

Inflation Drivers in Firms' Words: Bridging Micro Narratives and Macro Dynamics*

Chenyu (Sev) Hou[†]
Simon Fraser University

Jiannan (Jay) Jiang[‡]
UT Austin

Tao Wang[§]
Bank of Canada

June 18, 2026
(The latest draft)

Abstract

This paper develops a framework that links microeconomic narratives to structural macro models of firm pricing. Using the Federal Reserve Bank's Beige Book data, we employ an LLM to extract qualitative measures of firms' price adjustments and factor attributions at the micro level. These measures are then used to estimate a state-space model of aggregate inflation dynamics implied by a menu-cost model featuring lumpy price adjustments and reporting frictions. We demonstrate how micro price narratives provide useful information for correctly decomposing the macro drivers of inflation.

Keywords: Narratives, LLM, Firm Pricing, Inflation, Menu cost

JEL Codes: D84, E31, E37, E71

*Sev Hou thanks the Canadian Social Sciences and Humanities Research Council for the financial support for this project. We would like to thank Daniel Villar for sharing their data, and Olivier Coibion, Martin Eichenbaum, Michael Ehrmann, and Lu Han for useful comments. The views of this paper are those of the authors and not of their respective institutions.

[†]Chenyu Hou: Department of Economics, Simon Fraser University, chenyu_hou@sfu.ca.

[‡]Jiannan (Jay) Jiang: Department of Economics, University of Texas at Austin, jiannanjiang@utexas.edu.

[§]Tao Wang: Bank of Canada, TaoWang@bank-banque-canada.ca.

1 Introduction

Economic measurement goes beyond numbers. A growing literature ([Gentzkow et al., 2019](#)) emphasizes that qualitative information and textual data can provide insights that are difficult to extract from standard quantitative indicators alone. Among such soft data, economic narratives have received increasing attention. Narratives capture how economic agents describe their experiences and decisions in real time. The central premise of this paper is that narratives are informative not only about how economic outcomes evolve, but also about why they do so.

This paper uses *micro* narratives, the qualitative descriptions provided by firms about their pricing decisions and operating conditions, to understand both firm-level mechanisms and the contributors to aggregate outcomes such as inflation. Macroeconomists conventionally rely on aggregate time series data and structural restrictions to infer the drivers of inflation. We instead take a bottom-up approach by directly “listening” to firms. Using a large language model, we extract from survey-based narratives of firms’ own explanations for why prices changed or remained unchanged. Although doing so does not directly identify the underlying structural shocks, it reveals the channels relevant to firms via which such shocks propagate. We contribute to a long tradition of survey-based analysis of firm behavior on both the measurement and application fronts. On the former, we systematically structure information from free-text narratives. On the latter, we augment the aggregate time-series-based analysis with factor attributions derived from micro-narratives.

One key feature of narrative data is its discreteness. Firms describe price and condition changes using categorical language, such as “increased substantially” or “stayed about the same”, rather than continuous numeric values. Rather than viewing this as a limitation, we argue that it is an advantage. Price adjustment is inherently lumpy due to menu costs and adjustment frictions. Narrative data, therefore, aligns naturally with the extensive margin of price adjustment, making it particularly informative about the forces that trigger price changes. We validate that our narrative statistics are strongly correlated with the frequency of price adjustments calculated with microdata, produced by [Nakamura et al. \(2018\)](#) and extended by [Montag and Villar Vallenas \(2025\)](#).

Our approach uses textual narratives to directly recover how firms explain their own pricing decisions. Going beyond the use of textual data for measuring a latent variable, we focus on extracting attribution statements that link observed price outcomes to underlying economic factors. Leveraging the reading and reasoning capabilities of large language models, we identify explicit and implicit explanations embedded in firms’ narratives. For example, we identify statements indicating that prices rose because of higher input costs, or that prices remained stable despite strong demand. These explanations provide structured information on the per-

ceived drivers of price adjustment, allowing us to study not only which economic conditions are present, but also how firms connect those conditions to their pricing behavior. This focus on attribution transforms narrative text into a source of relational economic information that can be directly incorporated into quantitative analysis.

Correctly extracting such information requires context-sensitive interpretation. Identical phrases can have different economic meanings depending on the discussion. For example, what constitutes a final price in one sector may represent an input cost from another sector’s perspective. This distinction is critical for measuring things like demand, supply constraints, labor shortages, or excess demand. Large language models are well-suited for this task because they can flexibly interpret language in context. Expressions such as “hiring was difficult”, “difficulty finding qualified candidates,” or “struggled to find quality workers” all convey labor scarcity despite substantial variation in wording that simpler dictionary-based methods are likely to miss.

We validate our narrative measures in several ways. First, we assess labeling accuracy by comparing model-based extraction to human-coded benchmarks, evaluating both factor identification and attribution logic. Second, we benchmark our narrative price measures against real-time macroeconomic series. Despite their categorical nature, these measures closely track realized inflation. Third, we show that the narrative measures align with regional economic conditions, providing additional external validation.

Interpreting micro-level narratives as evidence for macroeconomic inflation drivers requires a model-based framework. Micro responses do not aggregate mechanically due to heterogeneity, idiosyncratic shocks, and general equilibrium feedback. Moreover, firms do not report all relevant factors, and narrative observations are discrete while underlying fundamentals are continuous. We address these challenges with a sticky-price model featuring menu costs, aggregate shocks, and idiosyncratic shocks. Firms adjust prices only when deviations from desired prices exceed factor-specific thresholds, generating discrete price changes that map naturally into narrative categories. We also allow for reporting frictions, whereby firms do not always report all experienced changes driven by each factor.

The inputs to our estimation consist of frequencies of unconditional mentions of price and factor movements, along with conditional frequencies of attribution statements linking price changes to specific factors. Intuitively, if many firms cite a particular factor as driving price changes, that factor is likely important at the aggregate level. Heterogeneity in price adjustment behavior—where some firms adjust prices while others do not, despite citing the same factor—is captured through idiosyncratic shocks. The dispersion of these shocks, together with factor-specific thresholds, identifies factor sensitivities or pass-through.

Our framework yields new insights into the time-varying importance of inflation drivers,

including non-labor input costs, demand conditions, and labor costs. We document a multi-phased evolution of inflation during the COVID period, with different forces dominating at different stages. In particular, over this period, inflation dynamics shifted from demand-driven to cost-push, with input prices becoming the dominant narrative as core inflation breached the 2% target. In the later phase, the wage pressure is shown to be a contributor to the inflation.

These narrative factors account for much of recent inflation: in the post-2019 period, the three together explain over 80% of the variation in observed inflation, with non-labor input costs alone accounting for more than 40%, followed by wages, while demand contributes only modestly. The contrast with demand is instructive—it moves substantially over this period, yet firms rarely cite it as the factor that actually triggered a price change. This is what narratives add: by recording which factors firms name as the *cause* of an adjustment, rather than merely which factors moved, they identify the conditional triggers that movements in the underlying series alone cannot reveal.

More broadly, the method developed in this paper enables a reassessment of macroeconomic dynamics using real-time micro narratives. Textual data are often available at higher frequency and over longer horizons than surveys designed to elicit causal explanations. Moreover, real-time narratives help discipline retrospective macroeconomic analysis. The underlying premise is that firms making contemporaneous assessments of their decentralized micro environments possess valuable information about the forces shaping aggregate outcomes. Although the factors deemed relevant by micro decision-makers do not directly reveal the underlying macroeconomic drivers, our model, which explicitly aggregates firms facing idiosyncratic shocks and micro-adjustment frictions, bridges the gap and enables the micro narratives to be used for macroeconomic analysis.

Literature

Our paper is closest to the literature on measuring the “top of the mind” and “whys ” in surveys in the context of firms’ pricing.¹ One particularly relevant paper is [Dogra et al. \(2023\)](#), which measures how firms set prices via open-ended surveys. This also aligns with the earlier literature on directly eliciting firms’ reasons for price adjustments and price stickiness. ([Blinder et al., 1998](#)) Our work complements theirs by taking a further step: we measure such data from texts using LLMs. Our structural model estimation can also be applied to survey data that directly elicit changes in factors and attribution logic. Meanwhile, our framework enables us to analyze narrative statistics over a much longer period. Another important yet subtle difference is that surveys often elicit strategic responses in hypothetical form, whereas textual narratives

¹See [Haaland et al. \(2024\)](#) for a comprehensive survey on this topic.

describe actual conditions and factors. In that regard, our paper is related to, but distinct from, studies and estimates of cost pass-through and the slopes of the Phillips curve using microdata, such as [Gagliardone et al. \(2025\)](#), [Hoeck and Renkin \(2025\)](#).

We contribute to the literature that uses textual analysis and LLM tools ² to study macroeconomic narratives.³ Most of these papers measure intensities of a factor. Our paper differs not only in measuring “how” but also “why”, by extracting attributions. (See [Yang et al. \(2020\)](#) and [Trebbs \(2025\)](#) as recent exceptions.) We also differentiate from the survey-based narrative literature, such as [Andre et al. \(2022\)](#), by focusing on micro mechanisms, whereas they directly elicit beliefs about the macro drivers of inflation. We explicitly filter the micro narrative data through the lens of a structural menu-cost model, admitting idiosyncratic shocks to firms and adjustment costs.

Our paper also provides a different perspective on the drivers of aggregate inflation dynamics, especially the recent post-COVID price surge. The most common empirical approach pursued in answering this question is structural VAR.⁴ We instead take a bottom-up approach by focusing on micro narratives from the firms’ perspectives, which we see as complementary instead of substitutes to the conventional methods. The factors reported to be important to firms are not necessarily the exogenous causes of the macro dynamics. For instance, firms rarely explicitly discuss fiscal/monetary policies in their narratives. Instead, the narratives describe the direct change in activities and prices, possibly stemming from these shocks.

