\alpha_j + \ell_j)\right]E_{ij,t}[\widehat{W}_t]. \tag{30}$$

Given the quadratic loss, the optimal price is the conditional expectation of the full-information target under the information structure chosen by the firm. Proposition 1 delivers an important simplification: conditional on the equilibrium response of sectoral prices, the firm's target depends on aggregate conditions only through the scalar nominal wage/expenditure shock $\widehat{W}_t$. The information problem therefore boils down to tracking a single aggregate state $\widehat{W}_t$.

It is also useful to note that the full-information exposure

$$\phi_{ij} \equiv \rho_{ij} + (1-\rho_{ij})(\alpha_j + \ell_j)$$

is increasing in ρ_{ij} whenever $\alpha_j + \ell_j < 1$, since

$$\frac{\partial \phi_{ij}}{\partial \rho_{ij}} = 1 - (\alpha_j + \ell_j) > 0.$$

Thus, the ordering of firms by partial-equilibrium cost pass-through is preserved in their full-information response to the aggregate shock:lower-pass-through firms also have lower full-information exposure. This property holds for the class of demand systems in which pricing/demand elasticity depend on the price relative to a price index. I next characterize the optimal information structure for a given information-processing capacity, before turning to the firm's optimal choice of capacity.

Optimal information structure

Lemma 2 (Optimal information structure)**.** *Suppose*

$$\widehat{W}_t \overset{\text{i.i.d.}}{\sim} N(0, \sigma_W^2),$$

and consider firm ij's quadratic tracking problem for a given information-processing capacity κ_{ij}. There exists an optimal memoryless information structure that can be represented by a scalar Gaussian signal

$$s_{ij,t} = \widehat{W}_t + \varepsilon_{ij,t}, \qquad \varepsilon_{ij,t} \sim N(0, \sigma_{\varepsilon,ij}^2), \qquad \varepsilon_{ij,t} \perp \widehat{W}_t, \tag{31}$$

where the signal errors can be taken to be independent over time. The posterior mean of the aggregate state is

$$E_{ij,t}[\widehat{W}_t] = \gamma_{ij} s_{ij,t}, \qquad \gamma_{ij} = \frac{\sigma_W^2}{\sigma_W^2 + \sigma_{\varepsilon,ij}^2}. \tag{32}$$

For an informative signal and a nondegenerate tracking problem, the information constraint binds:

$$\kappa_{ij} = \frac{1}{2} \log_2 \left(1 + \frac{\sigma_W^2}{\sigma_{\varepsilon,ij}^2} \right), \tag{33}$$

so that

$$\gamma_{ij} = 1 - 2^{-2\kappa_{ij}}. \tag{34}$$

This is a well-known result in the rational inattention literature. Nevertheless, a formal proof can be found in Appendix C.3. The result follows from the Gaussian-quadratic structure of the tracking problem. Any signal can be replaced by the scalar posterior mean it induces without

changing the firm's pricing decision. Gaussian optimality then implies that this scalar information structure can, up to a rescaling, be represented by the additive Gaussian signal in (31). Finally, because $\widehat{W}_t$ is i.i.d., past signals do not improve the prior for the current state, so an optimal information structure can be chosen to be memoryless. It now remains to determine the optimal capacity κ_{ij}.

2.1 Optimal information capacity

By lemma 2, the optimal information structure implies

$$\text{Var}\left(\widehat{W}_t \mid s_{ij,t}\right) = \sigma_W^2 \, 2^{-2\kappa_{ij}}.$$

Substituting this into the quadratic pricing loss, the firm's capacity choice problem becomes to

$$\min_{\kappa_{ij} \geq 0} \left\{ \frac{1}{2} \mathcal{V}_{ij} 2^{-2\kappa_{ij}} + \omega \kappa_{ij}^p \right\}, \tag{35}$$

where

$$\mathcal{V}_{ij} \equiv \underbrace{s_{ij}^* \frac{\eta_{ij} - 1}{\rho_{ij}}}_{\text{Profit curvature}} \left[\rho_{ij} + (1 - \rho_{ij})(\alpha_j + \ell_j)\right]^2 \sigma_W^2 \tag{36}$$

summarizes the stakes from tracking the aggregate shock. Under optimal pricing, the expected quadratic pricing loss is proportional to the expected variance of the firm's full-information target price.[5] The firm is essentially trying to minimize the conditional variance of its target price subject to a cost. This formulation also clarifies the components of $\mathcal{V}_{ij}$.

Information is more valuable when a given pricing error is more costly and when the firm's target price varies more with the aggregate state. The former is governed by $s_{ij}^*(\eta_{ij} - 1)/\rho_{ij}$: in particular, the same percentage pricing error generates a larger loss for a firm that is larger in steady state and when the curvature of profit function $(\eta_{ij} - 1)/\rho_{ij}$ is higher. The latter is governed by the firm's exposure to the aggregate shock, $\phi_{ij} \equiv \rho_{ij} + (1 - \rho_{ij})(\alpha_j + \ell_j)$, together with the variance σ_W^2 of that shock.

Proposition 2 (Optimal information capacity)**.** *1. Optimal capacity: For an interior solution, opti-*

[5] $\mathcal{V}_{ij} 2^{-2\kappa_{ij}}$ equals profit curvature times the conditional variance of the firm's full-information target price.

mal information capacity satisfies

$$\ln(2)\,\mathcal{V}_{ij}\,2^{-2\kappa_{ij}^*} = p\omega(\kappa_{ij}^*)^{p-1}. \tag{37}$$

For $p > 1$, the solution is unique. Moreover, κ_{ij}^*, and hence the Kalman gain

$$\gamma_{ij}^* = 1 - 2^{-2\kappa_{ij}^*},$$

is strictly increasing in $\mathcal{V}_{ij}$.

2. *Size gradient: For a given $\alpha_j + \ell_j$, within a sector, $\mathcal{V}_{ij}$ is monotonically increasing in firm relative size if*

$$\begin{aligned} &1 - \eta_{ij} + \frac{\eta_{ij}}{\eta_{ij} - 1}\xi_{ij} \\ &\qquad + \left[2\frac{\left[1 - (\alpha_j + \ell_j)\right]\rho_{ij}}{\rho_{ij} + (1 - \rho_{ij})(\alpha_j + \ell_j)} - 1\right]\frac{d\log\rho_{ij}}{d\log\left(\frac{P_{ij}}{\lambda_j P_j}\right)} < 0. \end{aligned} \tag{38}$$

If (38) *holds throughout the relevant support, then larger firms optimally acquire more information and have higher Kalman gains.*

Appendix C.4 proves the proposition. Rearranging (37),

$$2^{2\kappa_{ij}^*} = \frac{\ln(2)\mathcal{V}_{ij}}{p\omega(\kappa_{ij}^*)^{p-1}}. \tag{39}$$

Under the optimal information structure,

$$2^{2\kappa_{ij}} = \frac{Var(\widehat{W}_t)}{Var(\widehat{W}_t \mid s_{ij,t})},$$

so the left-hand side measures the factor by which the firm reduces uncertainty about the aggregate state. The numerator on the right-hand side summarizes the stakes/benefits from doing so, while the denominator is the marginal cost of producing information capacity.

Firm size enters these stakes directly through s_{ij}^*. All else fixed, a larger firm incurs proportionally larger losses from a given pricing error and therefore has a stronger incentive to resolve uncertainty. Firm size, demand elasticity, and pass-through, of course, are all functions of the steady-state relative price. The second part of the proposition gives the formal condition under

which the direct size effect gradient of capacity is monotonic in relative size.

2.2 Aggregation and monetary non-neutrality

The preceding analysis characterizes each firm's information acquisition and pricing decision conditional on the response of sectoral aggregates, $\alpha_j + \ell_j$. We now close the equilibrium by requiring these aggregate responses to be consistent with the prices chosen by firms. At the optimal information choice, firm ij's price response to the nominal shock is

$$\widehat{P}_{ij,t} = \gamma^*_{ij} \left[\rho_{ij} + (1-\rho_{ij})(\alpha_j + \ell_j)\right] \widehat{W}_t. \tag{40}$$

Thus, a firm's contribution to sectoral price adjustment depends on both how well it tracks the aggregate shock, through γ^*_{ij}, and how strongly its desired price responds to that shock, through its pricing type ρ_{ij}. I now aggregate firms' pricing decisions to characterize the response of sectoral and aggregate prices to the nominal shock. Throughout this section, let

$$s^*_{ij|j} \equiv \frac{P^*_{ij} y^*_{ij}}{P^*_j Y^*_j}, \qquad E^s_j[z_{ij}] \equiv \int_0^1 s^*_{ij|j} z_{ij}\, di$$

denote firm ij's steady-state sales share within sector j and the associated sales-weighted expectation. By construction, $\int_0^1 s^*_{ij|j}\, di = 1$.

Sectoral fixed point Recall the conjecture

$$\widehat{P}_{j,t} + \widehat{\lambda}_{j,t} = (\alpha_j + \ell_j)\widehat{W}_t.$$

Aggregating firms' optimal pricing rules through the Kimball price-index conditions yields a fixed-point equation for $\alpha_j + \ell_j \equiv x_j$:

$$x_j = \frac{E^s_j\left[\eta_{ij}\gamma^*_{ij}(x_j)\left\{\rho_{ij} + (1-\rho_{ij})x_j\right\}\right]}{E^s_j[\eta_{ij}]} \tag{41}$$

Note that γ^*_{ij} itself depends on $\alpha_j + \ell_j \equiv x_j$ through the firm's information-acquisition problem. (41) is a fixed-point problem. Appendix C.5 derives this equation, establishes existence, and provides a sufficient condition for uniqueness. This leads us to the most important aggregation result.

Sectoral price response Let x_j^* denote the equilibrium solution to (41), and evaluate firms' optimal Kalman gains at this solution, $\gamma_{ij}^* = \gamma_{ij}^*(x_j^*)$.

Given x_j^*, the linearized sectoral price index is given by $\widehat{P}_{jt} = E_j^s\left[\widehat{P}_{ij,t}\right]$ and the linearized aggregate price index is given by $\widehat{P}_t = E^S\left[\widehat{P}_{jt}\right]$. [6]

Proposition 3 (Price pass-through). *1. Sectoral pass-through: The equilibrium response of the sectoral price index is*

$$\widehat{P}_{j,t} = \alpha_j \widehat{W}_t, \tag{42}$$

where

$$\alpha_j = E_j^s[\gamma_{ij}^*] - K_j^+ E_j^s\left[\gamma_{ij}^*(1-\rho_{ij})\right], \tag{43}$$

and

$$K_j^+ \equiv \frac{E_j^s[\eta_{ij}] - E_j^s[\eta_{ij}\gamma_{ij}^*]}{E_j^s[\eta_{ij}] - E_j^s[\eta_{ij}\gamma_{ij}^*(1-\rho_{ij})]}. \tag{44}$$

2. Aggregate pass-through: The aggregate price index responds according to

$$\widehat{P}_t = \alpha \widehat{W}_t, \qquad \alpha = \int_0^1 S_j^* \alpha_j \, dj, \tag{45}$$

where

$$S_j^* \equiv \frac{P_j^* Y_j^*}{P^* Y^*} \tag{46}$$

is sector j's steady-state expenditure share.

Proof can be found in Appendix C.6.

The first part of Proposition 3 shows that sectoral price adjustment depends not only on firms' average responsiveness to information, but also on how information is distributed across firms with different pricing types. Under CES demand, $\rho_{ij} = 1$ for all firms and the second term in (43) disappears, yielding

$$\alpha_j^{CES} = E_j^s[\gamma_{ij}^*].$$

This is the familiar 'average attention effect' emphasized in the literature. More generally, under symmetry in $\eta_{ij} = \eta$ and $_{ij} = \rho$, the responsiveness of the aggregate price index α_j is given by $\frac{\rho \mathbb{E}_j^s[\gamma_i]}{1-(1-\rho)\mathbb{E}_j^s[\gamma_i]}$, with no role for the joint distribution of firm-level Kalman gain and pass-through. However, given the heterogeneity in firms' demand curvature, the allocation of information across

[6] Derived in Appendix C.6

firms matters for sectoral price rigidity. Another key aspect to note from the proposition is that, for a given distribution of Kalman gains, the price responsiveness α_j under CES is lower than under the baseline model. Therefore, the degree of price rigidity under the baseline is always higher than under CES. Furthermore, it is important to note that there are two separate channels that increase aggregate price rigidity relative to the baseline under CES. The following split of the sectoral price responsiveness α_j is particularly revealing here.

$$\alpha_j = E_j^s[\gamma_{ij}^*] - K_j^+ \underbrace{E_j^s[\gamma_{ij}^*]E_j^s[1-\rho_{ij}]}_{\text{average pass-through}} + K_j^+ \underbrace{\text{Cov}_j^s(\gamma_{ij}^*, \rho_{ij})}_{\text{Covariance of attention and pass-through}} . \tag{47}$$

The decomposition is useful for intuition, although it does not exactly separate the two forces: K_j^+ itself depends on the joint distribution of demand elasticities, pass-throughs, and Kalman gains. In the sufficient statistics section that follows, I will isolate the 'Covariance channel' separately from other distributional moments.

The term $E_j^s[\gamma_{ij}^*]E_j^s[1-\rho_{ij}]$ captures the effect of incomplete pass-through or the canonical "real rigidity channel". Even in the absence of distributional channels, incomplete pass-through dampens sectoral price adjustment. The term $\text{Cov}_j^s(\gamma_{ij}^*, \rho_{ij})$ captures how information is allocated across pricing types. When more informed firms have lower pass-through, $\text{Cov}_j^s(\gamma_{ij}^*, \rho_{ij}) < 0$ and this alignment reduces sectoral price responsiveness.

In the following section, I empirically quantify the magnitude of the reduction in price responsiveness in the baseline relative to the benchmark models. I use measures of firm-level steady-state measures $\left\{s_{ij|j}^*, \rho_{ij}, \eta_{ij}, \gamma_{ij}\right\}_{i\in j}$ that are sufficient to compute the degree of monetary non-neutrality under various models nested in this environment. Here, I isolate the two forces above using explicit counterfactual environments rather than relying on the algebraic decomposition in (47).

3 Measuring Rigidity

The aggregation result summarized by Proposition 3 in the previous section identifies a small set of firm-level objects that are sufficient to quantify the degree of monetary non-neutrality. Recall that the response of the sectoral price index to the nominal shock is

$$\alpha_j = E_j^s[\gamma_{ij}] - K_j^+ E_j^s\left[\gamma_{ij}(1-\rho_{ij})\right],$$

where

$$K_j^+ \equiv \frac{E_j^s[\eta_{ij}] - E_j^s[\eta_{ij}\gamma_{ij}]}{E_j^s[\eta_{ij}] - E_j^s[\eta_{ij}\gamma_{ij}(1-\rho_{ij})]}$$

Here, $E_j^s[\cdot]$ denotes a steady-state sales-weighted cross-sectional mean within sector j. These equations therefore imply that sectoral price rigidity is pinned down by the stationary joint distribution

$$\left\{ s^*_{ij|j}, \rho_{ij}, \eta_{ij}, \gamma_{ij} \right\}_{i \in j}, \tag{48}$$

where $s^*_{ij|j}$ is the firm's steady-state sales share within its sector, ρ_{ij} is its cost pass-through in the steady state, η_{ij} is its demand elasticity in the steady state, and γ_{ij} is its Kalman gain on the aggregate nominal shock [7]. The objective of this section is to construct empirical counterparts of these four objects and use their joint distribution to evaluate the implied rigidity of the aggregate price index.

3.1 From Textual Attention to Kalman Gains

The 10-K exercise in Section 1 delivers a firm-year text-based measure A_{it} of attention to inflation and aggregate nominal conditions. While A_{it} is useful for documenting cross-sectional heterogeneity in attention, it does not have an economically meaningful interpretable scale. The object that enters the aggregation result is instead the Kalman gain γ_{it}: the weight a firm places on new information when updating its belief about the aggregate nominal state. I therefore use a state-space model to map the textual attention measure into an implied firm-year Kalman gain. Importantly, this mapping is estimated independently of the structural model in Section 2: I do not impose the model-predicted relationship between the Kalman gain and firm characteristics such as size, demand elasticity, or pass-through when recovering γ_{it}. The resulting Kalman gains can therefore be used both as input to the quantification of aggregate rigidity and as independent evidence on the cross-sectional relationships implied by the model.

For the measurement exercise, I take inflation as a summary of the aggregate nominal state

[7] This is not a steady-state object per se. It represents the weight on firms' private information on the underlying nominal shocks that cause fluctuations around this steady state.

that the firms are trying to learn about. [8] I assume that it follows an AR(1) process,

$$\pi_t - \mu_\pi = \rho_\pi(\pi_{t-1} - \mu_\pi) + u_t, \qquad u_t \sim N(0, \sigma_\pi^2). \tag{49}$$

Firm i receives a noisy private signal about current inflation,

$$s_{it} = \pi_t + \zeta_{it}, \qquad \zeta_{it} \sim N(0, R_{it}), \tag{50}$$

where R_{it} is the variance of the signal noise. Let $P_{it|t-1}$ denote the firm's prior variance about inflation before observing s_{it}. Standard Bayesian updating gives

$$\widehat{\pi}_{it|t} = \widehat{\pi}_{it|t-1} + \gamma_{it}\left(s_{it} - \widehat{\pi}_{it|t-1}\right), \qquad \gamma_{it} = \frac{P_{it|t-1}}{P_{it|t-1} + R_{it}}. \tag{51}$$

Thus, firms with more precise signals, i.e., a lower R_{it}, place greater weight on new information and have a larger Kalman gain. This setup is nothing but a standard Bayesian learning model. The assumption is that each firm receives an independent private signal about the state and the precision of this signal is firm-specific. Both of these assumptions are consistent with the optimal information structure in the environment described in Section 2.

Next, I connect the precision of the firm's signal to the textual measure A_{it}. The assumption I make here is that a firm with high measured A_{it} acquires a signal with higher precision R_{it}^{-1}. Since textual measured attention A_{it} has no natural economic scale and is likely to be contaminated by random measurement error, I allow for a latent attention state a_{it},that maps into firm-level signal precision R_{it}^{-1}. This latent attention a_{it} is modeled as having a permanent component $\bar{a}_i$ that is firm-specific and a persistent component x_{it}.[9]

$$a_{it} = \bar{a}_i + x_{it}, \qquad x_{it} = \rho_a x_{i,t-1} + \omega_{it}, \tag{52}$$

where $\bar{a}_i$ captures persistent differences in attention across firms. The standardized text measure,

[8] I use the estimated inflation Kalman gains as proxies for gains on the nominal state $\hat{W}_t$. In the static model, this equivalence holds given known lagged price level: $E[\pi_t \mid \mathcal{I}_{i,t}] = \alpha E[\widehat{W}_t \mid \mathcal{I}_{i,t}] - \widehat{P}_{t-1}$. Its extension to dynamic learning is a measurement assumption.

[9] The rationale for this is that A_{it} is highly persistent within a firm and the assumption allows for consistent estimation of the persistence parameter of latent a_{it}.

denoted by $\widetilde{A}_{it}$, is a noisy measure of this latent state,

$$\widetilde{A}_{it} = c + a_{it} + \varepsilon_{it}. \tag{53}$$

I then link latent attention to signal precision through a monotone function

$$R_{it}^{-1} = g(a_{it}; \theta), \qquad \frac{\partial g(a;\theta)}{\partial a} > 0. \tag{54}$$

The exact parameterization of $g(\cdot)$, together with the laws of motion and measurement-error assumptions, is described in Appendix D.1. The only key assumption worth noting is that $g(\cdot)$ is assumed to be increasing in a_{it}. That is, a higher latent attention a_{it} translates into higher signal precision received by the firm.