Lastly, various studies ([Zavodny and Ginther, 2005](#); [Armesto et al., 2009](#); [Balke et al., 2017](#); [Kliesen and Werner, 2022](#); [Soto, 2023](#); [Gascon and Martorana, 2024](#)) have developed textual measures based on the Beige Book that exhibit predictive power for macroeconomic activity beyond other data. Meanwhile, [Aruoba and Drechsel \(2024\)](#) uses the Beige book to proxy for the Fed’s information set up to FOMC decisions. In some sense, our paper’s finding that micro narratives in Beige help assess price factors provides one explanation why the Beige Book, despite its anecdotal nature, has predictive power about the future economy.

²Our deployment of LLMs also makes us related to other applications of LLMs in macroeconomic research. For instance, [Bybee \(2023\)](#) “elicits” numeric macroeconomic forecasts made by LLMs by feeding them textual macroeconomic news. [Wu et al. \(2025\)](#) uses an LLM to simulate a retrospective sample of surveys of macroeconomic forecasts that also elicits reasoning that help understand expectation-formation mechanisms. Unlike these papers, this paper instead uses an LLM to extract structural information from texts that can be mapped to a model.

³For instance, [Ash et al. \(2021\)](#); [Larsen and Thorsrud \(2019\)](#); [Bybee et al. \(2024\)](#); [Macaulay and Song \(2022\)](#); [Flynn and Sastry \(2022\)](#).

⁴[Bernanke and Blanchard \(2025\)](#), [Giannone and Primiceri \(2024\)](#), etc.

2 Measuring Price Narratives from Texts

2.1 What do narrative statistics offer?

Before going into the measurement framework, it is useful to first establish the intuition why narrative data in the form of factor attribution is helpful to a correct identification of the macro drivers of inflation. The key insight is that factor-based narratives provide the correct conditioning of the triggers of price adjustment.

Let's assume that the aggregate inflation, which stems from the disaggregated pricing behaviors of individual firms in the economy, can be ultimately represented in a state-space form as follows. (Section 4 formalizes this.) Multiple factors (indexed by m) may simultaneously trigger micro price adjustments, leading to aggregate inflation changes. Due to adjustment frictions such as menu cost, the factor m is actually loaded onto aggregate inflation only conditional on it being strong enough to trigger price adjustment. Therefore, the seemingly linear factor model of inflation is actually nonlinear.

$$\underbrace{\Pi_t}_{\text{Aggregated inflation}} = \theta_0 + \sum_{m=1}^k \underbrace{\theta_m}_{\text{Factor } m \text{ loading}} \underbrace{\mathbf{1}(|X_{m,t}| > k_m)}_{\text{Conditional trigger}} X_{m,t} + \epsilon_t$$

If economists do not actually observe the relevant factor m that triggers the actual loading to the price adjustment, which means that $\mathbf{1}(|X_{m,t}| > k_m)$ are not observed, linear estimation of factor decomposition using movements of all factors $X_{m,t}$ implies a model misspecification, leading to biased estimates of factor loading θ_m . What narrative data can offer, the focus of this paper, is exactly the conditional trigger dynamics that can be attributed to some factors.

2.2 The choice of texts

Our baseline choice of macroeconomic texts is the Beige Book by the Federal Reserve Bank system of the United States. The Beige books are qualitative summaries of district economic conditions based on interviews with business contacts, economists, and market experts by teams from the 12 Federal Reserve Banks. Since 1970, Beige books have been published 8 times a year by the Federal Reserve branches. Our sample includes all regional reports that start from 1990 January and ends in April 2024, which yields 3300 articles. Each article will contain multiple paragraphs regarding different aspects of the economy.

There are several advantages to such types of texts. They are written by professional economists in a dense, structured language typical of economic commentaries. The texts provide rich narrative evidence on the overall and sector-specific economic conditions across different

districts. Across time and space, the Beige Book maintains a stable and consistent tone and writing style. They usually consist of short, self-contained sentences indicating directional changes in the assessed topic, with occasional causal/relational statements linking different descriptions. It is also worth noting that “contacts are not selected at random; rather, Banks strive to curate a diverse set of sources that can provide accurate and objective information about a broad range of economic activities.” This means that they are not a perfect substitute for representative surveys. It is also stated in the Fed Board’s webpage that “the Beige Book is not a commentary on the views of Federal Reserve officials.” Lastly, an extensive literature that is surveyed in the Literature review has shown that the Beige Book provides valuable insights into current macroeconomic conditions that go beyond what is captured by standard macroeconomic indicators. This work also documents that such anecdotal information is widely incorporated into monetary policy decisions.([Aruoba and Drechsel, 2024](#))

Sample Selection. First, although the Beige Book spans many areas of the economy, we omit the real estate and financial sectors from our analysis. These sectors are influenced by regulations, financial structures, and asset-pricing forces that are specific to them, resulting in price dynamics that differ qualitatively from those in goods- and service-producing industries. These differences are also reflected in the terminology and expressions used to describe their developments. Including these sectors would therefore blur the mechanisms we aim to study and introduce noise into our measurement. Second, when constructing aggregate statistics from the Beige Book, we rely solely on the regional reports and leave out the nationwide summaries. This approach prevents double-counting and guarantees that our aggregates capture independently reported regional conditions rather than a consolidated narrative written by the Federal Reserve. Third, we narrow our focus to narratives related to firms. While Beige Book interviews also draw on perspectives from other agents, such as labor organizations, trade groups, and community representatives, our analysis emphasizes firms because they are the main units that set prices and thus provide the most direct evidence on price adjustment behavior.

In each Beige Book report, there are usually one or two paragraphs that focus specifically on price developments. Beyond these, other sections address issues closely tied to pricing behavior, including input costs, labor market conditions, demand strength, and supply-side constraints. Our analysis draws on both categories of information. Finally, with respect to interpretation, we distinguish between mechanisms operating within sectors and those operating across sectors, since we add a sector flag to each identified factor. Unless we explicitly note otherwise, our analysis is primarily concerned with within-sector mechanisms—that is, how prices react to factors and constraints affecting firms in the same sector—rather than with compositional effects driven by changes in the sectoral mix. One exception to our within-sector analysis is the narratives linking energy prices in the energy sector to prices in other sectors, such as

transportation. It turns out that they can be reliably treated as links from the input cost to prices.

How do we think about the Beige Book? We interpret the Beige Book as a collection of firm-level survey responses, each associated with an implicit weight. It is important to emphasize that the anecdotal nature of the Beige Book is widely recognized and generally regarded as a feature rather than a limitation. Nonetheless, some potential caveats are worth noting. The Beige Book is not authored directly by the firms; instead, it is constructed from surveys and qualitative information gathered through contacts that the regional Federal Reserve Banks maintain with firms.

We assume that the Beige Book narratives provide an approximately balanced representation of the views of the surveyed firms. While an aggregation process inevitably occurs in moving from individual firm responses to regional narratives, we continue to treat each regional Beige Book report as a cross-sectional collection of firm-level responses. Furthermore, we assume that the Beige Book authors do not face systematic strategic incentives that would bias the document away from functioning as an unbiased summary of the underlying survey evidence.

2.3 From narratives to the measurement framework

In a standard Beige Book report, firm contacts describe observed changes in a range of relevant factors, operating conditions, and price dynamics. Of particular importance for our analysis is that these narratives not only document the occurrence and magnitude of such changes but also provide qualitative information on their underlying causes as perceived by the firms. For example, in one context respondents may attribute a price increase to higher input costs, whereas in another, they ascribe it primarily to shifts in demand conditions.

More formally, we represent a narrative as a collection of variables (e.g., $d, q, p, \dots$), corresponding to demand, input cost, and price, respectively, their temporal evolution or qualitative dynamics (denoted by --- , $-$, $=$, $+$, $++$), and, when present, potential causal relations between selected pairs of variables (indicated by $\rightarrow$). Consider, for example, the following two narratives. Narrative $A = \{q \uparrow; p \uparrow\}$ encodes that both q and p increase over the period under consideration. Narrative $B = \{q \uparrow \rightarrow p \uparrow\}$ likewise encodes that p rises, but in addition specifies a causal link whereby the change in p is attributed to a rise in q . Thus, while both A and B represent similar observational content (rising price and input cost), only B explicitly posits that the variation in p is caused by the variation in q .

Now consider the third narrative, $C = \{q \uparrow \rightarrow p =\}$. This representation indicates that, at the level of the underlying theoretical structure, a change in q is expected to create a pressure


for p to rise; however, the annotation $p =$ states that, empirically, no such change in p actually occurs, implying that the anticipated causal effect does not materialize in the observed period. By contrast, an alternative narrative $D = \{q \uparrow; p =\}$, in which no causal link between q and p is asserted, implies that the factor q is irrelevant for explaining the behavior of p within the narratives causal framework.

To develop intuition for why narratives containing explicit causal statements provide useful information for identification, consider the following example from the August 1999 report by the Federal Reserve Bank of San Francisco. The narrative states: “Retailers throughout the District noted that stiff competition and falling import prices have kept retail prices down, despite healthy consumer demand and rising labor costs and rising occupancy costs.” If the report merely described declining retail prices, robust demand, and increases in certain input costs, it would be difficult to isolate the underlying causal mechanisms driving these observed changes. Instead, the narrative characterizes a specific economic environment that rules out several alternative explanations for the fall in prices. First, the decline in prices is explicitly not attributed to a reduction in demand; that is, it rules out a leftward shift of the demand curve. Second, input costs are not uniformly declining. Indeed, some, such as labor and occupancy costs, are rising. This pattern suggests that any rightward shift in the supply curve is likely due to particularly falling import prices. Third, the reference to stiff competition is consistent with a relatively elastic (i.e., flat) demand curve faced by individual retailers, implying that small cost or competitive pressures translate into relatively small price adjustments.