The state-space model is estimated using via SMM using moments from both the 10-K attention measure and the Survey of Firms' Inflation Expectations (SoFIE). The two data sources play complementary roles. By itself the text does not pin down the economic scale of attention, that is, how differences in a_{it} translate into differences in firms' information precision. I use SoFIE to discipline this mapping. SoFIE is a quarterly survey of U.S. firms that elicits one-year-ahead CPI inflation expectations.

Let $\overline{F}_t$ denote the cross-sectional mean forecast and D_t the cross-sectional dispersion of forecasts, and define

$$FE_t \equiv \pi_t - \overline{F}_t, \qquad u_t \equiv (\pi_t - \mu_\pi) - \rho_\pi(\pi_{t-1} - \mu_\pi). \tag{55}$$

The moments used in the estimation and their fit are reported in Table 3.

The SoFIE moments give the latent attention measure an interpretable scale. The time-series variation in the error of the average forecast, $\mathrm{Var}_t(FE_t)$, helps pin down the overall precision of firms' information: more informative signals allow the average belief to track inflation more closely. The level and time variation of D_t discipline the heterogeneity in precision across firms. Any systematic variation by firm characteristics is determined by the text-based measure A_{it}. To better understand the role of each of the data sources, consider the following:

$$\frac{\partial R_{it}^{-1}}{\partial \log s_{ij|j,t}^{obs}} = \underbrace{\frac{\partial g(a_{it};\theta)}{\partial a_{it}}}_{\substack{\text{attention–precision mapping} \\ \text{disciplined by SoFIE}}} \underbrace{\frac{\partial a_{it}}{\partial \log s_{ij|j,t}^{obs}}}_{\substack{\text{attention–size gradient} \\ \text{disciplined by 10-Ks}}}.$$

Table 3: Moments Targeted in the Attention State-Space Model

Source	Moment	Data	Model	Data–Model
10-K	$\mathbb{E}[\widetilde{A}_{it}]$	0.0150	-0.0793	0.0944
10-K	$\mathrm{SD}(\widetilde{A}_{it})$	0.9554	0.9018	0.0536
10-K	$Q_{0.25}(\widetilde{A}_{it})$	-0.6700	-0.6913	0.0212
10-K	$Q_{0.75}(\widetilde{A}_{it})$	0.4167	0.5323	-0.1156
10-K	$\mathrm{Corr}(\widetilde{A}_{it}, \widetilde{A}_{i,t-1})$	0.1977	0.1952	0.0025
10-K	$\mathrm{Corr}_t(\overline{A}_t, \overline{A}_{t-1})$	0.8601	0.7801	0.0800
10-K	$\mathbb{E}_t[\mathrm{SD}_i(\widetilde{A}_{it})]$	0.9410	0.9018	0.0392
SoFIE	$\mathrm{Var}_t(FE_t)$	1.4000	1.3852	0.0148
SoFIE	$\mathbb{E}_t[D_t]$	1.4036	1.2079	0.1956
SoFIE	$\mathrm{Var}_t(D_t)$	0.3490	0.3030	0.0460
SoFIE	$\mathrm{Cov}_t(D_t, \lvert u_t \rvert)$	0.4163	0.3933	0.0230

Notes: $\overline{A}_t$ denotes the cross-sectional mean of the standardized 10-K attention measure in year t. FE_t is the forecast error of the cross-sectional mean SoFIE inflation forecast, D_t is cross-sectional forecast dispersion, and u_t is the innovation to the inflation state.

where R_{it}^{-1} is given by (54). The 10-K data determine the slope of the precision R_{it}^{-1} by firm characteristics, most importantly the relative size of firm i in industry j, $s^{obs}ij|j,t$, in this case, while the survey moments determine how these differences translate to differences in Kalman gains. Together, they determine how information precision varies with firm size.

Appendix D.1 lays out the identification arguments for each estimated parameter.

The estimated model delivers a sequence of firm-year Kalman gains $\{\widehat{\gamma}_{it}\}_t$. The procedure gives a time series of Kalman gains for each firm in the sample. Figure 4 plots the resulting $\widehat{\gamma}_i^*$ against firms' position in the within-industry steady-state size distribution. The size gradient documented in Section 2 persists under the transformation from text measure to Kalman gains: larger firms have systematically higher implied Kalman gains.

3.2 Steady-State Sales Shares and Pricing Characteristics

The remaining three objects in (48) are the firm's steady-state sales share, demand elasticity, and cost pass-through.

Steady-state sales shares. For the empirical exercise, I use the NAICS 6-digit industry classification to define the boundary of a market within which firms compete. This is also consistent with the definition of an industry used in Edmond, Midrigan and Xu (2023), which studies the welfare costs of markups at the macroeconomic level.

Computing the degree of sectoral price responsiveness α_j requires the steady-state objects in (48). However, Compustat provides annual realized sales shares. I therefore construct an empiri-

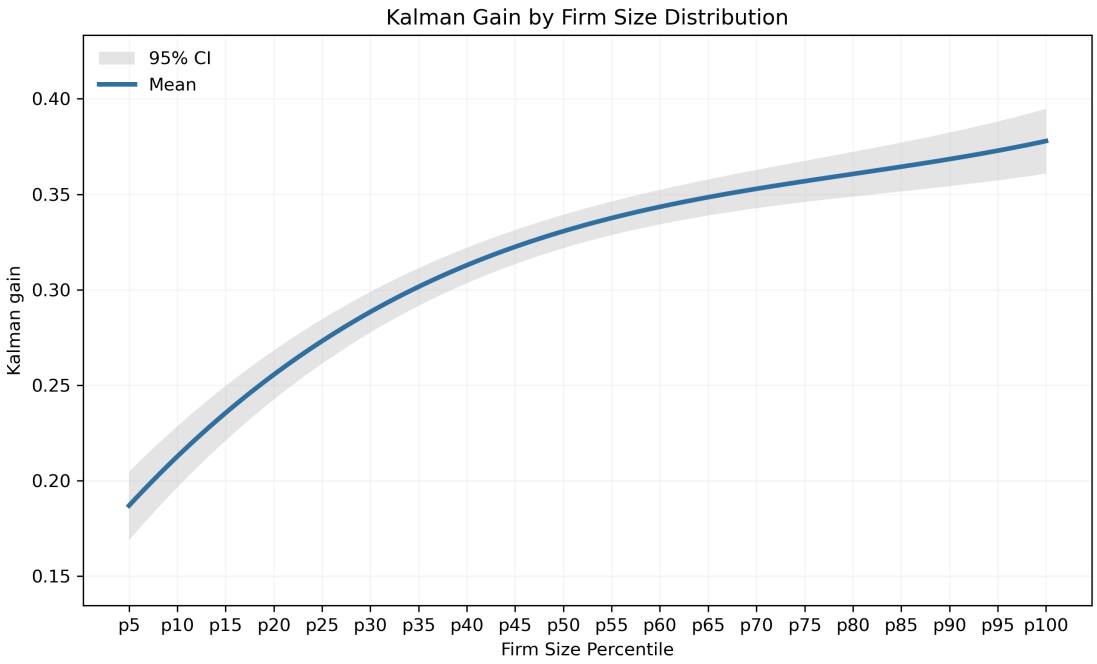


Figure 4: Kalman gains by relative size rank

Notes: The figure plots the fitted relationship between the estimated firm-year Kalman gain, $\widehat{\gamma}_{it}$, and firm size. Firm size is measured by the firm's percentile rank in the within-industry-year revenue-share distribution, where industries are defined at the NAICS 6-digit level [a] The solid line is a smooth fit of $\widehat{\gamma}_{it}$ on the size rank. A higher Kalman gain implies that the firm places greater weight on new information when updating its belief about inflation.

[a]For all the empirical exercises, I use NAICS 6-digit classification to define an industry.

cal approximation to the local steady-state cross-section using short rolling windows. Starting in 2013 and going up to 2017, for each year τ, I consider a five-year window, $(\mathcal{W}_\tau = \tau - 2, \ldots, \tau + 2)$. Within each NAICS-6 sector j, I keep firms that are active in the target year and observed at least twice within the window. For each such firm, let $s^{obs}_{ij|j,t}$ denote its observed sales share within sector j in year t. I estimate the firm's local steady-state log sales share around year τ as

$$\log \widehat{s}^{*}_{ij|j,\tau} = \frac{1}{T_{ij,\tau}} \sum_{t \in \mathcal{W}_{ij,\tau}} \log s^{obs}_{ij|j,t}. \tag{56}$$

I treat the resulting distribution as the local steady-state distribution of within-industry sales shares around year τ. [10] I end up with five steady-state cross-sections that I use to compute measures of price responsiveness α_j.

This construction relies on an assumption that the observed sales shares are just steady-state

[10]The shares are re-normalized to sum up to one for any given cross-section.

values with mean-zero disturbances over a given time window. Appendix D.2 derives this approximation within the model in Section 2. To a first order, realized log sales shares fluctuate around their steady-state values with mean-zero transitory deviations, so averaging over a short window recovers the local steady-state cross-section. I use a relatively short window because longer horizons endanger the assumption that the "steady-state" is not evolving. Long horizons may involve changes in the underlying steady state arising from persistent firm growth, demand shifts, entry and exit, or changes in industry structure.

Demand elasticities and cost pass-throughs Given the estimated steady-state sales-share distribution, I impute firm-level demand elasticities and cost pass-throughs using the estimates of Burya and Mishra (2022). Using Nielsen Retail Scanner data on firm-level prices and quantities, the authors estimate demand elasticities and super-elasticities using a granular supply-side instrument that exploits regional variation in firm-specific cost shocks. They estimate demand elasticity and cost pass-through as functions of firms' market shares within Nielsen product categories. The result of their estimation is a full schedule of firm-level demand elasticities and cost pass-throughs, showing how they vary with firm market share within a granular product category. These schedules are estimated conditional on industry HHI. I proxy for a firm's product-market position using its sales share within its 6-digit NAICS industry.[11] I impute firm elasticity and cost pass-through using steady-state sales share $\widehat{s}^*_{ij|j}$, the implied cross-sectional HHI_j and the Burya and Mishra (2022) estimated schedules:

$$\widehat{\eta}_{ij} = \eta^{BM}\left(\widehat{s}^*_{ij|j}, HHI^*_J\right), \qquad \widehat{\rho}_{ij} = \rho^{BM}\left(\widehat{s}^*_{ij|j}, HHI^*_J\right). \tag{57}$$

Alternative strategies for estimating firm-level steady-state elasticity and cost pass-throughs are explored in Appendix C.1. For elasticities, my primary alternative source is firm-level markups estimated via the production-function approach, à la De Loecker, Eeckhout and Unger (2020). For pass-throughs, Amiti, Itskhoki and Konings (2019) provide robust alternative estimates. They estimate partial equilibrium pass-throughs for Belgian manufacturing firms by firm size. I impute pass-throughs to firms in our sample using the sales-share rank within the industry as a linking variable.

[11] Nielsen product categories are roughly about the same granularity as NAICS-6-digit classification. In fact, the FDA/RTI product classification framework treats Nielsen ScanTrack categories and 6-digit NAICS industries as alternative product-category definitions Muth et al. (2015). Important caveats are assumptions of external validity made while exporting these estimates to all Compustat industries.

3.3 Results and Decomposition

For each target year τ I have the cross-section of four firm-level objects in (48). I also average the estimated firm-year Kalman gains over the same window, denoted by $\gamma_{ij,\tau}$. All sectoral and aggregate price responses in this subsection are computed separately for each target year using these local steady-state objects, and the reported results average the resulting gaps across target years. One thing to note is that the empirical joint distribution exhibits the same ordering emphasized by the theory. This is true for all the local steady states considered. Firms that are larger within their sectors tend to have higher Kalman gains, while the externally estimated pricing schedules assign them lower demand elasticities and lower cost pass-throughs. These two gradients imply a negative sales-weighted covariance between attention and cost pass-through within the average sector. An important point is that this covariance is not imposed by the structural model. The firm-level Kalman gains are recovered from the text and survey data independently of the model-implied relationship between attention and firm characteristics. The negative attention–pass-through covariance therefore provides independent empirical support for the model's information-acquisition mechanism: firms that optimally acquire more information are also those with lower pass-through.

Implications for price rigidity and the Phillips curve I now use the measured firm-level objects to quantify the aggregate price response under alternative environments nested in the model of Section 2. Aggregating the sectoral responses using steady-state Domar weights gives the economy-wide price response α. The corresponding static Phillips-curve relation is

$$\widehat{P}_t = \frac{\alpha}{1-\alpha}\widehat{Y}_t. \tag{58}$$

A lower α therefore implies a larger output response to a nominal shock and a flatter Phillips curve.

Recall that (47)

$$\alpha_j = E_j^s[\gamma_{ij}] - K_j^+ E_j^s[\gamma_{ij}]E_j^s[1-\rho_{ij}] + K_j^+ Cov_j^s(\gamma_{ij}, \rho_{ij}).$$

This expression motivates the quantitative decomposition below. Let

$$\bar{\gamma}_j \equiv E_j^s[\gamma_{ij}].$$

Table 4: Price Rigidity Relative to the CES Benchmark

	$\Delta\alpha$ (%)	Phillips-curve flattening (%)
Homogeneous attention, $\gamma_{ij} = \bar{\gamma}_j$	-22.30	32.61
Additional negative-covariance channel	-5.31	6.53
Baseline	-27.61	39.14

Notes: All changes are measured relative to the CES benchmark, which sets $\rho_{ij} = 1$. The homogeneous-attention economy introduces the empirical distributions of ρ_{ij} and η_{ij} while assigning all firms within a sector the common sales-weighted Kalman gain $\bar{\gamma}_j = E_j^s[\gamma_{ij}]$. The negative covariance row reports the additional change obtained by restoring the empirical within-sector distribution of γ_{ij}. The baseline combines both margins. The table reports averages across target years.

I consider three nested environments. The first sets $\rho_{ij} = 1$ for all firms, corresponding to the CES benchmark. In this case,

$$\alpha_j^{CES} = \bar{\gamma}_j, \tag{59}$$

and only average attention matters. Next, I restore the empirical distributions of ρ_{ij} and η_{ij} and assign every firm in sector j the common gain $\bar{\gamma}_j$. Denote the resulting price response by α_j^R. This counterfactual captures the contribution of "Kimball" real rigidity and asymmetry while shutting down within-sector heterogeneity in attention. For convenience, I call this the "real rigidity" benchmark. Finally, I restore the empirical firm-level Kalman gains γ_{ij} and obtain the baseline response α_j. Relative to the CES benchmark, the baseline economy's price responsiveness differs by $\alpha_j^{CES} - \alpha_j$. I decompose this gap

$$\alpha_j^{CES} - \alpha_j = \underbrace{\left(\alpha_j^{CES} - \alpha_j^R\right)}_{\text{real rigidity}} + \underbrace{\left(\alpha_j^R - \alpha_j\right)}_{\text{negative covariance channel}}. \tag{60}$$

The first component measures how much of the total gap relative to CES is accounted for by the "real rigidity" channel, when attention is homogeneous. The second is the additional contribution from the allocation of attention across pricing types. I compute this decomposition separately for each target year $\tau \in \{2013, \ldots, 2017\}$, using the corresponding steady-state distribution constructed from the five-year window centered around that year. I then aggregate the results by averaging the sector-level gaps across target years. Thus, the reported numbers summarize the average contribution of each channel across these steady states. Table 8 reports the change in aggregate price responsiveness and the Phillips-curve slope relative to the CES benchmark. In the baseline, aggregate price responsiveness is 27.61 percent lower than in the CES benchmark. Of this decline, about a fifth of it is attributed to the "negative covariance" channel. The baseline

Phillips curve is 39.14 percent flatter than under CES. Of this, again about a fifth is attributed to the "negative covariance" channel. More concretely, a 1% monetary shock increases the output response on impact by about 18% relative to the CES benchmark.

4 Calibration and Results

The preceding exercise quantifies monetary non-neutrality directly from the measured stationary joint distribution of firm relative size, pricing characteristics, and attention. I now turn to the structural model, calibrated to the U.S. economy around the mid-2010s. The calibration disciplines the parameters governing firm heterogeneity, demand curvature, and information acquisition. The purpose of the structural model is to study how counterfactual changes in within-industry market concentration affect monetary non-neutrality. As in Baqaee, Farhi and Sangani (2024), I vary the primitive governing within-industry firm heterogeneity while holding demand and information-cost parameters fixed, and study how aggregate price rigidity responds. Such a counterfactual requires the model because changes in the primitive determinant of market structure—the productivity distribution—jointly determine firms' relative sizes, demand elasticity, cost pass-through, and endogenous attention. Since monetary non-neutrality depends on the joint distribution of precisely these objects, the model provides a way to trace how changes in market structure alter aggregate price rigidity.

For transparency, the benchmark quantitative economy is a representative sector: the sector-level demand shifters v_j are normalized to 1, and the underlying distribution from which firm-level productivity is drawn is common across sectors. The aggregate price response therefore coincides with the representative-sector response. This benchmark isolates the within-sector mechanism and abstracts away from the effect of sectoral heterogeneity.

Calibration proceeds in two stages. In the first stage, I jointly calibrate the demand-side parameters $(\varepsilon, \theta, \psi)$ with the productivity distribution parameters via SMM using cross-sectional moments constructed from Compustat. Given the parameter estimates from Stage 1, Stage 2 then calibrates the two information-cost parameters (ω, p) to the moments from the empirical Kalman-gain schedule recovered in Section 3. I normalize the household preference parameter $\chi = 1$, so that in steady state $W^* = P^* Y^* = 1$.

4.1 Parameterization

Functional form for Kimball: I parameterize the Kimball aggregator $G(\cdot)$ as

$$G(q) \;=\; 1 \;+\; c\Big[\,\psi\,P_{+}(q) \;+\; (1-\psi)\,H(q)\,\Big], \qquad c \;\equiv\; \frac{\varepsilon}{\varepsilon-1+\psi}, \tag{61}$$

where the two component functions $P_{+}(\cdot)$ and $H(\cdot)$ are given in Appendix E.1. The specification (61) is a convex combination of the form used in Klenow and Willis (2016) piece P_{+}, and that in Dotsey and King (2005).The specification nests the Klenow and Willis (2016) and the Dotsey and King (2005) separately: $\psi = 1$ recovers pure Klenow and Willis (2016) and $\psi = 0$ recovers pure Dotsey and King (2005).

The mixture is useful because the empirical pass-through schedule varies substantially over the firm-size distribution. Small firms exhibit pass-through close to one, whereas pass-through is much lower among the largest firms (Amiti, Itskhoki and Konings, 2019; Burya and Mishra, 2022). With a reasonable elasticity schedule, the constant super-elasticity specification à la Klenow and Willis (2016) cannot generate this range. Appendix E.1 elaborates on this.

On the other hand, the Dotsey and King (2005) specification generates a reasonable range of cost pass-throughs by firm relative size but imposes a hard cut-off of 0.5 for fringe small firms, as their relative quantities tend to 0. Appendix E.1 provides more details and also shows that the parametric form used here successfully overcomes these limitations of the classic Kimball parametric demand functions. It does quite well in jointly matching the pass-through and elasticity gradients by firm size, as well as the sales share distribution.

The parameter ε determines the elasticity of demand as the relative size of the firm approaches zero. I set this equal to 16, which implies a markup of about 5% for the smallest of firms. The other parameters $(\theta,\, \psi)$ are calibrated jointly with the productivity distribution.

Productivity distribution: The firm-size distribution in the US displays two features: a substantial mass of similarly sized firms populating the middle of the distribution, and a much heavier right tail. This has been documented even in census data by Kondo, Lewis and Stella (2023). To capture these features parsimoniously, I parameterize the productivity distribution as the sum of a fringe block and a Pareto block.

A fraction $1 - \pi_{\text{tail}}$ of the blocks have a productivity that lies lies $\pm 1\%$ of mean body produc-

tivity z_{body}, while the remaining blocks draw productivity from a shifted Pareto distribution:

$$z_i = \begin{cases} z_{\text{body}}, & \text{body block,} \\ 1 + Y_i, \quad Y_i \sim \text{Pareto}(\alpha_p), & \text{tail block.} \end{cases} \tag{62}$$

Under this specification, the productivity distribution F is characterized by three parameters $(\pi_{\text{tail}}, z_{\text{body}}, \alpha_p)$. The mixture (62) accommodates both a right-skewed body and a heavy tail.

Information cost The technology for producing information-processing capacity is the one introduced in Section 2:

$$\kappa_i = \left(\frac{L_i^{\text{info}}}{\omega} \right)^{1/p}, \qquad \text{equivalently} \qquad L_i^{\text{info}} = \omega \kappa_i^p. \tag{63}$$

The parameters (ω, p) govern the level and curvature of information costs. Conditional on the Stage 1 demand and productivity parameters, these parameters determine the endogenous cross-sectional distribution of information capacity and Kalman gains through the sectoral fixed point.

4.2 Calibration Strategy and Target Moments

Stage 1: demand and productivity: Let

$$\Theta_1 = (\theta, \psi, \pi_{\text{tail}}, z_{\text{body}}, \alpha_p).$$

Stage 1 targets five moments related to firm-level size, markups, and pass-through to estimate the demand and productivity parameters outlined above. The parameters are estimated jointly via SMM. Table 5 summarizes the targets, their empirical values, and the primary role of each moment in identifying the six Stage 1 parameters. All within-industry moments are computed at the NAICS 6-digit level and averaged across sectors with at least 10 firms in the five-year window around the 2015 Compustat sample. Heuristically, the concentration moments are especially informative about productivity heterogeneity, whereas the pass-through and markup moments are especially informative about demand curvature. More elaborate identification arguments can be found in Appendix E.

Table 5: Stage-1 calibration targets

#	Moment	Value	Source
1	top-10% within-industry sales share	0.61	Compustat
2	ρ_{bot80}	0.802	Avg ρ for bottom 80th perecntile
3	ρ_{top20}	0.567	Avg ρ for top 20th percentile
4	β: sales-wtd. OLS of $\log \mu$ on $\log s^*$	$+0.03$	Compustat
5	$\overline{\mu}$: sales-wtd. mean markup	1.10	Basu (2019)

Notes. The top-10% within-industry share is computed within-industry at NAICS-6 and averaged across sectors with at least 10 firms in Compustat 2012-2018. Pass-through moments ρ_{bot80} and ρ_{top20} correspond to mean sales-weighted average pass-through for bottom 80th percentile and top 20th percentile in the sample. The markup gradient β is the coefficient from a sales-weighted regression of firm-level log markups on log within-industry sales shares with NAICS 6-digit fixed effects. Firm-level markups are constructed using the production-function approach in the spirit of De Loecker, Eeckhout and Unger (2020). All Stage 1 parameters are estimated jointly from the full vector of moments.

Stage 2: information costs: Conditional on $\widehat{\Theta}_1$, I estimate $\Theta_2 = (\omega, p)$ using moments of the Kalman-gain schedule recovered in Section 3. In particular, I target the Kalman gains (averaged across industries) at two points in the within-industry size distribution: the 5th and 95th.

$$\widehat{\gamma}(p_5) = 0.18, \quad \widehat{\gamma}(p_{95}) = 0.46 \tag{64}$$

Together, these moments help discipline (ω, p).

4.3 Parameters and Model Fit

The model is calibrated via the SMM procedure. Tables 6 and 7 report the calibrated parameter values and the moment fit at those values.

Table 6: Calibrated parameter values

Parameter	Value	Source
Panel A. Demand		
θ	1.07	Stage 1
ψ	0.19	Stage 1
Panel B. Productivity		
π_{tail}	0.4	Stage 1
z_{body}	1.0153	Stage 1
α_p	7.06	Stage 1
Panel C. Information		
ω	8.24×10^{-2}	Stage 2
p	1.92	Stage 2

Notes. The elasticity ε is fixed at 16 and other Stage-1 parameters are jointly estimated via SMM.

Table 7: Model fit at calibrated parameters

Moment	Data	Model	Std. resid.
Stage 1			
top-10% share	0.61	0.49	-1.2
ρ_{bot80}	0.802	0.858	$+1.114$
ρ_{top20}	0.567	0.561	-0.120
$\overline{\mu}$	1.100	1.139	$+0.781$
β	0.02718	0.03677	$+0.320$
Stage 2			
$\widehat{\gamma}(p_5)$	0.18	0.18	0
$\widehat{\gamma}(p_{100})$	0.460	0.460	0

Notes. Standardized residuals are $(m^{\text{model}} - m^{\text{data}})/\text{SE}(m)$ are reported. All five Stage-1 moments are hit within ± 1 standard error.

The average pass-through in the model closely aligns with measured pass-through for US firms reported in Dogra et al. (2023). Furthermore, Figure 5 shows the simulated model's Kalman gain schedule against the corresponding empirical schedule.

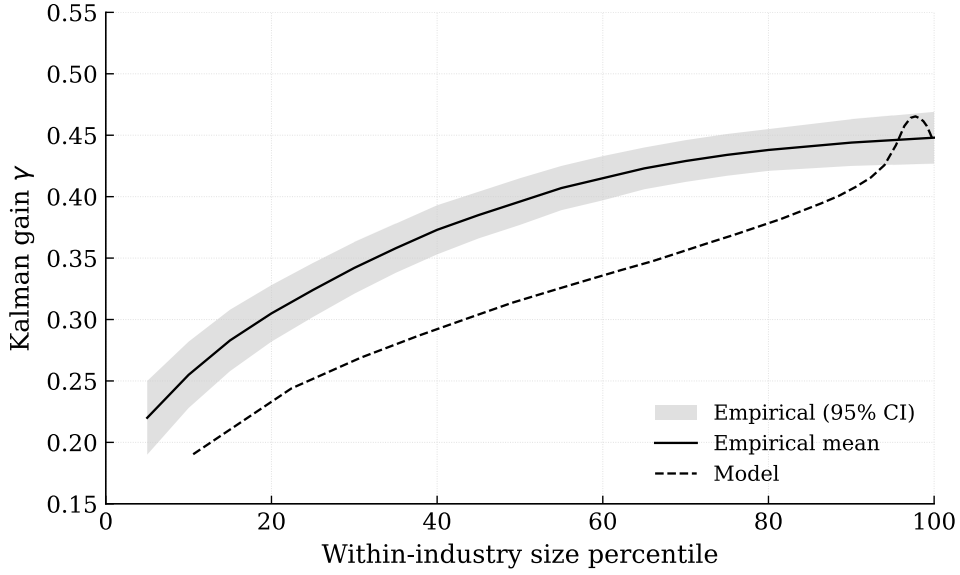


Figure 5: Kalman gain schedule: Model versus Data

4.4 Aggregate Price Rigidity

Using the calibrated model, I compare various nested models as in 3. To make the comparison transparent, I use the same three nested environments. First, the CES benchmark sets $\rho_i = 1$ for

all firms. In this case, only average Kalman gain matters and

$$\alpha^{\text{CES}} = E^s[\gamma_i].$$

Second, the homogeneous-attention economy restores the calibrated distributions of ρ_i and η_i, but assigns every firm the common sales-weighted Kalman gain $\bar{\gamma} = E^s[\gamma_i]$. Denote the resulting price response by α^R. Finally, the baseline restores the endogenous cross-sectional distribution of γ_i and yields α. Thus,

$$\alpha^{\text{CES}} - \alpha = \underbrace{\left(\alpha^{\text{CES}} - \alpha^R\right)}_{\text{real rigidity}} + \underbrace{\left(\alpha^R - \alpha\right)}_{\text{negative covariance channel}}. \tag{65}$$

Table 8: Price Rigidity Relative to the CES Benchmark

	$\Delta\alpha$ (%)	Phillips-curve flattening (%)
Homogeneous attention, $\gamma_{ij} = \bar{\gamma}_j$	-22.75	33.19
Additional negative covariance channel	-4.66	5.62
Baseline	-27.61	38.81

Notes. Changes are measured relative to the CES benchmark, which sets $\rho_{ij} = 1$. The homogeneous-attention economy restores the calibrated distributions of pass-through and elasticity but assigns all firms the common sales-weighted Kalman gain $\bar{\gamma}_j$. The negative covariance row reports the additional change obtained by restoring the endogenous within-sector distribution of Kalman gains. The baseline combines both margins.

The calibrated model generates magnitudes that are broadly comparable to those obtained directly from the empirical sufficient-statistics exercise. Relative to the CES benchmark, the calibrated economy implies a 38.81% flattening of the Phillips curve, compared with 39.14% in the data-based exercise. The additional contribution of the negative covariance channel is of a similar order of magnitude in both. Encouragingly, the calibrated model reproduces the quantitative importance of both the real-rigidity and the negative covariance between the information and attention channels.

With a calibrated model that closely reproduces the key features of the baseline economy, I now turn to the counterfactual exercise for which this model is particularly useful. Keeping demand and information-cost parameters fixed, I vary the underlying productivity distribution, which changes the level of market concentration, to trace how aggregate price rigidity moves with within-industry concentration.

Table 9: Concentration and price responsiveness across counterfactual economies

Pareto shape	Top 10%	Top 5%	Price response α
7.061, baseline	48.98	31.66	0.275
8, Economy 1	44.79	28.96	0.280
10, Economy 2	38.07	24.53	0.288

Notes: Sales shares are percentages within the representative industry. Only the Pareto shape changes; all other parameters remain fixed at baseline, including information costs. Firms' attention choices adjust endogenously.

4.5 Market Concentration Counterfactual

I now use the calibrated model to study how within-industry market concentration shapes monetary non-neutrality. I change only the Pareto tail index α_p, holding demand and information-cost parameters fixed, and re-solve firms' pricing and attention choices. Higher values of α_p produce a thinner tail and lower concentration. I consider two such economies: Economy 1, with $\alpha_p = 8$, and Economy 2, with $\alpha_p = 10$. In both cases, I describe the change *from the less concentrated economy back to baseline*, so concentration rises along the reported transition.

Concentration profile Table 9 shows that the top-10% sales share rises from 44.79% in Economy 1 (or 38.07% in Economy 2) to 48.98% at baseline. The table also shows that the aggregate price responsiveness α falls (or in other words prices become more rigid) as concentration rises from Economy 1 (or Economy 2) to the baseline. The change is mild, however.

To put these magnitudes in context, the experiments imply increases in the top-10% sales share of approximately 4-11 percentage points. These are comparable in scale to the increase in manufacturing and non-manufacturing CR8 concentration between the 1980s and 2010s reported in Covarrubias, Gutiérrez and Philippon (2020).

Three forces behind the price response Rising concentration affects price responsiveness α through changes in three forces that influence α, as noted in (60). First, sales shift toward larger, more attentive firms, raising average attention and tending to make aggregate prices more responsive to nominal shocks. Second, these larger firms now have lower cost pass-through: their desired prices move less when costs change. Giving them more weight therefore strengthens real rigidity, which dampens price responsiveness.

The third force concerns *who pays attention*, or the negative-covariance channel. Moving from a less concentrated economy (Economy 1 or Economy 2) to the baseline increases dispersion in firm size and pass-through. Because attention is endogenous, the attention gap between large and

small firms also widens. Larger firms are more attentive but have lower pass-through, so this greater dispersion makes the covariance between attention and pass-through more negative. The same distributional change that raises concentration and average attention also strengthens the negative-covariance channel, attenuating the resulting increase in aggregate price responsiveness. [12]

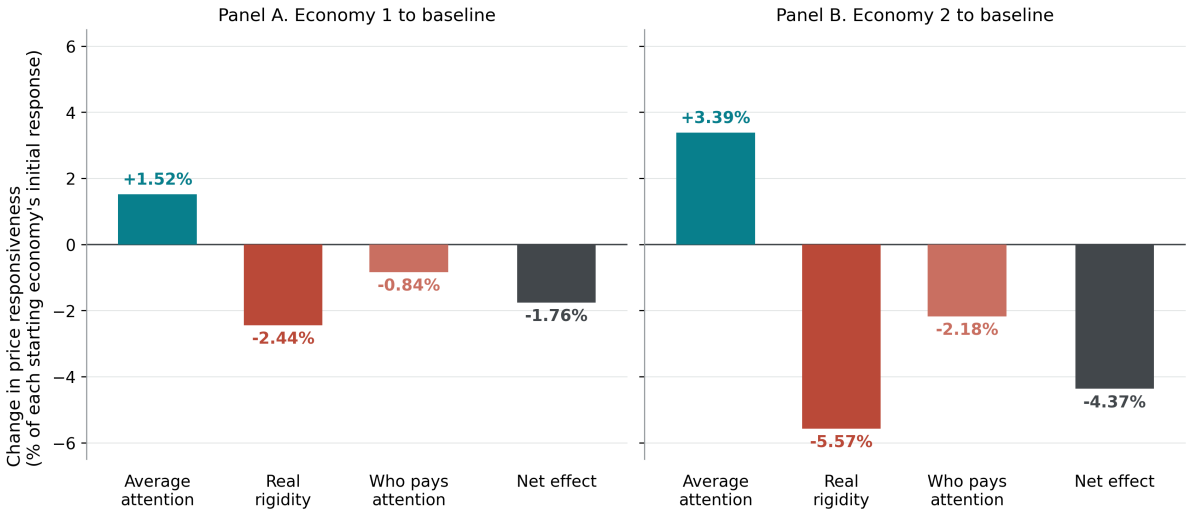


Figure 6: Why higher concentration lowers price responsiveness

Notes. Panel A shows Economy 1 ($\alpha_p = 8$) to baseline; Panel B shows Economy 2 ($\alpha_p = 10$) to baseline. Contributions are percentages of each starting economy's initial value. Positive values make prices more responsive and negative values less responsive.

Figure 6 illustrates the balance of these forces in the two counterfactual changes considered. For Economy 2, the increase in average attention contributes $+3.39\%$ to price responsiveness, but stronger real rigidity contributes -5.57%, and the change in who pays attention contributes a further -2.18%, all measured relative to the initial price response. The two negative forces more than offset the gain in attention, resulting in a net decline of 4.37% in α. The milder transition from Economy 1 yields the same pattern, with a net decline of 1.76%. This decline in price responsiveness implies greater monetary non-neutrality.

The key takeaway is that nominal rigidity responds endogenously to changes in the underlying firm distribution through all three forces. Greater concentration changes not only how much attention firms pay but also their price pass-through and the distribution of attention across firms with different pass-through. Understanding the resulting change in monetary non-neutrality therefore requires tracking the joint distribution of firm size, pass-through, and attention.

[12] In each economy, I compute the same decomposition reported in 8. This separates the real-rigidity and attention-allocation channels from the average attention channel. I then compare these channels across economies to understand their relative contributions.

5 Conclusion

This paper sheds light on a systematic pattern in the extent to which firms track aggregate nominal conditions. Using firms' 10-K filings, I construct a text-based measure of firms' attention to aggregate nominal conditions and map this measure into firm-level Kalman gains. A robust empirical pattern emerges: larger firms within narrowly defined industries are systematically more attentive to nominal conditions. This is a step forward in our understanding of how firm characteristics affect their information acquisition.

I develop a model of heterogeneous firms and endogenous information acquisition to understand this relationship and its aggregate consequences. Permanent productivity differences generate heterogeneity in firm size, while Kimball demand generates heterogeneity in demand elasticities and partial-equilibrium cost pass-through. Information is valuable because pricing mistakes are costly. The magnitude of these losses depends on firms' steady-state characteristics; larger firms optimally acquire more information. The model reveals that larger firms also have a lower pass-through; attention heterogeneity dampens the responsiveness of the aggregate price index and increases the degree of monetary non-neutrality. In the data, there is indeed a negative alignment between attention and pass-through. Accounting for this generates substantially more price rigidity than a benchmark where firm attention is homogeneous.

Broadly speaking, monetary non-neutrality therefore depends not only on how much firms know, but on which firms know it. This channel must be accounted for when studying the effects of changes in the underlying market structure and hence firm size distribution on effectiveness of monetary policy.

References

Afrouzi, Hassan. 2024. "Strategic Inattention, Inflation Dynamics, and the Nonneutrality of Money." *Journal of Political Economy*.

Afrouzi, Hassan, and Choongryul Yang. 2021. "Dynamic Rational Inattention and the Phillips Curve." CESifo CESifo Working Paper Series 8840.

Alexander, Patrick, Lu Han, Oleksiy Kryvtsov, and Ben Tomlin. forthcoming. "Markups and Inflation in Oligopolistic Markets: Evidence from Wholesale Price Data." *American Economic Journal: Macroeconomics*.

Amiti, Mary, Oleg Itskhoki, and Jozef Konings. 2019. "International Shocks, Variable Markups, and Domestic Prices." *The Review of Economic Studies*, 86(6): 2356–2402.

Angeletos, George-Marios, and Jennifer La'O. 2009. "Incomplete information, higher-order beliefs and price inertia." *Journal of Monetary Economics*, 56: S19–S37. Supplement issue: December 12-13, 2008 Research Conference on 'Monetary Policy under Imperfect Information' Sponsored by the Swiss National Bank (http://www.snb.ch) and Study Center Gerzensee (www.szgerzensee.ch).

Angeletos, George-Marios, Luigi Iovino, and Jennifer La'O. 2016. "Real Rigidity, Nominal Rigidity, and the Social Value of Information." *American Economic Review*, 106(1): 200–227.

Baqaee, David R., Emmanuel Farhi, and Kunal Sangani. 2024. "The Supply-Side Effects of Monetary Policy." *Journal of Political Economy*.

Bardóczy, Bence, Gideon Bornstein, and Sergio Salgado. 2026. "Monopsony Power and the Transmission of Monetary Policy." National Bureau of Economic Research Working Paper 35335.

Basu, Susanto. 2019. "Are Price-Cost Markups Rising in the United States? A Discussion of the Evidence." *Journal of Economic Perspectives*, 33(3): 3–22.

Burya, Anastasia, and Shruti Mishra. 2022. "Variable Markups, Demand Elasticity and Pass-through of Marginal Costs into Prices."

Candia, Bernardo, Olivier Coibion, and Yuriy Gorodnichenko. 2024. "The inflation expectations of U.S. firms: Evidence from a new survey." *Journal of Monetary Economics*, 145: 103569. Inflation: Expectations Dynamics October 14-15, 2022.

Coibion, Olivier, Yuriy Gorodnichenko, and Rupal Kamdar. 2018. "The Formation of Expectations, Inflation, and the Phillips Curve." *Journal of Economic Literature*, 56(4): 1447–91.

Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar. 2018. "How Do Firms Form Their Expectations? New Survey Evidence." *American Economic Review*, 108(9): 2671–2713.