When narratives are distilled to reflect a factor structure, they can often be intuitively represented as distinct scenarios within the supply-demand framework. Let’s illustrate with the following four examples.

- **Normal cost-pass-through to price.** May 1996 St. Louis Fed’s: “Contacts continue to report increasing costs for raw materials and increasing costs for labor. In contrast to previous reports, however, some contacts have indicated that these higher costs are being successfully passed along to consumers.” These narratives suggest the increases in good prices are caused by increasing costs. In supply-demand language, such a scenario can be described as a leftward shift of the supply curve along a downward-sloping demand curve, leading to a higher equilibrium price. (Figure 1 a)
- **Limited cost-pass-through to price due to uncertainty.** January 2026 Boston Fed’s Beige-book on prices: “On balance, the pace of selling price increases picked up further but remained moderate. Input price increases accelerated, except among manufacturers, where such increases remained elevated but slowed somewhat. Several contacts noted that price pressures from tariffs were pushing up selling prices and weighing heavily on profitability, while uncertainty was limiting their ability to set prices and plan ahead.” This passage

Figure 1: Examples of narratives



talks about increases in both manufactured goods prices and their input costs. It also reveals that the tariffs contributed (e.g., the use of “pushing up” to the price increase, whereas the factor “uncertainty” serves as a limiting factor of cost pass-through.) In the demand/supply framework, it essentially corresponds to a leftward shift in the supply curve for an unchanged demand curve, where the demand curve is flat due to “uncertainty”.(Figure 1 b)

- **Limited cost-pass-through to price due to weak demand and competition.** November 2011 San Francisco Fed’s: “Price inflation for final goods and services was limited during the reporting period. Contacts noted a recent uptick in the prices for energy inputs, particularly oil, and price increases for assorted food items at the retail level. However, the combination of robust supplier competition and lackluster final demand continued to hold down price pressures for most retail goods and services”. The narrative talked about increases in input cost, while competition and weak demand limit the pass-through. In the supply/demand framework, it corresponds to a leftward shift of the supply curve along a flattened demand curve. (Figure 1 b.)
- **Strong demand-pass-through due to supply constraint.** June 2021 St Louis Fed’s: “Prices have increased moderately overall since our previous report, but prices in some sectors such as automobile retailing and construction have increased sharply due to transitory supply chain constraints and a spike in consumer demand.” This passage attributes the price surge to both strong demand and supply constraint, which is well captured by Figure 1 c.

There are two important clarifications worth making here. First, micro-narratives reveal the factors that firms deem as directly relevant in their micro environments. They need not be exogenous to the economy. They are partial-equilibrium state variables that firms take as

given. We also do not view these descriptions as entirely representative of idiosyncratic changes in circumstances. Rather, they reflect the combined effect of aggregate and local dynamics.

Second, note that it is in line with the spirit of the narratives to represent such narratives with a factor structure, without attempting to differentiate the true exogenous shock versus the response functions. In equilibrium, both shocks and structural parameters enter the ultimate price function. They enter the function both as “factors”. Therefore, in the case of weak demand limiting the cost pass-through, we essentially view the input cost as a pushing factor of the price, while demand is viewed as a limiting factor. A weaker demand in the sense of a flattened demand curve versus a weak demand in the sense of an inward shift demand curve are both treated as a weak demand factor. We formalize these ideas in the latter section.

2.4 LLM-based measurement procedure

Table 1: Multi-dimensional labeling for short sample article entities.

Phrase Dimensions	a slight negative effect on consumer spending(V1)	the resumption of the full Social Security tax(V2)
General	variable	event
Sector	retail	NA
Economic Concept	good sales	tax
Realized or Expected	current realized	current realized
Dynamics	-	++

Our LLM-based measurement procedure involves two steps. First, we prompt the LLM to perform a concept-detection and classification task, identifying economics-related phrases in each paragraph and categorizing them according to a multi-dimensional labeling system. The multiple dimensions include the economic concept, the sector to which it belongs, its dynamics, as well as its temporal stance (about current/past, about future). The following “Economic Variable and Event Extraction Schema” is the multi-dimensional classification framework. In the second step, the LLM determines, for every consequence variable of our focus, i.e., price, production, etc, whether its dynamics are attributed to other identified variables in the surrounding text. Essentially, fixing the consequence as the price, we trace back the narratives in which a “causal attribution” is either explicitly or implicitly mentioned. Both steps rely on carefully designed prompts, which are validated with artificial data. We also provide positive and negative examples in the prompt to inform the LLM of the criteria. Our companion paper (Hou et al., 2025) documents a more general LLM procedure for structuring narratives from texts such as the Beige Book. Here, we only highlight the relevant elements.

Economic Variable and Event Extraction Schema

1. General Classification - Determine whether a phrase is a **Variable** or **Event**. Then follow either 2. Economic Variable Dimensions or 3. Economic Event Dimensions.

2. Economic Variable Dimensions If the entity is a **Variable**, classify it using the following dimensions:

1. **Sector:** *overall, manufacturing, retail, services, energy, agriculture, real estate, transportation*
2. **Economic Concept:** *good demand, good price, good sales, etc.*

3. Economic Event Dimensions If the entity is an **Event**, classify it using the following dimensions:

1. **Economic Concept:** *tax, government spending, rough weather, tariff, monetary policy, war, etc.*

4. Shared Dimensions for Variables and Events Both Variables and Events share the following dimensions:

1. **Realized or Expected:** *past realized, current realized, producer expectation, consumer expectation, expert expectation.*
2. **Dynamics and Value:**
 - *+++* (significant increases/expansion), *++* (moderate increases/expansion), *+* (slight increases/expansion),
 - *---* (significant decreases/shrinking), *--* (moderate decreases/shrinking), *-* (slight decreases/shrinking),
 - *=* (no change or unclear information), *?+* (uncertainty enlarges), *?-* (uncertainty reduces). Note: *?+* and *?-* phrases are identified but not included in the main results.

When identifying causal attributions in the narratives, we instruct the LLM to distinguish between factors that are explicitly linked to price changes and those that are only implicitly related. Explicit attributions typically appear in phrases such as “pass onto,” “attributed to,” “due to,” “to cover,” or “because of,” where firms clearly connect a factor to a price outcome. In contrast, implicit attributions arise when both the factor and the price outcome are reported to have changed, but the firm does not explicitly state a causal link. The degree of explicitness also depends on the reported dynamics. A price change attributed to a factor is conceptually

different from an unchanged outcome. For example, in the statement “Price remained stable, although input costs continued to increase in some industries and retail prices rose in several Districts,” the term “although” signals that input costs are treated as a potential driver of prices, even if the causal link is not stated in a fully direct form. Our LLM labels such logics as “explicit”. Our baseline results and estimation reported below are based on explicitly labeled logics.

Here is one concrete example of demand-to-price logic identified by our procedure and its output: “[Parcel companies have increased shipping prices](p) due to [unprecedented demand from online shopping overall and holiday demand specifically](d).” from 2021 January’s Beige report from St Louis’s Fed, was labeled as an explicit pass-through from demand to the price, both of which are labeled as increased. The narrative identified from this entry is $\{d \uparrow \rightarrow p \uparrow\}$.

The choice of the LLM. Our baseline LLM is GPT-5 (gpt-5-2025-08-07). The knowledge cutoff for the model is Sep 30th, 2024. We use JSON schema to constrain LLM to produce structured output. The JSON schema also restricts LLM to pre-specified categories.

3 Narrative Facts

3.1 Variables

We first show how different words are mapped into concepts such as demand, price, production, input cost, etc. The LLM reasonably labels the phrases according to the context of the discussion. Figure A.1, A.2, and A.3 in the Appendix show wordclouds of the most common phrases under price, demand, and input price within each sector, respectively. One can also see that the phrases used to describe the same concept vary across sectors. For instance, in the transportation sector, prices are surcharges, shipping rates, fares, etc, whereas in real estate, they are rents, rental rates, housing prices, sale prices, property values, etc.

Although certain economic variables are easily identified using common vocabulary, many others are much more context-specific. The linguistic meaning of certain words is not the same as their economic meaning. One good example of such versatility is “demand”. Here are some examples of the phrases that are successfully categorized as describing the demand in our procedure: “price resistance from the buyers”, “less impulse buying”, “intensifying pressure on prices at the retail level”, “sluggish consumer spending”, “increased price sensitivity of consumers”, “consumers remain very value conscious”, “consumer bargain hunting pressures at the retail level”, “consumer resistance to price hikes”, and “raised prices with little push back”. Some examples of the price phrases: “heavy discounting”, “promotions”, “incentives”, or “a show of

goodwill to troubled consumers”.

3.2 Dynamics

We report the dynamic scores of prices and related factors derived from narratives and compare them directly with corresponding realized macroeconomic series. This exercise serves as an external validation of our measurement approach. Despite being constructed from qualitative text, the narrative-based indices closely track aggregate price dynamics. We conduct the comparison at both the national and regional levels (see Figure 2 and 4). In the aggregate, we compare our price and input price narrative scores with the real-time inflation and producer price index growth, respectively. At the regional level, we use the five regional Feds’ manufacturing survey indices, in which firms report conditions, prices paid, prices received, and orders to approximate relevant variables.