Covarrubias, Matias, Germán Gutiérrez, and Thomas Philippon. 2020. "From Good to Bad Concentration? US Industries over the Past 30 Years." *NBER Macroeconomics Annual*, 34(1): 1–46.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *The Quarterly Journal of Economics*, 135(2): 561–644.

Dogra, Keshav, Sebastian Heise, Edward S. Knotek, Brent H. Meyer, Robert W. Rich, Raphael Schoenle, Giorgio Topa, Wilbert van der Klaauw, and Wändi Bruine de Bruin. 2023. "Estimates of Cost-Price Passthrough from Business Survey Data." Federal Reserve Bank of New York Staff Report 1062.

Dotsey, Michael, and Robert G. King. 2005. "Implications of state-dependent pricing for dynamic macroeconomic models." *Journal of Monetary Economics*, 52(1): 213–242.

Edmond, Chris, Virgiliu Midrigan, and Yi Xu. 2023. "How Costly Are Markups?" *Journal of Political Economy*, 131(7): 1839–1879.

Flynn, Joel P, and Karthik Sastry. 2024. "Attention Cycles." National Bureau of Economic Research Working Paper 32553.

Gagliardone, Luca, and Joris Tielens. 2026. "Dynamic Pricing under Information Frictions: Evidence from Firm-level Subjective Expectations." Working paper.

Golosov, Mikhail, and Robert E. Lucas Jr. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy*, 115(2): 171–199.

Hellwig, Christian, and Venky Venkateswaran. 2025. "Dispersed information, nominal rigidities and monetary business cycles: A Hayekian perspective." *Journal of Monetary Economics*, 155: 103843. Informational Frictions in Macroeconomics (In honor of Robert E. Lucas, Jr.).

Kimball, Miles S. 1995. "The Quantitative Analytics of the Basic Neomonetarist Model." *Journal of Money, Credit and Banking*, 27(4): 1241–1277.

Klenow, Peter J., and Jonathan L. Willis. 2016. "Real Rigidities and Nominal Price Changes." *Economica*, 83(331): 443–472.

Kohlhas, Alexandre N., and Ansgar Walther. 2021. "Asymmetric Attention." *American Economic Review*, 111(9): 2879–2925.

Kondo, Illenin O., Logan T. Lewis, and Andrea Stella. 2023. "Heavy tailed but not Zipf: Firm and establishment size in the United States." *Journal of Applied Econometrics*, 38(5): 767–785.

Lucas, Robert E. 1972. "Expectations and the neutrality of money." *Journal of Economic Theory*, 4(2): 103–124.

Mackowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review*, 99(3): 769–803.

Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *The Quarterly Journal of Economics*, 117(4): 1295–1328.

Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt. 2018. "Dynamic rational inattention: Analytical results." *Journal of Economic Theory*, 176: 650–692.

Mongey, Simon. 2021. "Market Structure and Monetary Non-neutrality." National Bureau of Economic Research Working Paper 29233.

Muth, Mary K., Samantha Bradley, Jenna E. Brophy, Kristen L. Capogrossi, Michaela C. Coglaiti, Shawn A. Karns, and Catherine L. Viator. 2015. "Reformulation Cost Model." RTI International. Final report prepared for the U.S. Food and Drug Administration, Center for Food Safety and Applied Nutrition.

Nimark, Kristoffer. 2008. "Dynamic pricing and imperfect common knowledge." *Journal of Monetary Economics*, 55(2): 365–382.

of Firms' Inflation Expectations (SoFIE), Survey. n.d.. "Fed Cleveland." Accessed October 29, 2026.

Pasten, Ernesto, Raphael Schoenle, and Michael Weber. 2020. "The propagation of monetary policy shocks in a heterogeneous production economy." *Journal of Monetary Economics*, 116: 1–22.

Rotemberg, Julio J., and Michael Woodford. 1996. "Imperfect Competition and the Effects of Energy Price Increases on Economic Activity." *Journal of Money, Credit and Banking*, 28(4): 549–577.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of Monetary Economics*, 50(3): 665–690. Swiss National Bank/Study Center Gerzensee Conference on Monetary Policy under Incomplete Information.

Song, Wenting, and Samuel Stern. 2025. "Firm Inattention and the Efficacy of Monetary Policy: A Text-Based Approach." *The Review of Economic Studies*, 92(5): 3438–3469.

Tanaka, Mari, Nicholas Bloom, Joel M. David, and Maiko Koga. 2020. "Firm performance and macro forecast accuracy." *Journal of Monetary Economics*, 114: 26–41.

Van Nieuwerburgh, Stijn, and Laura Veldkamp. 2006. "Learning asymmetries in real business cycles." *Journal of Monetary Economics*, 53(4): 753–772.

Wang, Olivier, and Iván Werning. 2022. "Dynamic Oligopoly and Price Stickiness." *American Economic Review*, 112(8): 2815–49.

Woodford, Michael. 2001. "Imperfect Common Knowledge and the Effects of Monetary Policy." National Bureau of Economic Research NBER Working Paper 8673.

Appendix

Who Pays Attention? Firm Size, Information, and Monetary Non-Neutrality

A Additional Checks on the Text-Based Measure

This appendix provides additional checks on the text-based measure used in Section 1. I first illustrate the language associated with different Stage-D precision scores. Then, I audit the content admitted by the Stage-C classifier manually to ensure accurate classification. I then show that the within-industry size gradient is unchanged when the measure is restricted to the narrowest nominal topics. Finally, I compare A_{it} with an independently constructed sentence-level measure of linguistic content and repeat this comparison on placebo topics.

A.1 Examples across precision scores

Recall that $p_{ij} = 1 - u_{ij}/10 \in [0, 1]$ denotes the passage-level precision score in (1). Table 10 reports representative sentences across five score bands. High-score passages tend to contain specific realized values, quantitative sensitivity exercises, or otherwise concrete statements. Low-score passages are more often conditional, heavily qualified, or explicit about the difficulty of characterizing future nominal conditions. The examples are meant to make transparent the textual variation captured by Stage D.

Band	Linguistic signature	Example sentences
Very high $p \in [0.9, 1.0]$ ($\approx 3\%$)	Numeric and factual; often reports specific realized values or quantified impacts.	*"This increase primarily resulted from significantly higher refined product prices, and to a lesser extent, an increase in our sales volume."* (refiner, 2008) *"The underlying junior subordinated debentures bore interest at a fixed rate of 6.98% paid quarterly until June 15, 2011 and then converted to floating rate (LIBOR plus 160 basis points) after June 15, 2011."* (financial, 2016)
High $p \in [0.7, 0.8]$ ($\approx 23\%$)	Quantitative sensitivity analysis or other numerically specific statements.	*"The gross margin rate increased due to favorable product category mix and margin improvement initiatives (70 bps), partially offset by lower product cost inflation (30 bps)."* (retailer, 2014)
Medium $p \in [0.5, 0.6]$ ($\approx 5\%$)	Directional or conditional statements with some concrete content but limited commitment to an outcome.	*"QNB actively manages the investment portfolio in an attempt to maximize earnings, while considering liquidity needs, interest rate risk, and credit risk."* (community bank, 2016)
Low $p \in [0.3, 0.4]$ ($\approx 45\%$)	Conditional or forward-looking language that emphasizes uncertainty about the nominal environment.	*"The economic environment is continuously changing, due to increased inflation, higher interest rates, global unrest, the war in Ukraine, the political environment, and trade issues."* (bank, 2023) *"If we are unable to enter into hedge agreements with respect to, or otherwise refinance, this indebtedness with acceptable terms, the ultimate impact of future interest rate increases could result in unanticipated reductions in our net operating income."* (REIT, 2023)
Very low $p \in [0.0, 0.2]$ ($\approx 22\%$)	Explicit statements that the relevant nominal conditions or their effects cannot be predicted.	*"The nature and impact of future changes in fiscal and monetary policies cannot be predicted."* (holding company, 2010) *"If business and economic conditions weaken generally or specifically in the principal markets in which we do business, more of our borrowers may fail to perform in accordance with the terms of their loans."* (bank, 2022)

Table 10: Representative passages across five Stage-D precision bands. The precision score is $p_{ij} = 1 - u_{ij}/10$, where u_{ij} is the Stage-D uncertainty score.

A.2 Hand-coded audit of Stage-C passages

I next examine the content admitted by the Stage-C classifier. I draw a random sample of 200 Stage C passages and hand-code each passage into one of eight categories: policy and market interest rates; inflation and the aggregate price level; monetary-policy stance; nominal credit and financial conditions; nominal exchange rates; broader macroeconomic commentary; firm-specific input costs or regulatory pricing; and off-topic content.

Table 11 reports the resulting composition. Sixty percent of the sampled passages concern narrowly nominal variables. Another 14% contain broader macroeconomic commentary, while 20% concern firm-specific input costs or regulatory pricing and 6% are off-topic. The audit therefore confirms that the broad Stage-C measure contains some content beyond the narrowest nominal

concepts. The following subsection checks whether this matters for the main size-gradient result.

Code	Topic category	Share (%)
1	Policy and market interest rates (fed funds, LIBOR / SOFR, yields)	32
2	Inflation / aggregate price level (CPI, PPI, PCE)	10
3	Monetary-policy stance (Fed / FOMC, QE / QT, forward guidance)	8
4	Nominal credit and financial conditions	6
5	Nominal exchange rates	4
	Subtotal: narrowly nominal (codes 1–5)	60
6	Broader macro / "economic conditions" commentary	14
7	Firm-specific input costs or regulatory pricing	20
8	Off-topic	6
	Total ($N = 50$)	100

Table 11: Hand-coded composition of a random sample of 50 Stage-C-admitted passages.

A.3 Narrow-topic robustness of the size gradient

Because Stage C intentionally uses a broad definition of aggregate nominal conditions, I repeat the main size-gradient exercise using only passages assigned to the two cleanest concept labels: `monetary_policy` and `inflation_price_level`. Let A_{it}^{narrow} denote the resulting filing-level measure, constructed using the same denominator $|\mathcal{B}_{it}|$ as the baseline index.

Figure 7 shows that the within-industry size gradient is preserved under this restriction. The measure rises steadily through most of the size distribution, peaks around the 80th percentile, and declines only modestly at the very top. The firm-year correlation between A_{it}^{narrow} and the baseline A_{it} is 0.85. Hence the main size-gradient result does not depend on the broader `nominal_cost_environment` category.

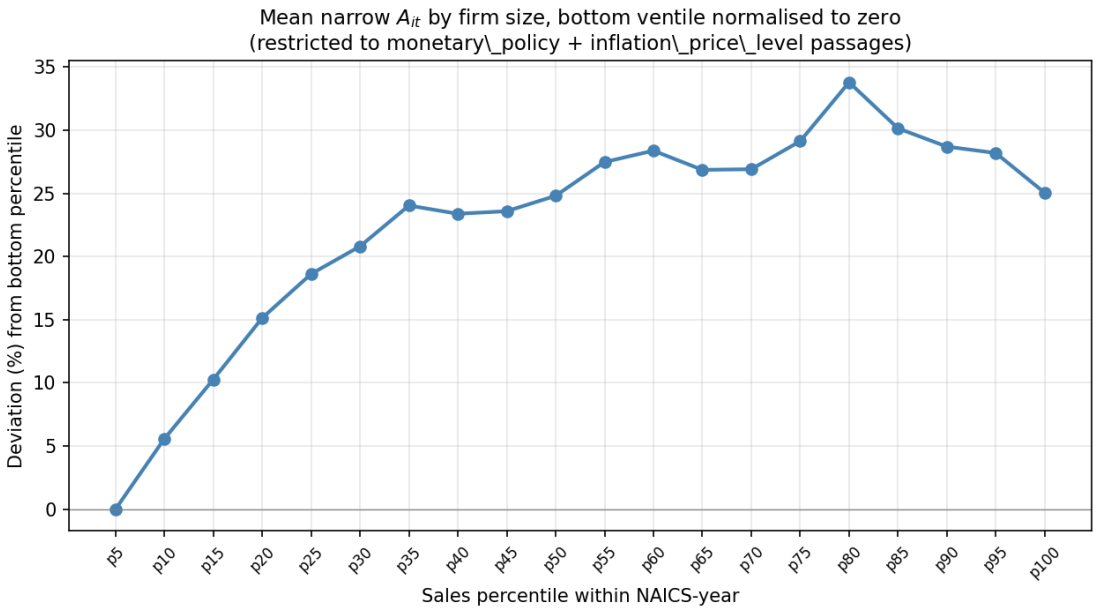


Figure 7: Mean A_{it}^{narrow}, restricted to passages labelled `monetary_policy` or `inflation_price_level`, residualized on NAICS 6-digit × fiscal-year fixed effects and plotted by within-industry revenue-share ventile. The bottom ventile is normalized to zero.

A.4 Independent sentence-level check

I construct a separate sentence-level measure to check whether the variation in A_{it} is associated with independently coded features of the underlying language. The exercise uses Items 1A, 7, and 7A and operates at the sentence level rather than on the Stage-C paragraphs used to construct A_{it}. A fine-tuned `gpt-4o-mini` classifier assigns each sentence to a macroeconomic or idiosyncratic bucket, records a topic labels including ..... The classifier also separately codes a small set of binary textual markers.

The markers capture simple features of content: whether a sentence contains a directional statement, a numeric value, an explicit time reference, a causal or conditional mechanism, a factual assertion, ambiguity about whether the content is macroeconomic or idiosyncratic, and generic boilerplate content. I also form a `concrete_forecast` indicator equal to one when a sentence contains direction, a numeric value, and a time reference. The label is mechanical; it is intended to identify relatively concrete, time-specific statements rather than to infer a firm's posterior directly.

For each firm-year, I compute the share of macro-nominal sentences carrying each marker, where macro-nominal sentences are those assigned to monetary policy, inflation, or financial conditions.

Figure 8 summarizes the results. The markers associated with more concrete language rise with A_{it}, while ambiguity falls. These patterns provide an independent check that high-A_{it} filings contain more specific and time-referenced nominal discussion, rather than simply more nominal words.

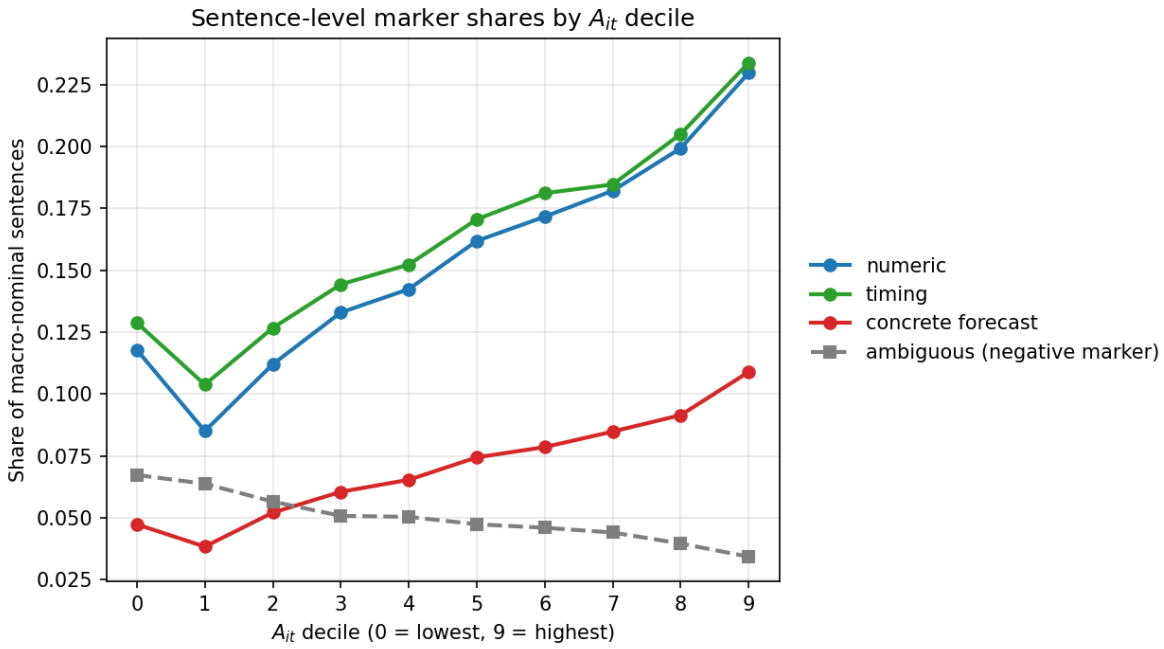


Figure 8: Sentence-level marker shares by A_{it} decile. The top panel reports markers associated with more concrete language; the bottom panel reports ambiguity and boilerplate.

A.5 Placebo topics

A remaining concern is that A_{it} may proxy for generic disclosure quality: firms with high A_{it} may simply write more concretely about every topic. I therefore repeat the sentence-level exercise separately for several non-nominal topics.

Figure 9 reports the results. The relation is essentially absent for technology operations and geopolitics: on these topics A_{it} shows almost no correlation with the numeric, timing, or concrete-statement markers. Labor shows smaller correlations with the numeric and timing markers. The macro-nominal category has the strongest correlation with the numeric and concrete-statement markers, as well as a similarly strong timing relationship. The placebo exercise therefore suggests that high A_{it} firms specifically show concrete and high-quality disclosure in nominal-relevant topics.

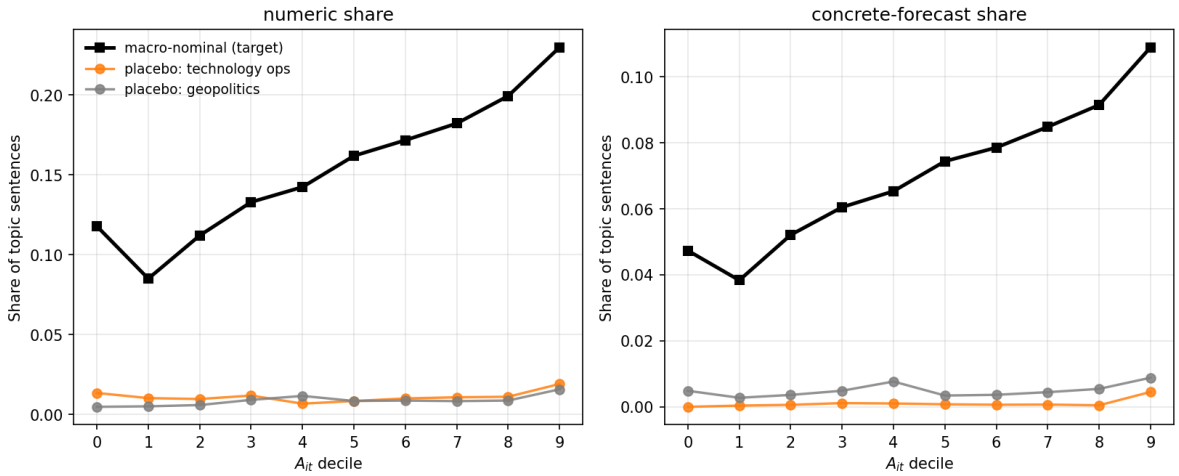


Figure 9: Sentence-level marker shares by A_{it} decile for the macro-nominal category and four comparison topics.