At the aggregate level, the narrative inflation index co-moves strongly with official inflation statistics, indicating that shifts in firms reported pricing conditions align well with realized price developments. A similar pattern emerges at the regional level, where cross-regional variation in narrative dynamics mirrors differences in observed inflation outcomes. The high correlations suggest that the dynamics extracted from micro narratives capture economically meaningful movements in price-setting behavior and underlying demand and cost conditions.

Narrative dynamics measure the extensive-margin price adjustment. By construction, narrative measures are discrete and event-based: they primarily capture whether firms report price changes or shifts in pricing pressure, rather than the magnitude of those changes. In this sense, they are more closely aligned with the extensive margin of price adjustment—the incidence of price changes—than with the intensive margin that determines the size of price movements.

Several studies emphasize that firms repricing frequencies are time-varying and state-dependent (Blanco et al., 2024). In particular, recent work documents a pronounced spike in the frequency of price adjustments during the inflation surge beginning in 2022, describing it as a salient yet often underappreciated feature of the episode (Cavallo et al., 2024). Our narrative-based dynamics corroborate this pattern, showing a clear increase in reported price changes during the same period. Figure 3 verifies that the narrative-based dynamics are indeed strongly correlated with the frequency of the price adjustments, based on the statistics of Nakamura et al. (2018) that is extended to 2023 by Montag and Villar Vallenias (2025). Narrative-based price increase measure tracked the frequency of upward price adjustments with a coefficient of 0.66, and the downward price adjustments with a coefficient of 0.31.

The narrative dynamics of several variables other than prices also exhibit intuitive patterns. For instance, upward adjustments in wages, and related variables occur more frequently than

Figure 2: Narrative Directions versus Economic Statistics



downward adjustments, consistent with the presence of downward nominal rigidity. Second, at any point in time, there is substantial dispersion in the measured dynamics for a given variable, reflecting heterogeneity in firms' reported conditions.

Over time, the narrative measures trace distributional shifts across upward, no-change, and downward movements in prices and demand (Figure 5). Notably, during the four recessionary episodes—1990, 2001, 2008, and 2020—the dynamics associated with demand conditions are clearly dominated by downward adjustments relative to upward ones. By contrast, during the persistent inflationary surge beginning in 2020, as well as the temporary inflationary episode preceding 2008, we observe a marked compression of downward dynamics and a pronounced expansion of upward movements. These patterns underscore the state-dependent nature of narrative dynamics across business-cycle episodes.

Figure 3: Narrative Directions versus Frequency of Price Adjustment





Figure 4: Correlation between narrative dynamics with survey indices




Figure 5: Distribution of dynamics of price and demand

(a) Distribution of price dynamics over time



(b) Distribution of demand dynamics over time



3.3 Logics

In this section, we focus on the subset of all identified economic variables in the Beige Book that are directly related to price dynamics, which we collectively refer to as price factors. Factors linked to non-price outcome variables are reported whenever relevant. Among all price factors, non-labor input prices constitute the most frequently observed one, followed in prevalence by demand conditions, competitive pressures, wages, and inventory outcomes (Figure 6). Non-labor input prices are most commonly reported as influencing price developments in the energy, manufacturing, and transportation sectors, as well as at the aggregate (overall economy) level, and appear less frequently in relation to the retail and services sectors. By contrast, demand-related factors play a more prominent role in the manufacturing, retail, and service sectors, in addition to being significant at the aggregate level.

Figure 6: Drivers of price changes by sectors




Figure 7 presents a comparative analysis of three major factors– non-labor input prices, demand, and labor market conditionsand their respective roles in driving changes in prices, output, profits, and wages. Several patterns emerge from this comparison. First, non-labor input costs are primarily cited as determinants of price adjustments, whereas demand-related factors are the predominant drivers of changes in production and other real economic activities. Second, labor market variables are reported as the main contributors to wage adjustments, although they are more often framed analytically as a component of price determination. Third, input price factors and labor market conditions are most frequently associated with upward or no adjustments (i.e., rigidity), while demand factors are invoked in explanations of both upward and downward adjustments.

Figure 7: Different narratives by dynamics






Figure 8: Narrative directions



By decomposing the bilateral relationships between causes and consequences by the direction of change (Figure 8), we can more precisely characterize the directional nature of the underlying impacts. For example, changes in non-labor input costs are much more frequently associated with price adjustments in the same direction, as evidenced by the concentration of positive diagonal entries in the non-labor input price $\rightarrow$ price heatmap. Demand-side factors similarly tend to induce price movements in the same direction as the underlying demand shock. By contrast, wage-to-price passthrough appears relatively infrequent and, when present, is predominantly associated with upward movements in both wages and prices. This pattern is indicative of pro-

nounced downward nominal rigidity in both wages and prices. When considering production (or quantity) as the outcome variable, demand emerges as the most common initiating factor. Increases in input prices, when they have any effect, are typically associated with reductions in output. Overall, these empirical patterns are broadly consistent with the canonical New Keynesian framework, in which firms predominantly adjust prices in response to cost shocks while accommodating demand by adjusting output.


Figure 9: Reasons for price changes



Prevalence of various price factors over time. Figure 9 illustrates the increasing discussion of inflation drivers over time, with a sharp rise following the 2020 pandemic, when inflation reached historically high and persistent levels. These discussions primarily focus on the causes of price hikes, with high input costs emerging as the most significant factor throughout this period, followed by demand and wages. Another period with extensive discussion on high input prices as a driver of inflation occurred between 2005 and 2008, when energy prices rose sharply. The reported drivers of inflation dynamics also exhibit asymmetry. Input prices and wages are primarily cited as reasons for price increases, while demand is most often mentioned as a factor in price declines.

Next, we focus on the pandemic period to examine the key drivers of inflation and wage changes at a monthly frequency. Figure 10 presents the decomposition of inflation and wage drivers mentioned in the Beige Book from January 2019 to November 2023.

Figure 10: Monthly reasons for price changes



In March and April 2020, when strict quarantine measures were imposed, discussions largely centered on declines in inflation and wages, primarily attributed to weak demand and layoffs. This was followed by a rebound in demand, which emerged as a key driver of price increases once restrictions were lifted.

Starting in March 2021, input prices replaced demand as the primary driver of price surges. Around the same time, core CPI began rising above its 2% target. Shortly after the surge in inflation discussions, wage pressure became a prominent topic, beginning in June 2021, with labor demand and shortages cited as the primary causes for wage increases.

Mentions of inflation and wage dynamics began to decline after June 2022. By November 2022, discussions on inflation increases and decreases had become more balanced, and several months later, discussions about wage pressure followed a similar pattern.

This analysis highlights a timing difference in the factors driving inflation. Demand initially received the most attention immediately after COVID restrictions were lifted. However, core CPI only began to surge when input prices became a major concern. Discussions about wage pressure lagged behind the rise in inflation, which is consistent with findings from [Lorenzoni and Werning \(2023\)](#).

4 Structural Model

This section introduces a parsimonious micro environment in which firms set prices subject to menu costs and face (i) a demand/markup shifter and (ii) input cost shifters. The goal is not to solve a fully-fledged general equilibrium; rather, it is to make precise (a) the economic meaning of the latent factors that appear in the estimation model, and (b) the mapping from

structural pass-through parameters to the threshold objects estimated from narrative data. The key ingredient is a non-homothetic demand system that allows aggregate demand conditions to shift markups even holding marginal costs fixed, unlike standard homothetic CES environments (Kimball, 1995; Klenow and Willis, 2016; Baqaee et al., 2024). Detailed derivations are in Appendix B.

4.1 Household's Problem

Time is discrete. A representative household chooses aggregate consumption C_t , labor N_t , and variety-level consumption $\{c_{i,t}\}_{i \in [0,1]}$ to maximize

$$\max_{\{c_{i,t}\}, C_t, N_t} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\ln C_t - e^{\eta_t} \frac{N_t^{1+\nu}}{1+\nu} \right) \quad (1)$$

subject to the budget constraint

$$\int_0^1 P_{i,t} c_{i,t} di = W_t N_t + e^{d_t} + \Pi_t, \quad (2)$$

and an aggregator that is *non-homothetic* across varieties:

$$u(C_t) = \int_0^1 \omega_{i,t} u(c_{i,t}) di \quad (3)$$

Where $u(\cdot)$ is non-homothetic utility function as in Cavallari and Etro (2020) with $u'(\cdot) > 0$ and $u''(\cdot) < 0$. The taste shifter $\omega_{i,t}$ is an idiosyncratic demand shock on variety i . The non-homothetic structure implies that the elasticity and optimal pricing of a firm can vary with aggregate demand shifts (see Appendix B.3). Let λ_t denote the lagrangian multiplier on household's budget constraint. The household problem implies an inverse demand schedule that follows:

$$P_{i,t} = \tilde{\lambda}_t \omega_{i,t} u'(c_{i,t}), \quad (4)$$

where $\tilde{\lambda}_t$ is an *aggregate demand factor* that the firms treat as exogenous in partial equilibrium that follows:

$$\tilde{\lambda}_t \equiv \frac{\beta^t}{C_t u'(C_t) \lambda_t}, \quad (5)$$

$\tilde{\lambda}_t$ is a function of aggregate consumption C_t , which is determined by total income. As a result, both aggregate income shocks d_t and idiosyncratic shock $\omega_{i,t}$ will shift firm's demand curve.⁵

4.2 Individual Firm's Problem

A continuum of monopolistically competitive firms produces using labor ($L_{i,t}$) and non-labor input ($M_{i,t}$):

$$Y_{i,t} = A_{i,t} L_{i,t}^\alpha M_{i,t}^{1-\alpha}, \quad (6)$$

and faces input prices $W_{i,t}$ and $Q_{i,t}$. Firms observe $\{W_{i,t}, Q_{i,t}, \tilde{\lambda}_t, \omega_{i,t}\}$ and take them as given from a partial equilibrium perspective, consistent with the narrative interpretation that firms infer “causes” from observables rather than from the full general equilibrium state.

Let $p_{i,t} \equiv \log P_{i,t}$ and let $p_{i,t}^0 \equiv \log P_{i,t}^0$ denote the (log) *flexible* optimal price absent adjustment frictions. Under standard regularity conditions, the flexible price can be expressed as a markup over marginal cost:

$$p_{i,t}^0 = \mu_{i,t} + mc_{i,t}, \quad (7)$$

where $\mu_{i,t}$ is the (log) desired markup and $mc_{i,t}$ is (log) marginal cost implied by (6). The crucial economic implication of non-homothetic demand is that $\mu_{i,t}$ varies with aggregate demand conditions through $\tilde{\lambda}_t$, so that aggregate demand can shift desired prices even holding marginal costs fixed (Kimball, 1995; Klenow and Willis, 2016; Baqaee et al., 2024).⁶

The linearized flexible price can then be written as:⁷

$$p_{i,t}^0 \equiv \Phi_t + p_{i,t}^* = \beta_\lambda \hat{\lambda}_t + \beta_w w_t + \beta_q q_t + \beta_\omega \hat{\omega}_{i,t} + \beta_w e_{i,t}^w + \beta_q e_{i,t}^q \quad (8)$$

where we assume $\hat{\lambda}_t, w_t, q_t$ are the (linearized) aggregate factors/shocks for demand, wage and non-labor input prices faced by the firms and their summation is denoted by Φ_t . $\{\hat{\omega}_{i,t}, e_{i,t}^w, e_{i,t}^q\}$ are the corresponding idiosyncratic shocks denoted with $p_{i,t}^*$. β 's are the corresponding pass-through parameters that govern how different factors affect flexible price decision. Consistent with our later estimation procedure, we assume for each linearized shock $x \in \{q, w, \hat{\lambda}\}$, it is a linear combination of an aggregate component that follows AR(1) and an idiosyncratic shock

⁵How will aggregate and idiosyncratic demand shocks *affect pricing decision* exactly depends on the form of $u(\cdot)$. When $u(\cdot)$ is homothetic (including standard CES and Kimball preference), aggregate demand factor typically cannot change pricing decision. We show this in the Appendix B.3.

⁶See Appendix B.3 for illustrative examples.

⁷See Appendix B.4 for detailed derivation.

$(e^q, e^w, \hat{\omega})$ that is i.i.d normal with mean zeros:

$$x_{i,t} = x_t + e_{i,t}^x, \quad x_t = \rho_x x_{t-1} + \epsilon_t^x, \quad x \in \{w, q, \hat{\lambda}\}$$

$$e^x \sim N(0, \sigma_x^2), \quad \epsilon_t^x \sim N(0, \sigma_{\epsilon,x}^2), \quad \hat{\omega}_{i,t} \sim N(0, \sigma_\omega^2)$$

4.3 Menu costs and the (S,s) policy

Each period, firm i may reset its nominal price by paying a fixed menu cost ξ . the individual firm solves dynamic loss minimization problem:⁸

$$\min_{\{p_{i,t}\}} E_0 \sum_{t=0}^{\infty} \left[\frac{1}{2} (p_{i,t} - p_{i,t}^0)^2 + \xi \mathbb{1}(p_{i,t} \neq p_{i,t-1}) \right] \quad (9)$$

where $p_{i,t}^0$ follows from (8). Following the canonical menu-cost literature (Sheshinski and Weiss, 1977; Caplin and Spulber, 1987; Auclert et al., 2024), as aggregate shocks (Φ_t) are stationarity and idiosyncratic shocks $(p_{i,t}^*)$ are mean zero and symmetrically distributed, the optimal policy admits an inaction region: there exist thresholds $g^L < 0 < g^U$ such that the optimal pricing for firm i is given by:⁹

$$p_{i,t} = \begin{cases} p_{i,t}^0 = \Phi_t + p_{i,t}^* & \text{if } p_{i,t-1} \notin [g^L + \Phi_t + p_{i,t}^*, g^U + \Phi_t + p_{i,t}^*] \\ p_{i,t-1} & \text{if } p_{i,t-1} \in [g^L + \Phi_t + p_{i,t}^*, g^U + \Phi_t + p_{i,t}^*] \end{cases} \quad (10)$$

5 Estimation Model, Reporting frictions, and Data

From the structural model above, we see multiple factors affect firm's pricing decisions and their pass-throughs into price adjustments vary. Moreover, many of the key factors are frequently mentioned as reasons for firms to adjust prices in our text logic data from Beigebook. In this section, we describe how we can utilize our pricing logic data to estimate a reduced-form structure implied by the previous structural pricing model and help to decompose different factors' contributions to inflation across time. Moreover, we incorporate explicit reporting frictions to accommodate the fact that text is an imperfect and selective measure of firms' decisions and attributions.

⁸Appendix B.4 provides the quadratic approximation that leads to the quadratic loss minimization problem.

⁹See Appendix B.5 for policy characterization.

5.1 From structural pass-through to empirical cutoffs

The empirical analysis classifies text into factor-specific narratives of the form “factor m increased/decreased and led to price adjustment.” To rationalize these statements, we use the factor-*alone* trigger concept: factor m is a sufficient reason for an upward adjustment if, holding other forces fixed at baseline, the component of desired inflation due to m alone pushes the firm outside the inaction region in the direction of an increase. For example, if a reported logic says “*price increase due to demand increase*”, it means the following condition is satisfied for the firm that reported this logic:

$$\beta_\lambda \hat{\lambda}_t + \beta_\omega \hat{\omega}_{i,t} - p_{i,t-1} > -g^L.$$

Now denote demand factor as d , non-labor input factor as c , and wage factor as w . We combine the aggregate and idiosyncratic components of each factor as:

$$s_{i,m,t} \equiv z_{m,t} + v_{i,m,t}, \quad (11)$$

and consider the following mapping:

$$\left\{ \begin{array}{l} z_{d,t} \equiv \hat{\lambda}_t, \quad v_{d,i,t} \equiv \frac{1}{\beta_\lambda} (\beta_\omega \hat{\omega}_{i,t} - p_{i,t-1}), \quad \kappa_{d,+} \equiv \frac{-g^L}{\beta_\lambda}, \quad \kappa_{d,-} \equiv \frac{-g^U}{\beta_\lambda} \\ z_{w,t} \equiv w_t, \quad v_{w,i,t} \equiv \frac{1}{\beta_w} (\beta_w e_{i,t}^w - p_{i,t-1}), \quad \kappa_{w,+} \equiv \frac{-g^L}{\beta_w}, \quad \kappa_{w,-} \equiv \frac{-g^U}{\beta_w} \\ z_{c,t} \equiv q_t, \quad v_{c,i,t} \equiv \frac{1}{\beta_q} (\beta_q e_{i,t}^q - p_{i,t-1}), \quad \kappa_{c,+} \equiv \frac{-g^L}{\beta_q}, \quad \kappa_{c,-} \equiv \frac{-g^U}{\beta_q} \end{array} \right. \quad (12)$$

The optimal pricing decision (10) from the micro-founded menu-cost model boils down to a state-space representation where firms attribute pricing decisions to factor m depending on observing a firm-specific “shock” $s_{i,m,t}$. In particular, the firm will report “*the change of factor m leads to increases of price*” if $s_{i,m,t} > \kappa_{m,+}$ and “*the change of factor m leads to decreases of price*” if $s_{i,m,t} < \kappa_{m,-}$, where the exact wording of the cause depends on the sign of β_m . Equations in (12) are the central links from the micro-founded menu-cost model to our estimation procedure. Holding the inaction band fixed, a larger structural pass-through β_m implies smaller empirical cutoffs $\kappa_{m,\pm}$ and hence faster pass-through on the extensive margin of adjustments.

5.1.1 Assumptions on factors

Consistent with the counter-parts in the menu-cost model, we impose the following assumptions on the shocks/factors in our estimation procedure.

Aggregate component. For $m \in \{d, w, c\}$, let $z_{m,t}$ follow an AR(1) process, where m stands for demand (d), non-labor input cost (c) and wage (w). So $z_{m,t}$ corresponds to the aggregate component x_t in the structural model.¹⁰

$$z_{m,t} = \rho_m z_{m,t-1} + \eta_{m,t}, \quad \eta_{m,t} \sim \mathcal{N}(0, 1), \quad (13)$$

Idiosyncratic component. At the firm level, the latent pricing index is

$$s_{i,m,t}^L = z_{m,t} + v_{i,m,t}, \quad v_{i,m,t} \sim \mathcal{N}(0, \sigma_{m,v}^2), \quad (14)$$

where we assumed the idiosyncratic component of the factor, $v_{i,m,t}$, is normally distributed. This is consistent with the micro-founded model where we will assume all idiosyncratic shocks are normal i.i.d and as a result, the stationary price distribution at any point of time will be Gaussian as well.

5.1.2 Narrative reporting frictions

We model phrase reporting and logic reporting separately. The key idea is that the logic and price blocks remain tied to the micro-consistent index $s_{i,m,t}^L$, while directional phrase mentions are allowed to be noisier measures of the same underlying force.