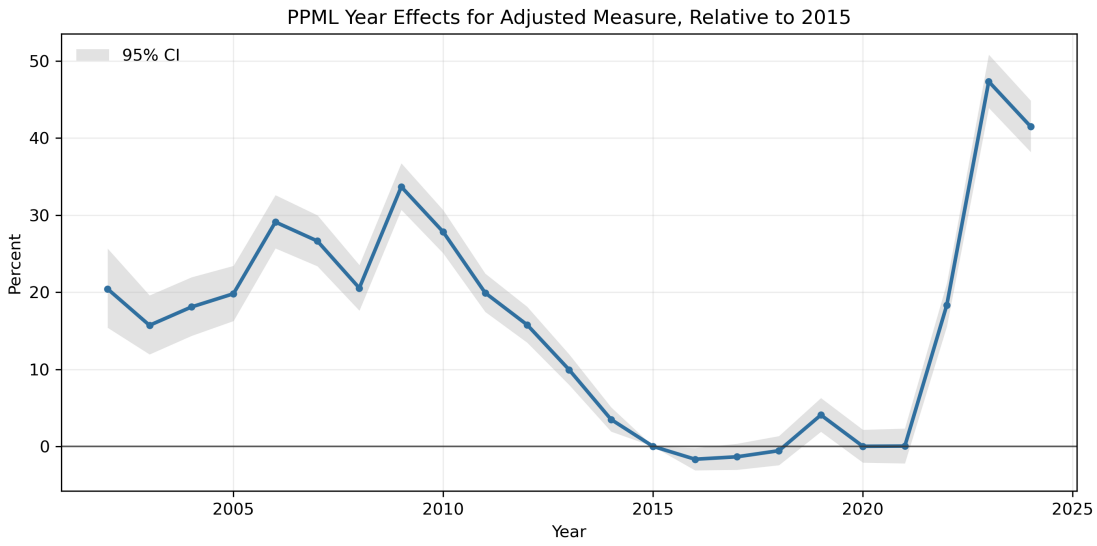
B Text-Based Measure Sensibility Checks

In this appendix, I show that the measured text-based attention meaningfully captures information about price changes and is correlated with price rigidity and inflation exposure at the industry level. For this analysis, I restrict the sample to firms with at least five observed years, yielding a sample of 68,388 firm-year observations from 5,283 firms.

For the time-series variation, I first residualize the measured attention by firm size and industry fixed effects and then plot annual averages. To do this, I estimate a PPML specification to obtain year effects in percentage units:

$$E[T_{it} \mid X_{it}] = \exp\left(\sum_{\tau \neq 2015} \beta_\tau 1[t = \tau] + \delta_{\text{size}(it)} + \gamma_{\text{NAICS6}(i)}\right).$$

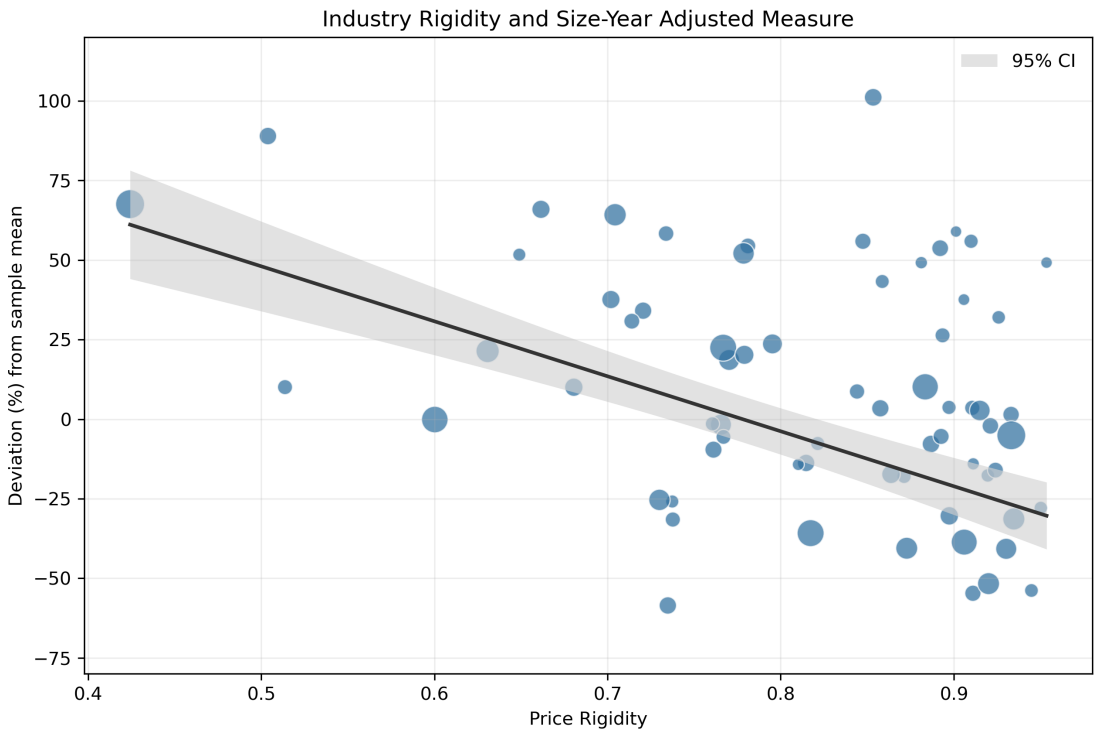
Figure **??** shows that attention is relatively close from 2002 to 2009, a period characterized by generally stable, low inflation. Then it starts declining over the next decade, at the onset of the Great Recession. The series suddenly spikes post-2021, a period that was clearly one of high inflation uncertainty (and level too). The reported coefficients are transformed to percentage deviations relative to 2015 using $100(\exp(\beta_\tau) - 1)$.



Next, I show how the measure actually captures the information contained in industry-level price changes and inflation. Specifically, I ask whether industries with more rigid prices have higher or lower text attention. I use the estimates of industry-level price rigidity from Pasten, Schoenle and Weber (2020). This paper uses micro-data underlying the producer price index (PPI) to compute the frequency of price adjustment at the industry level. I then calculate the rigidity of each industry j as:

$$\text{Rigidity}_j = 1 - \text{Frequency of Price Changes}_j.$$

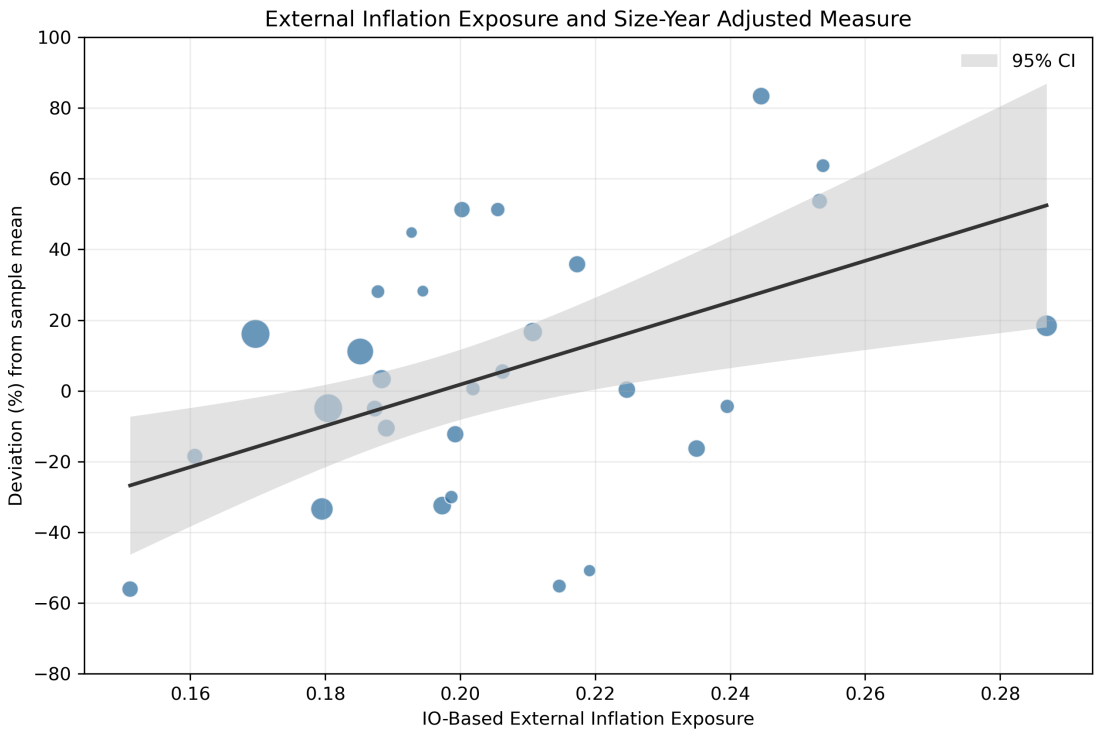
The measure is residualized by firm size and year fixed effects to obtain industry-level variation. In Figure **??**, the plot groups observations to the 3-digit NAICS level and reports each industry's mean adjusted text measure as a percent deviation from the matched-sample mean. Point size is proportional to industry sales. The plot also shows the fitted line with a coefficient of -0.0814 and a standard error of 0.0097 ($t = -8.39$). Thus, more rigid industries have lower adjusted text-measure levels after controlling for size deciles and year fixed effects.



I then ask whether industries that are more exposed to inflation pressure from their connected industries have higher or lower adjusted text measures. For each row industry, external exposure is the Input-Output table weighted average of connected industries' price-flexibility values:

$$\text{Exposure}_j = \frac{\sum_{k \neq j} w_{jk} \, \text{Flexibility}_k}{\sum_{k \neq j} w_{jk}}.$$

As in the rigidity plot, the measure is residualized with respect to firm size and year fixed effects, averaged to the IO-code level, and expressed as a percent deviation from the matched-sample mean. Point size is proportional to industry sales. The fitted line is a weighted linear fit with a 95 percent confidence interval. The estimated exposure coefficient is 0.2791 with standard error 0.1123 ($t = 2.48$). In contrast to the rigidity diagnostic, industries with greater external inflation exposure have higher adjusted text-measure levels in this descriptive specification.



C Proofs

C.1 Proof of Lemma 1

Proof of Lemma 1

In the deterministic steady state, firm ij's optimal price satisfies

$$P_{ij}^{*} = \mu_{ij} \frac{W^{*}}{Z_{ij}}, \qquad \mu_{ij} = \frac{\eta_{ij}}{\eta_{ij} - 1}. \tag{66}$$

From the definition of the superelasticity,

$$\frac{d \log \eta_{ij}}{d \log\left(P_{ij} / (\lambda_j P_j)\right)} = \xi_{ij}.$$

Since $\mu_{ij} = \eta_{ij} / (\eta_{ij} - 1)$,

$$\frac{d \log \mu_{ij}}{d \log\left(P_{ij} / (\lambda_j P_j)\right)} = -\frac{\xi_{ij}}{\eta_{ij} - 1}. \tag{67}$$

Log-differentiating (66) with respect to Z_{ij} therefore gives

$$\frac{d\log\left(P_{ij}/(\lambda_j P_j)\right)}{d\log Z_{ij}} = -1 - \frac{\xi_{ij}}{\eta_{ij}-1}\frac{d\log\left(P_{ij}/(\lambda_j P_j)\right)}{d\log Z_{ij}}.$$

Hence

$$\frac{d\log\left(P_{ij}/(\lambda_j P_j)\right)}{d\log Z_{ij}} = -\frac{\eta_{ij}-1}{\eta_{ij}-1+\xi_{ij}} = -\rho_{ij} < 0. \tag{68}$$

Higher-productivity firms therefore charge lower prices relative to the sectoral Kimball price index.

Firm demand satisfies

$$y_{ij} = Y_j F\left(\frac{P_{ij}}{\lambda_j P_j}\right).$$

Holding sectoral aggregates fixed, firm revenue therefore satisfies

$$\frac{d\log(P_{ij}y_{ij})}{d\log\left(P_{ij}/(\lambda_j P_j)\right)} = 1 - \eta_{ij} < 0.$$

Combining this with (68),

$$\frac{d\log s^*_{ij|j}}{d\log Z_{ij}} = \rho_{ij}(\eta_{ij}-1) > 0. \tag{69}$$

Thus higher-productivity firms are larger within their sectors.

Next,

$$\frac{d\log\eta_{ij}}{d\log Z_{ij}} = \xi_{ij}\frac{d\log\left(P_{ij}/(\lambda_j P_j)\right)}{d\log Z_{ij}} = -\xi_{ij}\rho_{ij} < 0, \tag{70}$$

so larger, more productive firms face lower demand elasticities. Since $\mu_{ij} = \eta_{ij}/(\eta_{ij}-1)$ is decreasing in η_{ij}, they therefore charge higher markups.

Finally, partial-equilibrium cost pass-through is

$$\rho_{ij} = \frac{\eta_{ij}-1}{\eta_{ij}-1+\xi_{ij}} = \left[1+\frac{\xi_{ij}}{\eta_{ij}-1}\right]^{-1}. \tag{71}$$

Assumption 1 implies that $\xi_{ij}/(\eta_{ij}-1)$ is decreasing in $P_{ij}/(\lambda_j P_j)$. Since (68) shows that higher productivity lowers this relative price, higher-productivity firms have a larger value of $\xi_{ij}/(\eta_{ij}-1)$ and therefore a lower ρ_{ij}.

Combining these results, higher-productivity firms are larger within their sectors, face lower demand elasticities, charge higher markups, and have lower partial-equilibrium cost pass-through.
□

C.2 Second-order approximation and proof for Proposition 1

This appendix derives the local quadratic pricing problem and the full-information target used in the main text. The approximation is taken around the deterministic asymmetric steady state.

Firm ij's operating profits, valued in household marginal-utility units, are

$$\widetilde{\pi}_{ij,t} \equiv \frac{1}{Q_t}\left(P_{ij,t} - \frac{W_t}{Z_{ij}}\right) y_{ij,t}, \tag{72}$$

where

$$y_{ij,t} = Y_{j,t} F\left(\frac{P_{ij,t}}{\lambda_{j,t} P_{j,t}}\right). \tag{73}$$

Recall also that

$$s^*_{ij} \equiv \frac{P^*_{ij} y^*_{ij}}{P^* Y^*}.$$

Second-order approximation. Differentiate normalized profits with respect to the firm's own log price. Using

$$\frac{\partial \log y_{ij,t}}{\partial \log P_{ij,t}} = -\eta_{ij,t},$$

we obtain

$$\frac{\partial \widetilde{\pi}_{ij,t}}{\partial \widehat{P}_{ij,t}} = \frac{P_{ij,t} y_{ij,t}}{Q_t}\left[1 - \eta_{ij,t}\left(1 - \frac{W_t / Z_{ij}}{P_{ij,t}}\right)\right]. \tag{74}$$

At the steady-state optimum this derivative is zero, implying

$$\frac{W^* / Z_{ij}}{P^*_{ij}} = \frac{\eta_{ij} - 1}{\eta_{ij}}, \qquad P^*_{ij} = \frac{\eta_{ij}}{\eta_{ij} - 1}\frac{W^*}{Z_{ij}}. \tag{75}$$

To obtain the local curvature, differentiate (74) once more with respect to $\widehat{P}_{ij,t}$. At the steady-state optimum, the term in square brackets in (74) is zero, so only its derivative matters. Since

$$\frac{\partial \log \eta_{ij,t}}{\partial \widehat{P}_{ij,t}} = \xi_{ij}, \qquad \frac{\partial}{\partial \widehat{P}_{ij,t}}\left(\frac{W_t / Z_{ij}}{P_{ij,t}}\right) = -\frac{W_t / Z_{ij}}{P_{ij,t}},$$

equation (75) gives

$$\begin{aligned}\left.\frac{\partial^2 \widetilde{\pi}_{ij,t}}{\partial \widehat{P}^2_{ij,t}}\right|_* &= -s^*_{ij}\left[\xi_{ij} + (\eta_{ij} - 1)\right] \\ &= -s^*_{ij}\frac{\eta_{ij} - 1}{\rho_{ij}},\end{aligned} \tag{76}$$

where

$$\rho_{ij} = \frac{\eta_{ij} - 1}{\eta_{ij} - 1 + \xi_{ij}}.$$

We next compute the cross derivatives that determine how the desired price moves with the state. Using (75),

$$\left.\frac{\partial^2 \widetilde{\pi}_{ij,t}}{\partial \widehat{P}_{ij,t} \partial \widehat{W}_t}\right|_* = s^*_{ij}(\eta_{ij} - 1) = s^*_{ij}\frac{\eta_{ij} - 1}{\rho_{ij}}\rho_{ij}. \tag{77}$$

A higher $\widehat{P}_{j,t}$ or $\widehat{\lambda}_{j,t}$ lowers the firm's relative price $P_{ij,t}/(\lambda_{j,t}P_{j,t})$ one-for-one. Hence

$$\frac{\partial \log \eta_{ij,t}}{\partial \widehat{P}_{j,t}} = \frac{\partial \log \eta_{ij,t}}{\partial \widehat{\lambda}_{j,t}} = -\xi_{ij}.$$

It follows that

$$\left.\frac{\partial^2 \widetilde{\pi}_{ij,t}}{\partial \widehat{P}_{ij,t} \partial \widehat{P}_{j,t}}\right|_* = \left.\frac{\partial^2 \widetilde{\pi}_{ij,t}}{\partial \widehat{P}_{ij,t} \partial \widehat{\lambda}_{j,t}}\right|_* = s^*_{ij}\xi_{ij} = s^*_{ij}\frac{\eta_{ij} - 1}{\rho_{ij}}(1 - \rho_{ij}). \tag{78}$$

The sectoral demand shifter does not enter the local target price. It scales the entire first derivative (74), which is zero at the steady-state optimum. The same is true of the normalization $1/Q_t$.

Collecting the terms of the second-order expansion that depend on the firm's chosen price gives

$$\begin{aligned}\widetilde{\pi}_{ij,t} \simeq \text{terms independent of } \widehat{P}_{ij,t} &- \frac{1}{2}s^*_{ij}\frac{\eta_{ij} - 1}{\rho_{ij}}\widehat{P}^2_{ij,t} \\ &+ s^*_{ij}\frac{\eta_{ij} - 1}{\rho_{ij}}\rho_{ij}\widehat{P}_{ij,t}\widehat{W}_t \\ &+ s^*_{ij}\frac{\eta_{ij} - 1}{\rho_{ij}}(1 - \rho_{ij})\widehat{P}_{ij,t}\left(\widehat{P}_{j,t} + \widehat{\lambda}_{j,t}\right).\end{aligned} \tag{79}$$

Completing the square,

$$\widetilde{\pi}_{ij,t} \simeq \text{terms independent of } \widehat{P}_{ij,t} - \frac{1}{2}s^*_{ij}\frac{\eta_{ij} - 1}{\rho_{ij}}\left(\widehat{P}_{ij,t} - \widehat{P}^f_{ij,t}\right)^2, \tag{80}$$

where the firm's full-information target is

$$\widehat{P}^f_{ij,t} = \rho_{ij}\widehat{W}_t + (1 - \rho_{ij})\left(\widehat{P}_{j,t} + \widehat{\lambda}_{j,t}\right). \tag{81}$$

The coefficient $s^*_{ij}(\eta_{ij} - 1)/\rho_{ij}$ is the curvature of normalized profits around that target.

Optimal pricing conditional on information For any information structure, after observing its signal the firm chooses $\widehat{P}_{ij,t}$ to minimize

$$E\left[\left(\widehat{P}_{ij,t}-\widehat{P}^{f}_{ij,t}\right)^{2}\,\middle|\,\mathcal{I}_{ij,t}\right].$$

The first-order condition is therefore

$$E\left[\widehat{P}_{ij,t}-\widehat{P}^{f}_{ij,t}\,\middle|\,\mathcal{I}_{ij,t}\right]=0,$$

which gives

$$\widehat{P}_{ij,t}=E_{ij,t}\left[\widehat{P}^{f}_{ij,t}\right]. \tag{82}$$

Using (81),

$$\widehat{P}_{ij,t}=\rho_{ij}E_{ij,t}[\widehat{W}_t]+(1-\rho_{ij})E_{ij,t}\left[\widehat{P}_{j,t}+\widehat{\lambda}_{j,t}\right]. \tag{83}$$

Finally, under the linear equilibrium conjecture

$$\widehat{P}_{j,t}=\alpha_j\widehat{W}_t,\qquad \widehat{\lambda}_{j,t}=\ell_j\widehat{W}_t,$$

the optimal price becomes

$$\widehat{P}_{ij,t}=\left[\rho_{ij}+(1-\rho_{ij})(\alpha_j+\ell_j)\right]E_{ij,t}[\widehat{W}_t]. \tag{84}$$

This establishes the optimal-pricing result used in the main text. □

C.3 Proof of Lemma 2

Conditional on the equilibrium coefficients, the full-information target is $\widehat{P}^{f}_{ij,t}=\phi_{ij}\widehat{W}_t$. Thus, for a given capacity κ_{ij}, the information problem is equivalent to minimizing the mean-squared error in tracking $\widehat{W}_t$ subject to the information constraint. *Step 1: The problem can be chosen to be memoryless.*

Because $\widehat{W}_t$ is i.i.d.,

$$\widehat{W}_t\perp\mathcal{I}_{ij,t-1},\qquad \widehat{W}_t\mid\mathcal{I}_{ij,t-1}\sim N(0,\sigma^2_W).$$

Hence, after every history, the firm faces the same one-period problem:

$$\min_{\pi_{ij,t}} E\Big[\big(\widehat{W}_t - E[\widehat{W}_t \mid S_{ij,t}, \mathcal{I}_{ij,t-1}]\big)^2\Big] \quad \text{s.t.} \quad I(\widehat{W}_t; S_{ij,t} \mid \mathcal{I}_{ij,t-1}) \leq \kappa_{ij}.$$

and

$$S_{ij,t} \sim \pi_{ij,t}\Big[. \mid \widehat{W}_t; \mathcal{I}_{ij,t-1})\Big]$$

Thus, the memoryless information structure. Below we suppress time and firm subscripts. *Step 2: A scalar signal is sufficient.* Let S be any (possibly multidimensional) signal and define its posterior mean

$$M \equiv E[\widehat{W} \mid S].$$

The pricing decision depends on S only through M. Moreover, by the law of iterated expectations,

$$E[\widehat{W} \mid M] = E\Big[E[\widehat{W} \mid S] \mid M\Big] = M,$$

so replacing S by the scalar statistic M leaves the mean-squared error unchanged. It also uses weakly less information. Since $M = m(S)$ is a deterministic function of S, the pair (S, M) contains exactly the same information as S, and the chain rule for mutual information gives

$$I(\widehat{W}; S) = I(\widehat{W}; S, M) = I(\widehat{W}; M) + I(\widehat{W}; S \mid M).$$

Conditional mutual information is nonnegative, so

$$I(\widehat{W}; M) \leq I(\widehat{W}; S).$$

Thus, any admissible multidimensional signal can be replaced by a scalar posterior mean that yields the same pricing decision with weakly less information. *Step 3: A Gaussian scalar information structure is optimal* Let

$$D \equiv E\Big[Var(\widehat{W} \mid S)\Big] = E\Big[\big(\widehat{W} - E[\widehat{W} \mid S]\big)^2\Big]$$

denote the posterior mean-squared error. Since $\widehat{W} \sim N(0, \sigma_W^2)$,

$$h(\widehat{W}) = \frac{1}{2}\log_2(2\pi e \sigma_W^2).$$

For each signal realization $S = s$, the Gaussian distribution maximizes entropy among all distributions with a given variance. Therefore

$$h(\widehat{W} \mid S = s) \leq \frac{1}{2}\log_2\Big(2\pi e\, Var(\widehat{W} \mid S = s)\Big).$$

Taking expectations over S,

$$\begin{aligned} h(\widehat{W} \mid S) &\leq \frac{1}{2}E\Big[\log_2\Big(2\pi e\, Var(\widehat{W} \mid S)\Big)\Big] \\ &\leq \frac{1}{2}\log_2\Big(2\pi e\, E[Var(\widehat{W} \mid S)]\Big) \\ &= \frac{1}{2}\log_2(2\pi e D). \end{aligned}$$

The second inequality is Jensen's inequality: $\log(\cdot)$ is concave, so $E[logX] \leq logE[X]$. Hence

$$I(\widehat{W}; S) = h(\widehat{W}) - h(\widehat{W} \mid S) \geq \frac{1}{2}\log_2\left(\frac{\sigma_W^2}{D}\right). \tag{85}$$

Combining (85) with $I(\widehat{W}; S) \leq \kappa$ implies

$$D \geq \sigma_W^2 2^{-2\kappa}. \tag{86}$$

Now consider the additive Gaussian signal

$$S = \widehat{W} + \varepsilon, \qquad \varepsilon \sim N(0, \sigma_\varepsilon^2), \qquad \varepsilon \perp \widehat{W}$$

For this signal, the conditional distribution of $\widehat{W}$ given S is Gaussian and its conditional variance is constant across signal realizations. Consequently, both inequalities above hold with equality. In particular,

$$I(\widehat{W}; S) = \frac{1}{2}\log_2\left(1 + \frac{\sigma_W^2}{\sigma_\varepsilon^2}\right), \qquad Var(\widehat{W} \mid S) = \frac{\sigma_W^2 \sigma_\varepsilon^2}{\sigma_W^2 + \sigma_\varepsilon^2}.$$

Choosing

$$\sigma_\varepsilon^2 = \frac{\sigma_W^2}{2^{2\kappa} - 1}$$

for $\kappa > 0$ makes the information constraint bind and attains the lower bound in (86). Thus an optimal scalar information structure can be represented by an additive Gaussian signal. (The scalar posterior mean from Step 2 is not itself $\widehat{W} + \varepsilon$; under this representation it is a rescaling of

that signal.) Since the same static problem is repeated independently over time, the signal errors can be taken to be independent across dates.

Step 4: Posterior gain and capacity. Joint Gaussianity implies

$$E[\widehat{W} \mid S] = \frac{\sigma_W^2}{\sigma_W^2 + \sigma_\varepsilon^2} S \equiv \gamma S$$

Using the binding information constraint,

$$\kappa = \frac{1}{2} \log_2 \left(1 + \frac{\sigma_W^2}{\sigma_\varepsilon^2} \right),$$

so

$$\gamma = \frac{\sigma_W^2}{\sigma_W^2 + \sigma_\varepsilon^2} = 1 - 2^{-2\kappa}$$

The limiting case $\kappa = 0$ corresponds to an uninformative signal ($\sigma_\varepsilon^2 \to \infty$).

C.4 Proof of Proposition 2

Recall that the second-order approximation to the firm's can be written as

$$\widetilde{\pi}_{ij,t} = C_{ij,t} - \frac{1}{2} s_{ij}^* \frac{\eta_{ij} - 1}{\rho_{ij}} \left(\widehat{P}_{ij,t} - \widehat{P}_{ij,t}^f \right)^2,$$

where $C_{ij,t}$ is independent of the firm's price choice. Given its information, optimal pricing implies

$$\widehat{P}_{ij,t} = E_{ij,t} \left[\widehat{P}_{ij,t}^f \right].$$

Hence the expected quadratic pricing loss is

$$\begin{aligned} &\frac{1}{2} s_{ij}^* \frac{\eta_{ij} - 1}{\rho_{ij}} E \left[\left(\widehat{P}_{ij,t}^f - E_{ij,t}[\widehat{P}_{ij,t}^f] \right)^2 \right] \\ &\quad = \frac{1}{2} s_{ij}^* \frac{\eta_{ij} - 1}{\rho_{ij}} E \left[\mathrm{Var} \left(\widehat{P}_{ij,t}^f \mid s_{ij,t} \right) \right]. \end{aligned} \tag{87}$$

Also note,

$$\widehat{P}_{ij,t}^f = \phi_{ij} \widehat{W}_t, \qquad \phi_{ij} \equiv \rho_{ij} + (1 - \rho_{ij})(\alpha_j + \ell_j).$$

The optimal information structure implies that, for capacity κ_{ij},

$$\mathrm{Var}(\widehat{W}_t \mid s_{ij,t}) = \sigma_W^2 2^{-2\kappa_{ij}}.$$

It follows that

$$\mathrm{Var}(\widehat{P}^f_{ij,t} \mid s_{ij,t}) = \phi_{ij}^2 \sigma_W^2 2^{-2\kappa_{ij}}.$$

Therefore, the firm's capacity problem is

$$\min_{\kappa_{ij} \geq 0} \left\{ \frac{1}{2} \mathcal{V}_{ij} 2^{-2\kappa_{ij}} + \omega \kappa_{ij}^p \right\}, \tag{88}$$

where

$$\mathcal{V}_{ij} = s_{ij}^* \frac{\eta_{ij} - 1}{\rho_{ij}} \phi_{ij}^2 \sigma_W^2. \tag{89}$$

Part 1: Optimal capacity. For an interior solution, differentiating (88) with respect to κ_{ij} gives

$$-\ln(2) \mathcal{V}_{ij} 2^{-2\kappa_{ij}} + p\omega \kappa_{ij}^{p-1} = 0,$$

or

$$\ln(2) \mathcal{V}_{ij} 2^{-2\kappa_{ij}^*} = p\omega (\kappa_{ij}^*)^{p-1}. \tag{90}$$

For $p > 1$ and $\mathcal{V}_{ij} > 0$, the objective in (88) is strictly convex and the solution is interior and unique.

Monotonicity in $\mathcal{V}_{ij}$ is straightforward. The optimal Kalman gain is also strictly increasing in $\mathcal{V}_{ij}$.

Part 2: Size gradient. Fix $\alpha_j + \ell_j$ and consider variation across firms within sector j. From (89),

$$d \log \mathcal{V}_{ij} = d \log s_{ij}^* + d \log(\eta_{ij} - 1) - d \log \rho_{ij} + 2\, d \log \phi_{ij}. \tag{91}$$

Holding sectoral aggregates fixed, firm demand implies

$$\frac{d \log s_{ij}^*}{d \log \big(P_{ij} / (\lambda_j P_j) \big)} = 1 - \eta_{ij}. \tag{92}$$

Using the definition of the superelasticity,

$$\xi_{ij} = \frac{d \log \eta_{ij}}{d \log\big(P_{ij}/(\lambda_j P_j)\big)},$$

we also obtain

$$\frac{d \log(\eta_{ij} - 1)}{d \log\big(P_{ij}/(\lambda_j P_j)\big)} = \frac{\eta_{ij}}{\eta_{ij} - 1} \xi_{ij}. \tag{93}$$

Finally, since

$$\phi_{ij} = \rho_{ij} + (1 - \rho_{ij})(\alpha_j + \ell_j),$$

holding $\alpha_j + \ell_j$ fixed gives

$$\frac{d \log \phi_{ij}}{d \log\big(P_{ij}/(\lambda_j P_j)\big)} = \frac{[1 - (\alpha_j + \ell_j)]\rho_{ij}}{\rho_{ij} + (1 - \rho_{ij})(\alpha_j + \ell_j)} \frac{d \log \rho_{ij}}{d \log\big(P_{ij}/(\lambda_j P_j)\big)}. \tag{94}$$

Substituting (92)–(94) into (91) yields

$$\begin{aligned} \frac{d \log \mathcal{V}_{ij}}{d \log\big(P_{ij}/(\lambda_j P_j)\big)} &= 1 - \eta_{ij} + \frac{\eta_{ij}}{\eta_{ij} - 1} \xi_{ij} \\ &\quad + \left[2 \frac{[1 - (\alpha_j + \ell_j)]\rho_{ij}}{\rho_{ij} + (1 - \rho_{ij})(\alpha_j + \ell_j)} - 1 \right] \frac{d \log \rho_{ij}}{d \log\big(P_{ij}/(\lambda_j P_j)\big)}. \end{aligned} \tag{95}$$

Since $\eta_{ij} > 1$, (92) implies that firm size is decreasing in the relative price. Consequently, $\mathcal{V}_{ij}$ is locally increasing in firm size whenever the right-hand side of (95) is negative. If this condition holds throughout the relevant support, $\mathcal{V}_{ij}$ is monotonically increasing in firm size. Part 1 then implies that κ_{ij}^* and γ_{ij}^* are also monotonically increasing in firm size. □

C.5 Sectoral Fixed Point, Existence, and Uniqueness

This appendix derives the sectoral fixed-point equation used in the main text and establishes existence and a sufficient condition for uniqueness. The key object is the response of the sectoral strategic price index,

$$X_{j,t} \equiv \lambda_{j,t} P_{j,t},$$

to the nominal shock. The logic of the proof is simple. We first *conjecture* a response of $X_{j,t}$ to the nominal shock. Firms take that response as given when choosing information and prices. We then *aggregate* those individual pricing decisions to obtain the collective response of $X_{j,t}$. An equilibrium is a conjecture that reproduces itself.

Fix a sector j and define

$$x \equiv \alpha_j + \ell_j, \qquad \phi_{ij}(x) \equiv \rho_{ij} + (1 - \rho_{ij})x. \tag{96}$$

Under the conjecture

$$\widehat{X}_{j,t} = x\widehat{W}_t, \tag{97}$$

firm ij's full-information target price satisfies

$$\widehat{P}^f_{ij,t} = \phi_{ij}(x)\widehat{W}_t. \tag{98}$$

The coefficient $\phi_{ij}(x)$ is the firm's general-equilibrium target exposure to the nominal shock. It combines the direct marginal-cost channel, ρ_{ij}, with the response through the sectoral strategic index, $(1 - \rho_{ij})x$.

Let $\kappa^*_{ij}(x)$ denote the firm's optimal information capacity conditional on the conjectured response x, and define the associated Kalman gain

$$\gamma^*_{ij}(x) \equiv 1 - 2^{-2\kappa^*_{ij}(x)}. \tag{99}$$

Writing the dependence of γ^*_{ij} on x explicitly is useful in this appendix because the endogenous information choice is one of the two feedback channels that determine uniqueness.

A. Deriving the sectoral fixed point

Within sector j, relative firm demand is

$$q_{ij,t} \equiv \frac{y_{ij,t}}{Y_{j,t}} = F\left(\frac{P_{ij,t}}{\lambda_{j,t}P_{j,t}}\right) = F\left(\frac{P_{ij,t}}{X_{j,t}}\right). \tag{100}$$

Let η_{ij} denote the steady-state elasticity of this demand curve. A first-order approximation around the asymmetric steady state gives

$$\widehat{q}_{ij,t} = -\eta_{ij}\left(\widehat{P}_{ij,t} - \widehat{X}_{j,t}\right). \tag{101}$$

The Kimball aggregation constraint is

$$\int_0^1 G(q_{ij,t})\, di = 1.$$

Log-linearizing we have,

$$E_j^s[\widehat{q}_{ij,t}] = 0. \tag{102}$$

Substituting (101) into (102) gives the linearized Kimball strategic-index condition

$$\widehat{X}_{j,t} = \frac{E_j^s[\eta_{ij}\widehat{P}_{ij,t}]}{E_j^s[\eta_{ij}]}. \tag{103}$$

This is the aggregation equation that closes the sector. Notice that $X_{j,t}$ is not simply the sales-weighted average firm price. The Kimball constraint places an additional weight on each firm's local demand elasticity, η_{ij}.

Conditional on x, the firm's optimal price can be written as

$$\widehat{P}_{ij,t} = \gamma_{ij}^*(x)\phi_{ij}(x)\left(\widehat{W}_t + \varepsilon_{ij,t}\right), \tag{104}$$

where $\varepsilon_{ij,t}$ is the idiosyncratic signal noise. Assume that these errors are independent across firms and law of large numbers within each sector. Then

$$E_j^s\left[\eta_{ij}\gamma_{ij}^*(x)\phi_{ij}(x)\varepsilon_{ij,t}\right] = 0.$$

Substituting (104) into (103),

$$\widehat{X}_{j,t} = T_j(x)\widehat{W}_t, \tag{105}$$

where

$$T_j(x) \equiv \frac{E_j^s\left[\eta_{ij}\gamma_{ij}^*(x)\phi_{ij}(x)\right]}{E_j^s[\eta_{ij}]}. \tag{106}$$

Combining the conjectured response (97) with the response generated by firms in (105), equilibrium requires

$$x = T_j(x). \tag{107}$$

Equivalently,

$$x = \frac{E_j^s\left[\eta_{ij}\gamma_{ij}^*(x)\left\{\rho_{ij} + (1-\rho_{ij})x\right\}\right]}{E_j^s[\eta_{ij}]}. \tag{108}$$

Thus $T_j(x)$ is the sectoral response generated by firms when they conjecture that the sectoral strategic index responds by x. A fixed point is precisely a self-consistent conjecture.

B. Existence for $p > 1$

Proposition 4 (Existence of a sectoral fixed point)**.** *Suppose $p > 1$, $0 < \rho_{ij} \leq 1$, $\eta_{ij} > 0$, and $0 < E_j^s[\eta_{ij}] < \infty$. Then the mapping T_j defined in* (106) *has at least one fixed point in $[0,1]$. Under finite information costs, the fixed point is interior.*

Proof. We proceed in two steps.

Step 1: T_j maps $[0,1]$ into itself. For any $x \in [0,1]$,

$$\rho_{ij} \leq \phi_{ij}(x) = \rho_{ij} + (1-\rho_{ij})x \leq 1. \tag{109}$$

Moreover, a Kalman gain satisfies $0 \leq \gamma_{ij}^*(x) \leq 1$. Hence

$$0 \leq \eta_{ij}\gamma_{ij}^*(x)\phi_{ij}(x) \leq \eta_{ij}.$$

Taking the sales-weighted expectation and dividing by $E_j^s[\eta_{ij}] > 0$ gives

$$0 \leq T_j(x) \leq 1. \tag{110}$$

Thus $T_j : [0,1] \to [0,1]$.

Step 2: T_j is continuous. Conditional on x, firm ij solves

$$\min_{\kappa \geq 0}\left\{\frac{1}{2}V_{ij}(x)2^{-2\kappa} + \omega\kappa^p\right\}, \qquad V_{ij}(x) \equiv \Lambda_{ij}\sigma_W^2\phi_{ij}(x)^2. \tag{111}$$

For $p > 1$, the objective is strictly convex in κ. Since $\phi_{ij}(x) > 0$ on $[0,1]$, we have $V_{ij}(x) > 0$, and the derivative of the objective at $\kappa = 0$ is $-\ln(2)V_{ij}(x) < 0$. The optimum is therefore unique and interior. Because $V_{ij}(x)$ is continuous in x, the unique solution $\kappa_{ij}^*(x)$ is continuous in x. It follows from (99) that $\gamma_{ij}^*(x)$ is continuous, and therefore $T_j(x)$ is continuous.

Now define

$$H_j(x) \equiv T_j(x) - x.$$

From (110),

$$H_j(0) = T_j(0) \geq 0, \qquad H_j(1) = T_j(1) - 1 \leq 0.$$

Continuity and the intermediate value theorem imply that there exists at least one $x_j^* \in [0,1]$ such that $H_j(x_j^*) = 0$, or equivalently $x_j^* = T_j(x_j^*)$.