Noisy phrase index. A firm is recorded as reporting a directional phrase $R_{i,m,t} \in \{+, -, 0\}$ according to a noisy phrase index

$$s_{i,m,t}^R = z_{m,t} + v_{i,m,t} + \xi_{i,m,t}, \quad \xi_{i,m,t} \sim \mathcal{N}(0, \sigma_{m,\xi}^2), \quad (15)$$

¹⁰We normalize $\text{Var}(\eta_{m,t}) = 1$ because scales will be absorbed by $(\sigma_{m,v}, \sigma_{m,\xi}, \kappa_{m,\pm}, \psi_m^\pm)$ below and are not separately identifiable. The level of the latent factor $z_{m,t}$ also does not matter for the decomposition exercise we consider.

with $v_{i,m,t} \perp \xi_{i,m,t}$. Phrase reports are generated by a factor-specific salience threshold $d_m > 0$:

$$R_{i,m,t} = \begin{cases} + & \text{if } s_{i,m,t}^R > d_m, \\ - & \text{if } s_{i,m,t}^R < -d_m, \\ 0 & \text{otherwise.} \end{cases} \quad (16)$$

The parameter d_m captures how strong a factor must be to be mentioned in text, while $\sigma_{m,\xi}$ allows phrases to be noisier than the underlying pricing logic.

Logic within phrase branches. Logic is observed only within the two directional phrase branches $R_{i,m,t} = +$ and $R_{i,m,t} = -$. Within each branch, the logic outcome is binary:

- *A*: active with desired direction;
- *No*: not-*A*.

Structural cutoffs satisfy $\kappa_{m,+} > d_m$ and $\kappa_{m,-} > d_m$. In the positive branch, a true active logic occurs if $s_{i,m,t}^L \geq \kappa_{m,+}$; in the negative branch, a true active logic occurs if $s_{i,m,t}^L \leq -\kappa_{m,-}$. Because phrase selection is determined by $s_{i,m,t}^R = z_{m,t} + v_{i,m,t} + \xi_{i,m,t}$ while logic status is determined by $s_{i,m,t}^L = z_{m,t} + v_{i,m,t}$, the correct conditional logic probabilities are bivariate-normal selection probabilities. Appendix C.2 gives the closed-form expressions.

Let $N_{+,m,t}, N_{0,m,t}, N_{-,m,t}$ denote the observed phrase counts, and let $M_{A,+,m,t}$ and $M_{A,-,m,t}$ denote the observed *A*-counts within the positive and negative phrase branches. Then the phrase block is multinomial, while the logic block is binomial conditional on branch exposure.

5.2 Inflation measurement and narrative-based decomposition

This subsection links the latent factors filtered from narrative data to observed aggregate inflation. A key motivation for introducing inflation as an additional measurement is that the narrative layer is informative about the *extensive margin* of pricing, but it does not by itself pin down the inflation units implied by these logics. We therefore use the micro-founded pricing rule in Section 4 to motivate an aggregate inflation equation that depends on the same latent threshold-crossing objects identified from the narrative data.

Inflation as a state-space measurement. Conditional on the latent factor states $z_t \equiv \{z_{m,t}\}_{m \in \{d,w,c\}}$, the cross-sectional signal structure in (14) implies a well-defined notion of factor-*m* “active adjustment intensity”: the mass of firms whose latent index crosses the active-logic cutoff,

weighted by the magnitude of the latent index when it crosses. We summarize these objects by the truncated-moment functions

$$M_{m,t}^+(z_{m,t}) \equiv \mathbb{E}[s_{i,m,t}^L \mathbf{1}\{s_{i,m,t}^L \geq \kappa_{m,+}\} \mid z_{m,t}], \quad (17)$$

$$M_{m,t}^-(z_{m,t}) \equiv \mathbb{E}[s_{i,m,t}^L \mathbf{1}\{s_{i,m,t}^L \leq -\kappa_{m,-}\} \mid z_{m,t}]. \quad (18)$$

Appendix C.5 derives these objects from the micro pricing rule and provides their closed-form expressions under normality.

The link to inflation is as follows. Under the fixed-band approximation to the menu-cost model in (10), a firm's realized price change is the desired price gap multiplied by an adjustment indicator. Equation (12) shows that, holding other forces at baseline, the factor- m component of the desired price gap can be written in terms of the same latent index $s_{i,m,t}^L$ and the cutoff values $\kappa_{m,+}$ and $\kappa_{m,-}$. As shown formally in Appendix C.5, the exact joint inflation map implied by the micro model admits an additive anchored approximation in which the factor-specific main effect is given by the factor-alone counterfactual. This yields a benchmark inflation contribution proportional to $M_{m,t}^+(z_{m,t}) + M_{m,t}^-(z_{m,t})$, so that the benchmark loading is symmetric across upward and downward threshold crossings for a given factor.

Motivated by this micro-to-aggregation link, we write the narrative-implied inflation component as

$$P_t^{\text{nar}} = \theta_0 + \sum_{m \in \{d,w,c\}} \psi_m [M_{m,t}^+(z_{m,t}) + M_{m,t}^-(z_{m,t})]. \quad (19)$$

The parameter ψ_m should be interpreted as the reduced-form inflation loading associated with factor m . In the exact factor-alone benchmark it coincides with the structural pass-through coefficient up to normalization. In the empirical specification it also absorbs approximation error from replacing the full joint pricing problem by the additive factor-alone representation.

Residual inflation dynamics. Observed inflation also reflects forces not captured by our three narrative factors. From our menu-cost model, these include productivity and other residual markup/cost movements. Practically, they also include any discrepancy between the factor-alone approximation and the full multi-factor pricing rule. We capture these components with a latent residual inflation factor u_t :

$$u_t = \rho_u u_{t-1} + \eta_{u,t}, \quad \eta_{u,t} \sim \mathcal{N}(0, \sigma_u^2), \quad (20)$$

and allow for measurement noise in inflation:

$$P_t = P_t^{\text{nar}} + u_t + \varepsilon_t^P, \quad (21)$$

$$\varepsilon_t^P \sim \mathcal{N}(0, R_P). \quad (22)$$

Decomposition given filtered factors. Let $\hat{z}_{m,t|t}$ denote the filtered estimate of $z_{m,t}$ from the narrative measurement equations in Section 5.1.2. The model-implied inflation contribution attributed to factor m is

$$P_{m,t} = \psi_m [M_{m,t}^+(\hat{z}_{m,t|t}) + M_{m,t}^-(\hat{z}_{m,t|t})], \quad (23)$$

so that the fitted inflation and decomposition satisfy

$$\hat{P}_t = \theta_0 + \sum_{m \in \{d, w, c\}} P_{m,t} + \hat{u}_{t|t}. \quad (24)$$

The decomposition is economically interpretable by construction: $\sum_m P_{m,t}$ is the part of inflation disciplined by the narrative-implied threshold-crossing structure of the three factors, while u_t absorbs the inflation component attributable to non-classified forces and residual measurement error.

6 Estimation Results

We estimate the structure described in the previous sections as a nonlinear state-space model whose unknowns consist of both model parameters and time-varying latent states/factors that affect price adjustments. The observations entering estimation are the phrase and logic frequencies extracted from Beige Book narratives, aggregate inflation, and macroeconomic variables that serve as linear Gaussian measurements of the latent factors. Let X_t denote the vector of aggregate states/factors and Y_t the observed measurements. Then

$$X_t = AX_{t-1} + \varepsilon_t^X \quad (25)$$

$$Y_t = CX_t + \varepsilon_t^Y, \quad (26)$$

where

$$\varepsilon_t^X \sim \mathcal{N}(0, Q), \quad \varepsilon_t^Y \sim \mathcal{N}(0, R_Y) \quad (27)$$

and (25) follows from (13), so Q is an identity matrix.¹¹

To jointly estimate the parameters and filter the latent states, we employ a generalized EM algorithm (GEM) to perform maximum likelihood estimation following Dempster et al. (1977) and Wu (1983).¹²

We combine information from linear Gaussian signals, phrases, logics, and the aggregate inflation index to construct the likelihood and infer the latent states. Moreover, as shown in (15), the phrase index contains additional reporting noise relative to the logic block, so phrase counts are treated as an over-dispersed composition block. To account for this, we use a weighted composite likelihood approach in the sense of Varin et al. (2011). Specifically, we first estimate the model using the state law of motion and the logic block alone; this first step has a conditional-likelihood interpretation in the sense discussed by Reid (2024). We then fit an auxiliary Dirichlet-multinomial calibration factor by factor and map the resulting concentration parameter into an effective sample size for the phrase block, following Thorson et al. (2017). This yields a factor-specific weight $\lambda_{gate,m} \in (0, 1]$ on the phrase likelihood that determines how strongly phrase counts enter the final objective.

The parameters to be estimated include the factor-specific parameters $\{\rho_m, \sigma_{m,v}, \sigma_{m,\xi}, d_m, \kappa_{m,+}, \kappa_{m,-}, \psi_m\}$ for each factor $m \in \{d, w, c\}$, as well as the inflation- and Gaussian-signal-specific parameters $\{\theta_0, \rho_u, \sigma_u^2, R_P, C, R_Y\}$. Denote the full parameter vector by θ . The final weighted composite likelihood is

$$\ell_{CL}(\theta) = \ell_X(\theta) + \ell_u(\theta) + \ell_Y(\theta) + \ell_L(\theta) + \ell_P(\theta) + \sum_{m \in \{d, w, c\}} \lambda_{gate,m} \ell_{R,m}(\theta), \quad (28)$$

where $\ell_X(\theta)$ is the state-transition contribution, $\ell_u(\theta)$ is the likelihood contribution of the residual inflation state, $\ell_Y(\theta)$ is the Gaussian signal block, $\ell_L(\theta)$ is the logic block, $\ell_P(\theta)$ is the inflation measurement block, and $\ell_{R,m}(\theta)$ is the phrase likelihood contribution for factor m . Appendix C provides the derivations of the likelihood blocks and the construction of the weights used in the composite likelihood. Using the parameter estimates and smoothed states/factors from the GEM procedure, we then perform a model-based decomposition exercise to assess how each factor contributes to observed aggregate inflation.