For the baseline $p > 1$ case, the fixed point is in fact interior. Because $V_{ij}(0) > 0$ and the marginal information cost at zero capacity is zero, every firm chooses strictly positive capacity, so $\gamma_{ij}^*(0) > 0$ and hence $T_j(0) > 0$. At $x = 1$, $\phi_{ij}(1) = 1$. With finite information costs, optimal capacity is finite, so $\gamma_{ij}^*(1) < 1$ and therefore $T_j(1) < 1$. Hence $H_j(0) > 0$ and $H_j(1) < 0$, which implies $x_j^* \in (0,1)$. □

C. A sufficient condition for uniqueness

Existence requires only continuity and the fact that T_j maps the unit interval into itself. With startegic complemenetraities and endogenous infromation feedback, uniqueness is not immediate. It requires that the strength of the feedback from a conjectured sectoral response x back into the response generated by firms to be 'controlled".

Differentiate (106). Since

$$\phi_{ij}'(x) = 1 - \rho_{ij},$$

we obtain

$$T_j'(x) = \frac{E_j^s\left[\eta_{ij}\left\{\gamma_{ij}^{*\prime}(x)\phi_{ij}(x) + \gamma_{ij}^*(x)(1-\rho_{ij})\right\}\right]}{E_j^s[\eta_{ij}]}. \tag{112}$$

Equation (112) makes the two feedback channels transparent.

(i) *Pricing complementarity.* Holding attention fixed, a larger conjectured x raises the firm's target exposure because

$$\phi_{ij}'(x) = 1 - \rho_{ij}.$$

The resulting contribution to $T_j'(x)$ is

$$\gamma_{ij}^*(x)(1-\rho_{ij}).$$

Firms with lower partial-equilibrium pass-through have a stronger direct response to the strategic price index and therefore generate a larger pricing feedback.

(ii) *Endogenous-attention complementarity.* A larger conjectured x also raises $\phi_{ij}(x)$, which raises the stakes of learning about the nominal shock. Firms therefore acquire more information, increasing $\gamma_{ij}^*(x)$. This additional feedback contributes

$$\gamma_{ij}^{*\prime}(x)\phi_{ij}(x)$$

to $T_j'(x)$.

The second channel is the new one created by endogenous information acquisition. The feedback can be summarized as

$$x \uparrow \implies \phi_{ij}(x) \uparrow \implies V_{ij}(x) \uparrow \implies \kappa_{ij}^*(x) \uparrow \implies \gamma_{ij}^*(x) \uparrow \implies T_j(x) \uparrow .$$

We next characterize this response explicitly.

The first-order condition associated with (111) is

$$\ln(2) V_{ij}(x) 2^{-2\kappa_{ij}^*(x)} = p\omega \left(\kappa_{ij}^*(x)\right)^{p-1}. \tag{113}$$

Log-differentiating (113) and using

$$\frac{d\gamma_{ij}^*}{d\kappa_{ij}^*} = 2\ln 2 \left(1 - \gamma_{ij}^*\right),$$

we obtain

$$\gamma_{ij}^{*\prime}(x) = \frac{4\ln 2 \left(1 - \gamma_{ij}^*(x)\right)(1 - \rho_{ij})}{\phi_{ij}(x)\left[2\ln 2 + (p-1)/\kappa_{ij}^*(x)\right]}. \tag{114}$$

In particular, $\gamma_{ij}^{*\prime}(x) \geq 0$: a stronger conjectured sectoral response weakly raises the firm's optimal attention to the nominal shock.

Substituting (114) into (112) also gives the useful representation

$$T_j'(x) = \frac{E_j^s\left[\eta_{ij}(1-\rho_{ij})\left\{\gamma_{ij}^*(x) + \frac{4\ln 2\left(1-\gamma_{ij}^*(x)\right)}{2\ln 2 + (p-1)/\kappa_{ij}^*(x)}\right\}\right]}{E_j^s[\eta_{ij}]}. \tag{115}$$

The first term in braces is the pricing-complementarity feedback, while the second is the additional feedback generated by endogenous information acquisition.

Proposition 5 (A sufficient condition for uniqueness)**.** *Under the assumptions of Proposition 4, the sectoral fixed point is unique if*

$$\sup_{x\in[0,1]} T_j'(x) < 1. \tag{116}$$

Equivalently, a sufficient condition is

$$\sup_{x\in[0,1]} \frac{E_j^s\left[\eta_{ij}\left\{\gamma_{ij}^*(x)(1-\rho_{ij}) + \gamma_{ij}^{*\prime}(x)\phi_{ij}(x)\right\}\right]}{E_j^s[\eta_{ij}]} < 1. \tag{117}$$

Proof. Again define $H_j(x) = T_j(x) - x$. Under (116),

$$H_j'(x) = T_j'(x) - 1 < 0$$

for every $x \in [0,1]$. Hence H_j is strictly decreasing. Proposition 4 establishes that H_j has at least one zero in $[0,1]$. A strictly decreasing function can cross zero at most once, so the fixed point is unique. □

The condition in Proposition 5 is sufficient rather than necessary. Economically, it requires the total feedback from a conjectured increase in the sectoral response—through both firms' desired prices and their endogenous information choices—to remain weaker than the initial increase in the conjecture itself. If this feedback were sufficiently strong, a high-response conjecture could induce enough additional information acquisition and price adjustment to become self-confirming, opening the door to multiple fixed points.

Remark 1 (Exogenous attention)**.** *The role of endogenous information acquisition is especially clear if the Kalman gains are instead held fixed. Then $\gamma_{ij}^{*\prime}(x) = 0$, and*

$$T_j'(x) = \frac{E_j^s[\eta_{ij}\gamma_{ij}(1-\rho_{ij})]}{E_j^s[\eta_{ij}]}.$$

Since $0 \le \gamma_{ij} \le 1$ and $0 < \rho_{ij} \le 1$, this derivative is strictly below one whenever there is positive sales weight on firms with $\rho_{ij} > 0$. Hence the fixed point is automatically unique with exogenous attention. The nontrivial part of the uniqueness condition in the baseline model comes from the additional term $\gamma_{ij}^{*\prime}(x)\phi_{ij}(x)$ generated by endogenous information acquisition.

D. Linear information costs: $p = 1$

The existence argument continues to go through when information costs are linear, but the information choice can now be at a corner. For $p = 1$, firm ij solves

$$\min_{\kappa \geq 0} \left\{ \frac{1}{2} V_{ij}(x) 2^{-2\kappa} + \omega\kappa \right\}.$$

The unique optimum is

$$\kappa_{ij}^*(x) = \begin{cases} 0, & V_{ij}(x) \ln 2 \leq \omega, \\ \frac{1}{2} \log_2 \left(\frac{V_{ij}(x) \ln 2}{\omega} \right), & V_{ij}(x) \ln 2 > \omega. \end{cases} \tag{118}$$

The corresponding Kalman gain is

$$\gamma_{ij}^*(x) = \begin{cases} 0, & V_{ij}(x) \ln 2 \leq \omega, \\ 1 - \dfrac{\omega}{V_{ij}(x) \ln 2}, & V_{ij}(x) \ln 2 > \omega. \end{cases} \tag{119}$$

Although $\gamma_{ij}^*(x)$ is kinked when a firm switches from zero to positive information acquisition, it is continuous at the threshold. Hence $T_j(x)$ remains continuous and still maps $[0, 1]$ into itself, so a sectoral fixed point exists.

The qualitative difference is that, with linear information costs, some firms may optimally acquire no information. In particular, $x = 0$ can itself be a fixed point if all firms have sufficiently low information stakes at $x = 0$. Thus the extensive margin of information acquisition creates an additional source of equilibrium feedback.

Away from an information-acquisition threshold, an attentive firm satisfies

$$\gamma_{ij}^*(x) = 1 - \frac{\omega}{\Lambda_{ij} \sigma_W^2 \ln 2 \, \phi_{ij}(x)^2}.$$

It follows that

$$\frac{d}{dx} \left[\gamma_{ij}^*(x) \phi_{ij}(x) \right] = (1 - \rho_{ij}) \left[2 - \gamma_{ij}^*(x) \right]. \tag{120}$$

Thus the slope of T_j is piecewise given by

$$T_j'(x) = \frac{E_j^s\left[\eta_{ij}\mathbf{1}\{V_{ij}(x)\ln 2 > \omega\}(1-\rho_{ij})\left(2-\gamma_{ij}^*(x)\right)\right]}{E_j^s[\eta_{ij}]}, \tag{121}$$

away from threshold points. A convenient stronger sufficient condition for uniqueness is therefore

$$2\,\frac{E_j^s[\eta_{ij}(1-\rho_{ij})]}{E_j^s[\eta_{ij}]} < 1. \tag{122}$$

This condition is stronger than necessary, but it guarantees that all one-sided slopes of the fixed-point mapping are less than 1. Relative to the $p > 1$ baseline, the $p = 1$ case therefore introduces an extensive margin of attention and can generate stronger self-reinforcing feedback between aggregate price adjustment and firms' incentives to become informed.

C.6 Proof of Proposition 3

This section derives the linearized sectoral price indices, the sectoral fixed point, and the sectoral and aggregate price responses.

Linearized sectoral price indices. Recall the two exact Kimball aggregation conditions:

$$P_{j,t} = \int_0^1 P_{ij,t} F\left(\frac{P_{ij,t}}{\lambda_{j,t}P_{j,t}}\right) di, \tag{123}$$

$$1 = \int_0^1 G\left[F\left(\frac{P_{ij,t}}{\lambda_{j,t}P_{j,t}}\right)\right] di. \tag{124}$$

Log-linearizing (123) gives

$$\widehat{P}_{j,t} = E_j^s[\widehat{P}_{ij,t}] - E_j^s\left[\eta_{ij}\left(\widehat{P}_{ij,t} - \widehat{\lambda}_{j,t} - \widehat{P}_{j,t}\right)\right]. \tag{125}$$

The second Kimball condition provides the remaining restriction.
Since $F = (G')^{-1}$, we have $G'\left(F\left(\frac{\widehat{P}_{ij,t}}{\widehat{P}_{jt}\widehat{\lambda}_{jt}}\right)\right) = \left(\frac{\widehat{P}_{ij,t}}{\widehat{P}_{jt}\widehat{\lambda}_{jt}}\right)$.
Together with $\left(\frac{\widehat{P}_{ij,t}}{\widehat{P}_{jt}\widehat{\lambda}_{jt}}\right) F'\left(\frac{\widehat{P}_{ij,t}}{\widehat{P}_{jt}\widehat{\lambda}_{jt}}\right) = -\eta_{ij}F\left(\frac{\widehat{P}_{ij,t}}{\widehat{P}_{jt}\widehat{\lambda}_{jt}}\right)$, log-linearizing (124) yields

$$E_j^s\left[\eta_{ij}\left(\widehat{P}_{ij,t} - \widehat{\lambda}_{j,t} - \widehat{P}_{j,t}\right)\right] = 0. \tag{126}$$

Substituting (126) into (125) gives

$$\widehat{P}_{j,t} = E_j^s[\widehat{P}_{ij,t}]. \tag{127}$$

Equation (126) also implies

$$\widehat{P}_{j,t} + \widehat{\lambda}_{j,t} = \frac{E_j^s[\eta_{ij}\widehat{P}_{ij,t}]}{E_j^s[\eta_{ij}]} \tag{128}$$

Equivalently, using (127),

$$\widehat{\lambda}_{j,t} = \frac{\mathrm{Cov}_j^s(\eta_{ij}, \widehat{P}_{ij,t})}{E_j^s[\eta_{ij}]} \tag{129}$$

Sectoral fixed point: Let

$$x_j \equiv \alpha_j + \ell_j.$$

Under firms' optimal information and pricing decisions,

$$\widehat{P}_{ij,t} = \gamma_{ij}^*(x_j)\left[\rho_{ij} + (1-\rho_{ij})x_j\right]\widehat{W}_t. \tag{130}$$

Substituting (130) into (128), and imposing $\widehat{P}_{j,t} + \widehat{\lambda}_{j,t} = x_j\widehat{W}_t$, gives

$$x_j = \frac{E_j^s\left[\eta_{ij}\gamma_{ij}^*(x_j)\left\{\rho_{ij} + (1-\rho_{ij})x_j\right\}\right]}{E_j^s[\eta_{ij}]}. \tag{131}$$

Let x_j^* denote an equilibrium solution and, with a slight abuse of notation, write $\gamma_{ij}^* \equiv \gamma_{ij}^*(x_j^*)$. Conditional on the equilibrium information choices, (131) can be rearranged as

$$x_j^* = \frac{E_j^s[\eta_{ij}\gamma_{ij}^*\rho_{ij}]}{E_j^s[\eta_{ij}] - E_j^s[\eta_{ij}\gamma_{ij}^*(1-\rho_{ij})]}. \tag{132}$$

Sectoral price pass-through: Using (127) and firms' optimal price responses,

$$\begin{aligned}\widehat{P}_{j,t} &= E_j^s[\widehat{P}_{ij,t}] \\ &= \left\{E_j^s[\gamma_{ij}^*\rho_{ij}] + x_j^* E_j^s[\gamma_{ij}^*(1-\rho_{ij})]\right\}\widehat{W}_t.\end{aligned} \tag{133}$$

Hence

$$\alpha_j = E_j^s[\gamma_{ij}^*\rho_{ij}] + x_j^* E_j^s[\gamma_{ij}^*(1-\rho_{ij})]. \tag{134}$$

Since

$$E_j^s[\gamma_{ij}^*\rho_{ij}] = E_j^s[\gamma_{ij}^*] - E_j^s[\gamma_{ij}^*(1-\rho_{ij})],$$

we obtain

$$\alpha_j = E_j^s[\gamma_{ij}^*] - (1-x_j^*)E_j^s[\gamma_{ij}^*(1-\rho_{ij})]. \tag{135}$$

From (132),

$$\begin{aligned} 1 - x_j^* &= \frac{E_j^s[\eta_{ij}] - E_j^s[\eta_{ij}\gamma_{ij}^*(1-\rho_{ij})] - E_j^s[\eta_{ij}\gamma_{ij}^*\rho_{ij}]}{E_j^s[\eta_{ij}] - E_j^s[\eta_{ij}\gamma_{ij}^*(1-\rho_{ij})]} \\ &= \frac{E_j^s[\eta_{ij}] - E_j^s[\eta_{ij}\gamma_{ij}^*]}{E_j^s[\eta_{ij}] - E_j^s[\eta_{ij}\gamma_{ij}^*(1-\rho_{ij})]} \equiv K_j^+. \end{aligned} \tag{136}$$

Substituting into (135) yields

$$\alpha_j = E_j^s[\gamma_{ij}^*] - K_j^+ E_j^s[\gamma_{ij}^*(1-\rho_{ij})] \tag{137}$$

Aggregate price pass-through. The CES price index across sectors is

$$P_t = \left[\int_0^1 \vartheta_j P_{j,t}^{1-\sigma}\, dj\right]^{\frac{1}{1-\sigma}}. \tag{138}$$

Its log-linearization around the asymmetric steady state is

$$\widehat{P}_t = \int_0^1 S_j^* \widehat{P}_{j,t}\, dj, \tag{139}$$

where

$$S_j^* \equiv \vartheta_j \left(\frac{P_j^*}{P^*}\right)^{1-\sigma} = \frac{P_j^* Y_j^*}{P^* Y^*} \tag{140}$$

is sector j's steady-state expenditure share and $\int_0^1 S_j^*\, dj = 1$. Substituting $\widehat{P}_{j,t} = \alpha_j \widehat{W}_t$ into (139) gives

$$\widehat{P}_t = \left(\int_0^1 S_j^* \alpha_j\, dj\right) \widehat{W}_t.$$

Therefore,

$$\alpha = \int_0^1 S_j^* \alpha_j\, dj \tag{141}$$

This establishes the sectoral and aggregate price pass-through results. □

D Data and Decomposition

D.1 Textual Attention to Kalman Gains

This section estimates firm-level attention to inflation and the implied Kalman gain with which firms update inflation beliefs. The exercise has two connected goals. First, it uses firms' text-based discussion of macroeconomic conditions to recover a latent state of attention. Second, it maps this latent attention state into the precision of firms' inflation signals and, therefore, into the Kalman gain that governs belief updating.

The key empirical difficulty is that attention and signal precision are not separately observed. Text data reveal a noisy measure of attention, while survey data from SoFIE give aggregate properties of expectations and forecast errors. The estimator therefore combines these two sources of information. The firm text panel disciplines the stochastic process for latent attention and the measurement error in text. The aggregate survey moments discipline how attention maps into signal precision and hence Kalman gains.

Model

Inflation follows a mean-reverting process,

$$\pi_t - \mu_\pi = \rho_\pi(\pi_{t-1} - \mu_\pi) + u_t, \qquad u_t \sim N(0, \sigma_\pi^2). \tag{142}$$

The parameters $(\mu_\pi, \rho_\pi, \sigma_\pi)$ summarize the unconditional level, persistence, and innovation volatility of inflation. Firms are assumed to care about inflation as the relevant macroeconomic state.

Firm i has latent attention a_{it}, which combines a permanent firm component and a persistent transitory component:

$$a_{it} = \eta_i + x_{it}, \qquad \eta_i \sim N(0, \sigma_\eta^2), \tag{143}$$

$$x_{it} = \rho_a x_{i,t-1} + \omega_{it}, \qquad \omega_{it} \sim N(0, \sigma_a^2). \tag{144}$$

The permanent component η_i captures stable cross-firm differences in attention. The persistence parameter ρ_a governs how slowly the transitory component x_{it} moves over time, and σ_a governs the size of new within-firm movements in attention. This standardized text outcome is a noisy

signal of total latent attention:

$$\widetilde{T}_{it} = \alpha + a_{it} + \epsilon_{it}, \qquad \epsilon_{it} \sim N(0, \sigma_\epsilon^2). \tag{145}$$

The measurement equation contains no macro controls: the identifying assumption is that changes in the adjusted text indicator reflect changes in attention plus measurement noise.

Firms observe noisy signals about inflation,

$$s_{it} = \pi_t + \zeta_{it}, \qquad \zeta_{it} \sim N(0, R_{it}), \tag{146}$$

where R_{it} is signal-noise variance. To make attention comparable across parameter values, define the stationary standard deviation of latent attention,

$$\tilde{\sigma}_a = \sqrt{\sigma_\eta^2 + \frac{\sigma_a^2}{1-\rho_a^2}}, \tag{147}$$

and standardized attention

$$\tilde{a}_{it} = \frac{a_{it}}{\tilde{\sigma}_a}. \tag{148}$$

Signal precision is then specified as

$$\log R_{it}^{-1} = \operatorname{sgn}(\theta_0 + \theta_1 \tilde{a}_{it}) \left|\theta_0 + \theta_1 \tilde{a}_{it}\right|^{\theta_2}, \qquad \theta_1 > 0, \quad \theta_2 > 0. \tag{149}$$

The parameter θ_0 determines the baseline precision index, θ_1 determines how strongly attention raises precision, and θ_2 determines curvature. When θ_2 is close to one, the mapping is approximately linear in standardized attention. The signed-power specification keeps the mapping real-valued without requiring the index $\theta_0 + \theta_1 \tilde{a}_{it}$ to be positive for all firms and periods.