¹¹For identification, we normalize the innovation covariance matrix of the latent factors to the identity, $Q = I$, which is standard in the factor-model literature following Bai and Ng (2002). We also restrict the Gaussian signal loading matrix C to be diagonal. This is a stronger exclusion restriction that assigns each Gaussian proxy to a single latent factor and thereby fixes the rotation in a way that is consistent with our measurement interpretation.

¹²Because our model features a nonlinear state-space representation, we implement the ECM algorithm following Meng and Rubin (1993).

Data In our baseline results, we focus on three factors that may affect inflation in the menu-cost model in Section 4: demand, wage, and non-labor input cost. The estimation uses:

- frequencies of phrases about the three factors: for each m and t , the counts or fractions of texts classified as $R_{m,t} \in \{+, -, 0\}$;
- logic frequencies within directional phrase branches: for each m and t , we label the logic as A (“activate”) if the logic indicates that this factor affects inflation. Otherwise, it is labeled as not- A . We then count their frequencies conditional on $R_{m,t} \in \{+, -\}$;
- aggregate inflation: P_t , where we use monthly headline CPI;
- linear Gaussian measurements of factors: we use the monthly change in PCE as the measurement of demand, PPI as the measurement of non-labor input cost, and wage inflation as the measurement of wage/labor cost.

Estimated Parameters Table 2 reports the estimated parameters. We group them according to whether they are relevant for logic and phrase reports, the inflation equation (21), or the linear Gaussian measurements (26).

Table 2: Estimated Parameters

factor	Logic/State block						Price block			Signal block
	ρ_m	$\sigma_{v,m}$	$\sigma_{\xi,m}$	d_m	κ_m^+	κ_m^-	ψ_m	θ_0	R_p	c_m
demand	0.6069 (0.063)	3.4970 (0.249)	5.0695 (0.375)	0.9808 (0.053)	8.8480 (0.571)	7.7287 (0.520)	0.4200 (0.171)	0.1468 (0.071)	0.2869 (0.025)	0.1433 (0.035)
non-labor input cost	0.8315 (0.032)	2.4874 (0.164)	0.0238 (0.241)	0.5501 (0.039)	5.5716 (0.362)	5.6167 (0.373)	0.7024 (0.073)	0.1468 (0.071)	0.2869 (0.025)	0.2189 (0.023)
wage	0.9517 (0.018)	9.1218 (1.267)	18.1277 (1.606)	6.7867 (0.410)	25.1656 (3.460)	N/A N/A	0.8114 (0.514)	0.1468 (0.071)	0.2869 (0.025)	0.3279 (0.004)

Bootstrapped standard errors are reported in brackets.

Several patterns emerge from the estimates. First, the thresholds κ differ sharply across factors. Recall that in the estimation model, holding the aggregate factor $z_{m,t}$ at mean zero, $[-\frac{\kappa_m^-}{\sigma_{v,m}}, \frac{\kappa_m^+}{\sigma_{v,m}}]$ captures the effective “inaction region” of price sensitivity to factor m . The estimates of κ_m and $\sigma_{v,m}$ suggest that price adjustment is relatively more sensitive to non-labor input costs than to demand and wages when these factors exert upward pressure on prices. This captures the *extensive margin* of each factor’s pass-through into inflation. Note that we cannot estimate κ_m^- for wages because in our dataset there are no logic reports indicating that falling wages lead to price declines.

Second, the estimates of ψ_m capture the *intensive margin* of pass-through, as they measure the average factor-alone impact on inflation conditional on that factor triggering price movements.

The estimates suggest that the intensive margin of pass-through is strongest for wages, similar in magnitude for non-labor input costs, and weakest for demand.¹³

Third, the estimates $\sigma_{\xi,m}$ capture how “noisy” the phrase reports are relative to the logic reports. This reveals two important aspects of the text data. First, phrase counts can overstate the intensity of the underlying factor of interest, and this overstatement varies across factors. In our setting, the object of interest is the aggregate and idiosyncratic factors that matter for firms’ price adjustments. The parameter $\sigma_{\xi,m}$ captures the discrepancy between phrases and logic data in measuring those factors. A large positive estimate of $\sigma_{\xi,m}$ implies noisier phrase reports and therefore lowers the informativeness of phrase counts for inferring the latent states/factors. Second, $\sigma_{\xi,m}$ also captures how the raw frequency of logic reports produced by our text measurement procedure may understate the contribution of some factors, which represents a specific type of “reporting friction.” In our estimation model, we observe price-related logic reports only conditional on observing directional phrases. A larger $\sigma_{\xi,m}$ makes directional phrase mentions more likely, but many of these mentions do not translate into price logics because they are triggered by $\xi_{i,m,t}$ rather than $v_{i,m,t}$. This mechanically inflates the denominator when measuring conditional logic frequencies. Accordingly, the large estimates of $\sigma_{\xi,m}$ for demand and wage factors capture these reporting frictions and properly down-weight the phrase counts, making the price decomposition robust to such frictions in text measurement. Moreover, the sharp contrast in estimated $\sigma_{\xi,m}$ between non-labor input costs and the other two factors suggests that text measurements of logic and phrases about input costs are much more consistent and informative about the underlying factor that matters for price adjustment.


Finally, the parameters ρ_m capture the estimated persistence of each factor, while c_m captures the informativeness of the corresponding Gaussian signal.

Uncovered States/Factors Another object recovered from the estimation is the set of smoothed states/factors inferred from the joint likelihood of the observed data. These states represent the aggregate factors in the menu-cost model that serve as the time-specific center of the idiosyncratic factor each firm receives. Intuitively, the higher a given state is at a point in time, the larger the mass of firms that will adjust prices accordingly, and the more logic reports will be observed. The interaction between the aggregate state and the inaction thresholds κ governs the extensive margin of factor-alone pass-through into inflation. Figure 11 depicts the smoothed states of the three factors from the estimated model. We see clearly that non-labor input costs exhibit persistently positive deviations during 2005–2008 and the pandemic episode, whereas the demand factor is centered around zero, drops sharply during the COVID-19 lockdowns, and

¹³Note that from (19) the menu-cost model imposes symmetric loadings ψ_m on positive and negative factor-alone pass-through. Empirically, we can allow for asymmetric loadings, but the asymmetric component is usually only weakly identified. For this reason, we use symmetric loadings in the baseline results.

turns positive again after the reopening. These patterns are consistent with the phrase and logic reports in the text data presented in Sections 3.2 and 3.3.

Figure 11: Uncovered smoothed states/factors




Inflation Decomposition We now turn to the main exercise of inflation decomposition under the estimated model. Using equation (24), we attribute observed month-to-month headline inflation to the model’s three factors. Figure 12 reports the results. Panel (a) covers the full sample from 1990 to 2024, while panel (b) focuses on the COVID-19 episode after 2019. The red, yellow, and blue bars denote the contributions of non-labor input costs, wages, and demand to observed inflation, while the purple bars denote the residual component u_t in (24). During the pandemic, demand explains most of the decline in inflation in the three lockdown months. Once the lockdowns ended, demand contributed to the inflation increases from July 2020 to April 2022; however, its contribution was not comparable to that of non-labor input costs and wages. At the onset of the inflation surge, up to November 2021 when inflation peaked, non-labor input costs were the dominant contributor to inflation increases, averaging more than 1% month-to-month inflation. After the summer of 2022, their contribution declined rapidly as inflation began to fall, while the slower-moving wage component became the dominant force keeping inflation elevated for a period. Meanwhile, demand’s contribution turned negative after the end of 2022.

Figure 12: Decomposition of Inflation



(a) Full Sample



(b) Post 2019 Sample

To quantify the relative importance of these three factors more precisely, Table 3 reports the covariance shares of the factors and the residual on the full sample and the post-2019 subsample. Over the full sample, non-labor input costs account for 25.85% of the variation in observed inflation, almost five times the contribution of demand and roughly twice that of wages. After 2019, the three factors together explain more than 80% of the variation in observed inflation, with non-labor input costs accounting for more than 40% of total inflation. These findings are consistent with both the estimated parameters and the smoothed states discussed above: the narrative data suggest that price adjustment is most sensitive to non-labor input costs, and that this factor was the most important contributor to post-pandemic inflation. Wages also play a significant role, especially after the pandemic, whereas demand has only a

limited direct effect on inflation.¹⁴

Table 3: Covariance shares explained by each factor

sample	demand	non-labor input cost	wage	residual
full sample	0.0522	0.2585	0.1217	0.5676
post-2019	0.1386	0.4196	0.2583	0.1836


Finally, our inflation decomposition explicitly takes into account the reporting friction that can understate the fraction of firms adjusting prices in response to demand factors because of the over-dispersed phrase-report distribution. It also incorporates both the extensive and intensive margins of price adjustment in the estimation model. To illustrate these points more clearly, we use the estimated model to predict the frequency with which each factor triggers price adjustments, purging the reporting friction captured by $\sigma_{\xi,m}$ and highlighting the role of the extensive margin in the inflation decomposition. Figure 13 shows the model-predicted frequency with which factor m leads to price adjustment. Both demand and wage are triggered relatively infrequently because they have larger thresholds κ_m and aggregate states that are, on average, closer to zero, as shown in Figure 11. Moreover, over time, the demand factor is often triggered in a fairly symmetric way: many firms increase prices while many others decrease prices at the same time, which further dampens the aggregate contribution of demand to inflation. Finally, the model-predicted logic frequencies in panel (b) line up closely with the narrative evidence reported in Figure 10.