Kalman Updating

Let $\widehat{\pi}_{it|t-1}$ denote firm i's prior belief about inflation and $P_{it|t-1}$ its prior variance. The Kalman gain is

$$K_{it} = \frac{P_{it|t-1}}{P_{it|t-1} + R_{it}}. \tag{150}$$

Beliefs update according to

$$\widehat{\pi}_{it|t} = \widehat{\pi}_{it|t-1} + K_{it}\left(s_{it} - \widehat{\pi}_{it|t-1}\right). \tag{151}$$

Posterior and next-period prior variances evolve as

$$P_{it|t} = (1 - K_{it})P_{it|t-1}, \tag{152}$$

$$P_{i,t+1|t} = \rho_\pi^2 P_{it|t} + \sigma_\pi^2. \tag{153}$$

Thus, attention affects belief updating only through signal precision. This is the mechanism that links text-based attention to estimated Kalman gains.

Estimation

The estimation uses three empirical objects. The first is an annual firm text panel containing the text measure, transformed into $\widetilde{T}_{it}$ as described above. The baseline sample keeps firms observed for at least five years, which leaves 6,069 firms and 78,458 nonmissing firm-year gain estimates. The second object is realized inflation. The second object is annual inflation data used to estimate the AR(1) process, and the third object is aggregate survey information: average expected inflation, dispersion in expected inflation across firms.

For a candidate parameter vector, the model simulates annual latent text observations on the observed firm-year mask and annual survey-style belief moments on the annual inflation path. The SMM objective is

$$\widehat{\Theta} = \arg\min_{\Theta} \left[m^{data} - \bar{m}^{model}(\Theta) \right]' W \left[m^{data} - \bar{m}^{model}(\Theta) \right]. \tag{154}$$

Parameter Estimates

Table C1 reports the active standardized-text estimates.

Table C1: Estimated Parameters: Active Standardized Text Specification

Parameter	Estimate
Panel A: Annual Inflation Process	
μ_π	2.5785
ρ_π	0.4178
σ_π	1.4639
Panel B: Attention and Precision Mapping	
ρ_a	0.6834
σ_a	0.0412
σ_η	0.9078
α	0.0114
σ_ϵ	0.0216
θ_0	-2.3092
θ_1	3.8582
θ_2	1.2105
SMM objective	0.1228
SMM converged	Yes

Model Fit

Table C2 compares targeted data moments with model moments. The estimator matches the external forecast-error variance target closely and fits the text moments reasonably well. The largest discrepancy is average expectation dispersion, which is lower in the model than in the survey data.

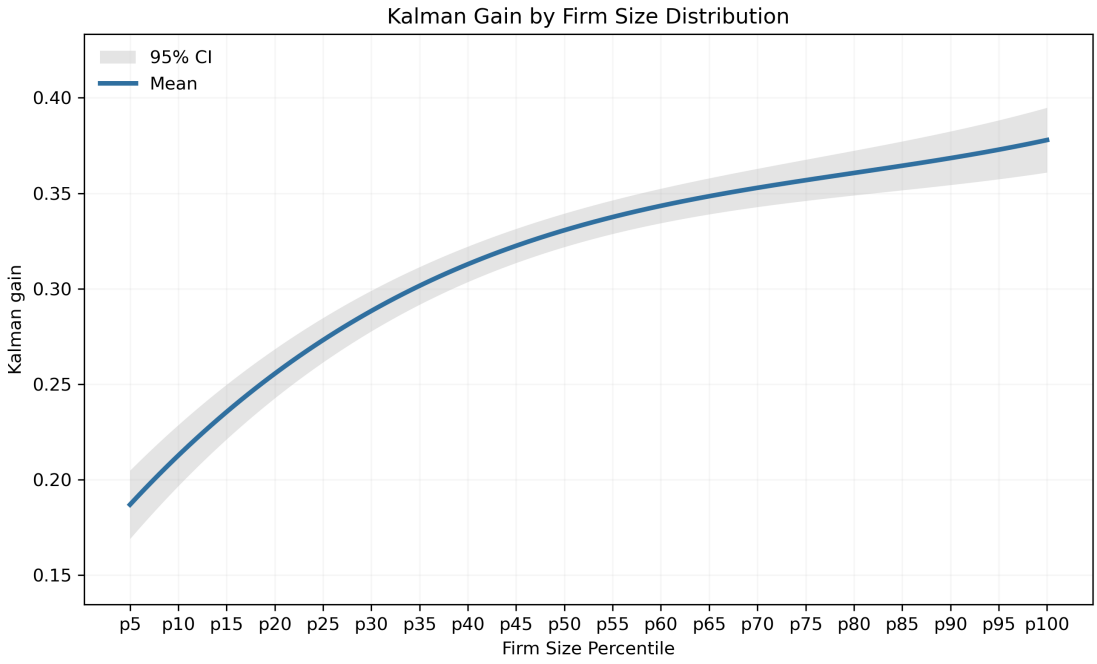
Table C2: Targeted Moment Fit: Active Standardized Text Specification

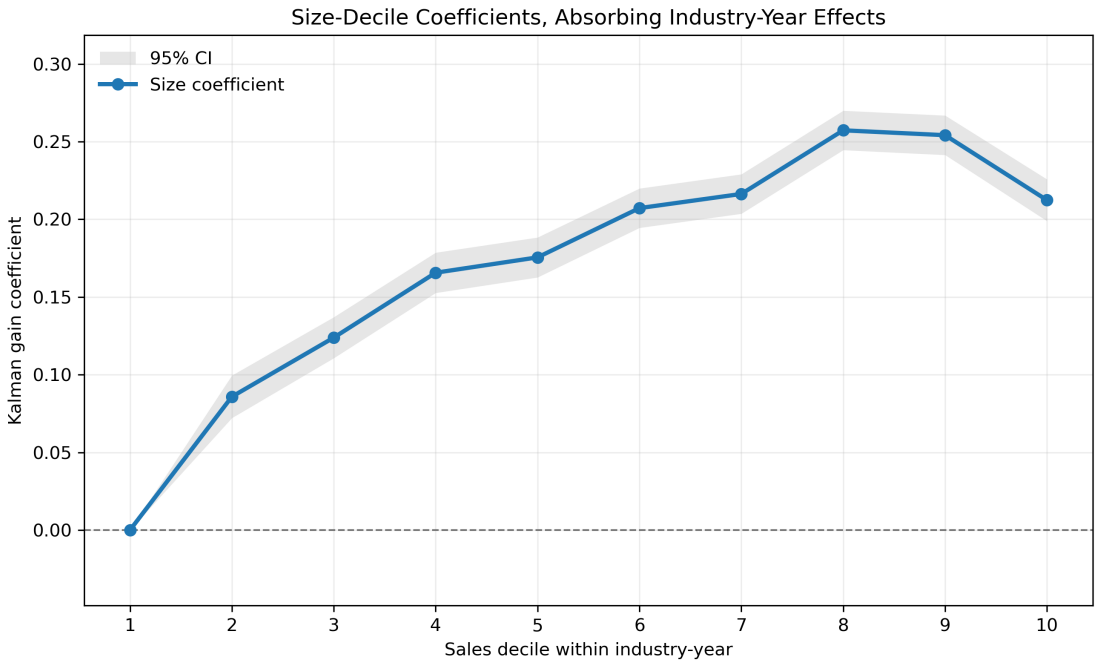
Source	Moment	Data	Model	Data–model
Text	Mean	0.0197	0.0134	0.0063
Text	Std. dev.	0.9757	0.9068	0.0689
Text	25th percentile	-0.7132	-0.6022	-0.1110
Text	75th percentile	0.5173	0.6278	-0.1106
Text	Firm autocorr.	0.2721	0.2725	-0.0003
Text	Year-mean autocorr.	0.8402	0.7401	0.1001
Text	Cross-firm dispersion	0.9610	0.9066	0.0543
Survey	Forecast-error variance	1.4000	1.3922	0.0078
Survey	Mean dispersion	1.4036	1.1871	0.2165
Survey	Dispersion variance	0.3490	0.3337	0.0153
Survey	Dispersion-absolute innovation covariance	0.4163	0.4278	-0.0115

Table C3: Goodness of Fit

Sample	Source	Moments	RMSE	Relative RMSE	Fit R^2
Targeted	All	11	0.0903	0.1057	0.9769
Targeted	Text	7	0.0777	0.1173	0.9812
Targeted	Survey	4	0.1087	0.0814	0.9546
Nontargeted	All	15	0.7582	0.8651	0.6691
Nontargeted	Text	8	0.6994	0.7144	0.7972
Nontargeted	Survey	7	0.8202	1.0101	0.2960

Estimated Kalman Gains





D.2 Local Steady-State Approximation for Sales Shares

The empirical exercise requires the steady-state within-industry sales share $s^*_{ij|j}$, while Compustat provides annual realized shares. This appendix gives the local approximation that motivates using a short time average of log sales shares around a reference year.

Model implication Let

$$s^{obs}_{ij|j,t} \equiv \frac{P_{ij,t} y_{ij,t}}{P_{j,t} Y_{j,t}}$$

denote firm ij's realized revenue share within sector j, and let hats denote log deviations from the asymmetric steady state. Kimball demand implies, to a first order,

$$\widehat{s}_{ij|j,t} = (1 - \eta_{ij})\left(\widehat{P}_{ij,t} - \widehat{P}_{j,t}\right) + \eta_{ij}\widehat{\lambda}_{j,t}. \tag{155}$$

Recall firm's pricing rule

$$\widehat{P}_{ij,t} = \phi_{ij}\gamma_{ij}(\widehat{W}_t + \varepsilon_{ij,t}), \qquad \phi_{ij} \equiv \rho_{ij} + (1 - \rho_{ij})(\alpha_j + \ell_j).$$

Combining these expressions gives

$$\widehat{s}_{ij|j,t} = \underbrace{\Big[(1-\eta_{ij})(\phi_{ij}\gamma_{ij} - \alpha_j) + \eta_{ij}\ell_j\Big]}_{\equiv B_{ij}} \widehat{W}_t + (1-\eta_{ij})\phi_{ij}\gamma_{ij}\varepsilon_{ij,t}. \tag{156}$$

Both $\widehat{W}_t$ and the signal noise $\varepsilon_{ij,t}$ are mean zero. Hence, conditional on the firm type,

$$E\Big[\widehat{s}_{ij|j,t} \mid i, j\Big] = 0. \tag{157}$$

Since

$$\log s^{obs}_{ij|j,t} = \log s^{*}_{ij|j} + \widehat{s}_{ij|j,t} + O(2),$$

it follows that

$$E\Big[\log s^{obs}_{ij|j,t} \mid i, j\Big] = \log s^{*}_{ij|j} + O(2). \tag{158}$$

Thus, to a first-order approximation, time-averaging log realized shares recovers the local steady-state log share.

Local time averaging For a reference year τ, let $\mathcal{W}_{ij,\tau}$ denote the years in a short window around τ in which firm ij is observed and let $T_{ij,\tau} = |\mathcal{W}_{ij,\tau}|$. The local-constant estimator is

$$\widehat{x}^{*}_{ij|j,\tau} = \frac{1}{T_{ij,\tau}} \sum_{t \in \mathcal{W}_{ij,\tau}} \log s^{obs}_{ij|j,t} \tag{159}$$

The level shares are obtained by exponentiating and renormalizing within the reference-year cross-section,

$$\widehat{s}^{*}_{ij|j,\tau} = \frac{\exp(\widehat{x}^{*}_{ij|j,\tau})}{\sum_{k \in j} \exp(\widehat{x}^{*}_{kj|j,\tau})}. \tag{160}$$

The argument requires only that, over this local window, the underlying firm type is approximately fixed, and the deviations around it are transitory and mean zero. [13]

Why use a short window? A longer window reduces transitory noise but does not necessarily improve the steady-state approximation. Over longer horizons, persistent firm growth, changes in product demand or quality, mergers and acquisitions, entry and exit, and changes in industry

[13] The argument does not require shocks to be exactly i.i.d. If the disturbances driving realized sales shares are mildly persistent with mean-zero disturbances, then realized log shares still fluctuate around the same local steady-state component.

composition can move the underlying steady state itself. Averaging over such episodes would mix observations generated by different local steady states.

For the baseline exercise, I therefore take 2015 as the reference year and use the 2012–2018 window. The cross-section is defined by firms active in 2015; a firm need not be observed in every year of the window.

E Calibration and Counterfactual

E.1 A Parametric Demand

Empirical evidence suggests substantial heterogeneity in marginal-cost pass-through across firms. Amiti, Itskhoki and Konings (2019) document nearly complete pass-through among small firms but considerably lower pass-through among large firms. Similarly, Burya and Mishra (2022) find that pass-through declines sharply with firm market share at the product-level using Neilsen Retail Scanner data. Accommodating these patterns while maintaining empirically plausible markups presents a challenge for commonly used non-CES demand specifications.

To illustrate this challenge, recall that the pass-through of an idiosyncratic marginal-cost shock to the firm's optimal price is given by

$$\rho_i = \left(1 + \frac{\xi_i}{\eta_i - 1}\right)^{-1}.$$

Klenow-Willis Demand Specification The Klenow and Willis (2016) specification imposes a particular relationship between demand curvature and markups,

$$\frac{\xi_i}{\eta_i - 1} = \widetilde{b}\mu_i,$$

where $\widetilde{b}$ is a common parameter. Consequently, firm-level pass-through satisfies

$$\rho_i^{KW} = \frac{1}{1 + \widetilde{b}\mu_i}.$$

This restriction makes it difficult to jointly rationalize near-complete pass-through among small firms and very low pass-through among large firms without generating implausibly high markups. For example, suppose small firms have pass-through of 0.95 and gross markups of 1.2. Matching these moments requires a very small calibrated $\widetilde{b} \simeq 0.044$. Given this curvature, generat-

ing pass-through of say 0.2 among large firms would require a gross markup of approximately 91. Thus, calibrating the specification to match the high pass-through of small firms makes it impossible to generate lower pass-through among large firms without imposing an implausible elasticity schedule. And given that the aggregate price responsiveness α_j depends on the joint distribution of sales hares, pass-through, elasticity and Kalman gains, getting the elasticity schedule to stay in the plausible range is important. This difficulty has also been explicitly noted in Baqaee, Farhi and Sangani (2024).

The Dotsey and King Specification The Dotsey and King (2005) specification imposes a different restriction. Let x_i denote the firm's relative output. The demand elasticity and superelasticity are given by

$$\eta_i(x_i) = \frac{\omega}{\omega - 1}\left(1 + \psi - \frac{\psi}{x_i}\right), \qquad \xi_i(x_i) = -\frac{\omega\psi}{(\omega - 1)x_i}$$

and $\omega > 1$. The pass-through is therefore

$$\rho_i^{DK}(x_i) = \frac{(1 + \omega\psi)x_i - \omega\psi}{(1 + \omega\psi)x_i - 2\omega\psi}.$$

where ω, ψ are demand parameters to be calibrated. At $\psi = 0$, this collapses to a CES demand with elasticity ω. To be consistent with the empirical evidence of declining pass-through with size, $\psi < 0$ i.e., positive super-elasticity is imposed.

Under the restrictions, pass-through decreases with relative firm output. However, the specification imposes an upper bound on pass-through. As relative output approaches zero, pass-through is bounded above by 0.5. Therefore, although the Dotsey and King (2005) specification can accommodate very low pass-through among large firms, it effectively caps pass-through among the smallest firms at 0.5, preventing it from matching the very high pass-throughs documented empirically.

Taken together, these limitations motivate a more flexible demand specification capable of jointly accommodating high pass-through among small firms, low pass-through among large firms, and economically plausible markups.

A New Parametric Kimball Demand Specification To overcome the restrictions noted above, we introduce a more flexible demand specification. The specification combines the Klenow and

Willis (2016) and Dotsey and King (2005) demand aggregators through a convex combination:

$$G_J(x) = \alpha G_{KW}(x) + (1-\alpha)G_{DK}(x), \qquad \alpha \in [0,1],$$

where $G_{KW}(x)$ and $G_{DK}(x)$ denote the Klenow–Willis and Dotsey–King aggregators, respectively, x denotes relative firm output, and α governs their relative weights. Aggregate output Y is implicitly defined by

$$\int_0^1 G_J\left(\frac{y_i}{Y}\right) di = 1.$$

The proposed specification constitutes a valid Kimball aggregator. The aggregator functions are strictly increasing and strictly concave for positive relative output wherever both constituent functions are well-defined. Consequently,

$$G_J'(x) = \alpha G_{KW}'(x) + (1-\alpha)G_{DK}'(x) > 0,$$

$$G_J''(x) = \alpha G_{KW}''(x) + (1-\alpha)G_{DK}''(x) < 0.$$

The additional restriction on the feasible range of relative output arises from the Dotsey and King (2005) component, which requires $(1+\psi)x - \psi > 0$ to be real-valued. For $\psi < -1$, it requires $0 < x < \psi/(1+\psi)$. Moreover, we have $G_J(1) = 1$.

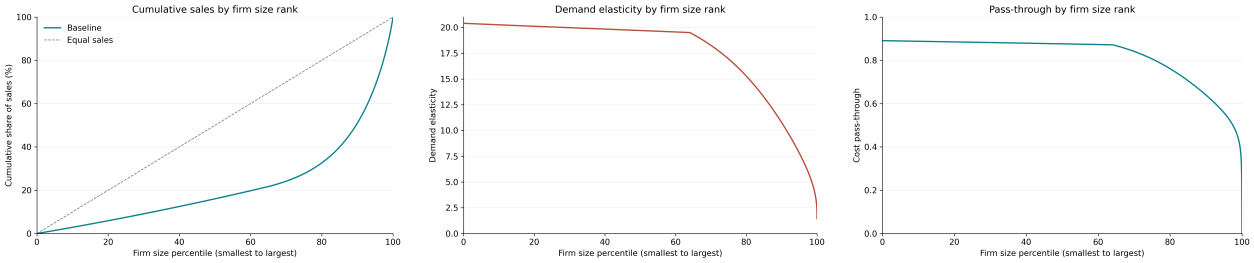
This specification, by construction, nests both Klenow and Willis (2016) and Dotsey and King (2005) demand. Setting $\alpha = 1$ recovers the Klenow–Willis aggregator, while setting $\alpha = 0$ recovers the Dotsey–King aggregator. For $\alpha \in (0,1)$, the specification combines the curvature properties of both demand systems. This is particularly useful when trying to match the joint distribution of firm size, markups, and marginal-cost pass-through observed in the data. Intuitively, for a given mixing weight α, the larger the relative output, the greater is the contribution of the power component to the local demand curvature (the Dotsey-King component).

I next examine the quantitative implications of the new specification under our baseline calibration. Figure D1 shows that the specification simultaneously captures key features of firm heterogeneity, including sales concentration, declining demand elasticities with firm size while keeping them in a reasonable range, and substantial variation in marginal-cost pass-through. In particular, the model generates a very high pass-through among small firms (almost 0.9) and substantially lower pass-through among large firms while maintaining economically plausible

markups.

While the model generates substantial sales concentration, it understates the degree observed in the data. Using a common rank convention, the top 10% of firms account for 49% of aggregate sales, compared with about 60% empirically, while the bottom 50% account for 16%, compared with 8% in the data.

Taken together, these results show that the new parametric demand specification can accommodate salient patterns of firm heterogeneity in sales, demand elasticities, and cost pass-through within a single parametric demand framework. This stands as another contribution of this paper.



(a) Cumulative sales **(b)** Demand elasticity **(c)** Cost pass-through

Figure D1: Firm Heterogeneity under the new parametric demand specification. The figure presents the implications of the baseline calibration across the firm-size distribution. Panel (a) plots the cumulative share of aggregate sales. Panels (b) and (c) show the demand elasticity and marginal-cost pass-through, respectively. The specification generates substantial sales concentration, declining demand elasticities, and pronounced heterogeneity in cost pass-through, with near-complete pass-through among small firms and substantially lower pass-through among large firms.