7 Conclusion


In this paper, we develop a framework for quantifying micro-level narratives that admit a factor-model representation and employ these narratives to estimate a structural model of firm-level price setting. Using a model that explicitly aggregates individual firms' pricing decisions in the presence of menu-cost frictions, we isolate the underlying macroeconomic determinants of inflation. This approach offers a complementary perspective on decomposing macroeconomic dynamics, alongside conventional time-series methods based on structural identification restrictions.

¹⁴One caveat of our results is that we impose symmetric loadings on positive and negative price adjustments for each factor. This contributes to the fact that when inflation falls sharply, our factors do not explain those declines as well. We do have a set of results allowing for asymmetric loadings. However, the asymmetric estimates are only weakly identified, so we use the symmetric specification as the baseline. Our qualitative inflation decomposition results remain unchanged when asymmetric loadings are allowed.

Figure 13: Frequency of price adjustment triggered by each factor



(a) Full Sample



(b) Post 2019 Sample

We conceptualize micro narratives—firms qualitative descriptions of changes in their own operating conditions and price outcomes—as distinct from macro narratives, which pertain directly to the evolution of aggregate macroeconomic variables and the propagation of exogenous shocks. (Andre et al., 2022) Micro narratives are more immediately pertinent to firms decision-making processes. While factor representations of micro price narratives can, in principle, be directly elicited through survey instruments, we instead employ a large language model (LLM) to systematically extract and structure such information from a long-run historical corpus of economic texts, namely the Federal Reserves Beige Book. Our application thereby demonstrates the substantial informational content and analytical value of narrative data for the estimation of structural macroeconomic models.

References

- Andre, Peter, Ingar Haaland, Christopher Roth, and Johannes Wohlfart**, “Narratives about the Macroeconomy,” *Working Paper*, 2022.
- Armesto, Michelle T, Rubén Hernández-Murillo, Michael T Owyang, and Jeremy Piger**, “Measuring the information content of the beige book: A mixed data sampling approach,” *Journal of Money, Credit and Banking*, 2009, *41* (1), 35–55.
- Aruoba, S Borağan and Thomas Drechsel**, “Identifying monetary policy shocks: A natural language approach,” Technical Report, National Bureau of Economic Research 2024.
- Ash, Elliott, Germain Gauthier, and Philine Widmer**, “Text semantics capture political and economic narratives,” *arXiv preprint arXiv:2108.01720*, 2021.
- Auclert, Adrien, Rodolfo Rigato, Matthew Rognlie, and Ludwig Straub**, “New pricing models, same old Phillips curves?,” *The Quarterly Journal of Economics*, 2024, *139* (1), 121–186.
- Bai, Jushan and Serena Ng**, “Determining the Number of Factors in Approximate Factor Models,” *Econometrica*, 2002, *70* (1), 191–221.
- Balke, Nathan S, Michael Fulmer, and Ren Zhang**, “Incorporating the beige book into a quantitative index of economic activity,” *Journal of Forecasting*, 2017, *36* (5), 497–514.
- Baqae, David R, Emmanuel Farhi, and Kunal Sangani**, “The supply-side effects of monetary policy,” *Journal of Political Economy*, 2024, *132* (4), 1065–1112.
- Bernanke, Ben and Olivier Blanchard**, “What caused the US pandemic-era inflation?,” *American Economic Journal: Macroeconomics*, 2025, *17* (3), 1–35.
- Bertoletti, Paolo and Federico Etro**, “Monopolistic Competition when Income Matters,” *The Economic Journal*, 2017, *127* (603), 1217–1243.
- Blanco, Andrés, Corina Boar, Callum J Jones, and Virgiliu Midrigan**, “The inflation accelerator,” Technical Report, National Bureau of Economic Research 2024.
- Blinder, Alan, Elie RD Canetti, David E Lebow, and Jeremy B Rudd**, *Asking about prices: a new approach to understanding price stickiness*, Russell Sage Foundation, 1998.
- Bybee, J Leland**, “The ghost in the machine: Generating beliefs with large language models,” *arXiv preprint arXiv:2305.02823*, 2023.
- Bybee, LELAND, BRYAN KELLY, ASAF MANELA, and DACHENG XIU**, “Business News and Business Cycles,” *The Journal of Finance*, 2024, *79* (5), 3105–3147.

- Caplin, Andrew S and Daniel F Spulber**, “Menu costs and the neutrality of money,” *The Quarterly Journal of Economics*, 1987, 102 (4), 703–725.
- Cavallari, Lilia and Federico Etro**, “Demand, markups and the business cycle,” *European Economic Review*, 2020, 127, 103471.
- Cavallo, Alberto, Francesco Lippi, and Ken Miyahara**, “Large shocks travel fast,” *American Economic Review: Insights*, 2024, 6 (4), 558–574.
- Dempster, A. P., N. M. Laird, and D. B. Rubin**, “Maximum Likelihood from Incomplete Data Via the EM Algorithm,” *Journal of the Royal Statistical Society: Series B (Methodological)*, 09 1977, 39 (1), 1–22.
- Dogra, Keshav, Sebastian Heise, Edward S Knotek, Brent Meyer, Robert W Rich, Raphael Schoenle, Giorgio Topa, Wilbert van der Klaauw, and Wändi Bruine de Bruin**, “Estimates of Cost-Price Passthrough from Business Survey Data,” *FRB of New York Staff Report*, 2023, (1062).
- Flynn, Joel P and Karthik Sastry**, “The Macroeconomics of Narratives,” *Available at SSRN 4140751*, 2022.
- Gagliardone, Luca, Mark Gertler, Simone Lenzu, and Joris Tielens**, “Micro and macro cost-price dynamics in normal times and during inflation surges,” Technical Report, National Bureau of Economic Research 2025.
- Gascon, Charles S and Joseph Martorana**, “The Beige Book and the Business Cycle: Using Beige Book Anecdotes to Construct Recession Probabilities,” Technical Report 2024.
- Gentzkow, Matthew, Bryan Kelly, and Matt Taddy**, “Text as data,” *Journal of Economic Literature*, 2019, 57 (3), 535–74.
- Giannone, Domenico and Giorgio Primiceri**, “The drivers of post-pandemic inflation,” Technical Report, National Bureau of Economic Research 2024.
- Haaland, Ingar, Christopher Roth, Stefanie Stantcheva, and Johannes Wohlfart**, “Measuring what is top of mind,” Technical Report, ECONtribute Discussion Paper 2024.
- Hoeck, Christian Philip and T Renkin**, “Demand shocks and prices—Micro evidence and macro implications,” Technical Report, Working Paper 2025.
- Hou, Chenyu, Jiannan Jiang, and Tao Wang**, “Reading between Lines: Measuring Macroeconomic Narratives from Texts using Large Language Models,” Technical Report, mimeo 2025.

- Kimball, Miles S**, “The Quantitative Analytics of the Basic Neomonetarist Model,” *Journal of Money, Credit and Banking*, 1995, 27 (4), 1241–1277.
- Klenow, Peter J and Jonathan L Willis**, “Real rigidities and nominal price changes,” *Economica*, 2016, 83 (331), 443–472.
- Kliesen, Kevin L and Devin Werner**, “Using Beige Book text analysis to measure supply chain disruptions,” *Economic Synopses*, 2022, (18).
- Larsen, Vegard and Leif Anders Thorsrud**, “Business cycle narratives,” 2019.
- Lorenzoni, Guido and Iván Werning**, “Wage-Price Spirals,” *Brookings Papers on Economic Activity*, 2023.
- Macaulay, Alistair and Wenting Song**, “Narrative-Driven Fluctuations in Sentiment: Evidence Linking Traditional and Social Media,” *Available at SSRN 4150087*, 2022.
- Matsuyama, Kiminori**, “Non-CES Aggregators: A Guided Tour,” *Annual Review of Economics*, 2023, 15 (1), 235–265.
- Meng, Xiaoli and Donald B. Rubin**, “Maximum likelihood estimation via the ECM algorithm: A general framework,” *Biometrika*, 06 1993, 80 (2), 267–278.
- Montag, Hugh and Daniel Villar Vallenás**, “Post-pandemic price flexibility in the US: evidence and implications for price setting models,” 2025.
- Nakamura, Emi, Jón Steinsson, Patrick Sun, and Daniel Villar**, “The elusive costs of inflation: Price dispersion during the US great inflation,” *The Quarterly Journal of Economics*, 2018, 133 (4), 1933–1980.
- Reid, N**, “On partial likelihood,” *Journal of the Royal Statistical Society Series A: Statistics in Society*, 08 2024, 187 (3), 567–577.
- Sheshinski, Eytan and Yoram Weiss**, “Inflation and costs of price adjustment,” *The Review of Economic Studies*, 1977, 44 (2), 287–303.
- Soto, Paul E**, “Measurement and effects of supply chain bottlenecks using natural language processing,” 2023.
- Thorson, James T., Kelli F. Johnson, Richard D. Methot, and Ian G. Taylor**, “Model-based estimates of effective sample size in stock assessment models using the Dirichlet-multinomial distribution,” *Fisheries Research*, 2017, 192, 84–93. Data conflict and weighting, likelihood functions, and process error.

- Trebbi, Giovanni**, “Inflation narratives and expectations,” 2025.
- Varin, Cristiano, Nancy Reid, and David Firth**, “AN OVERVIEW OF COMPOSITE LIKELIHOOD METHODS,” *Statistica Sinica*, 2011, *21* (1), 5–42.
- Wu, C. F. Jeff**, “On the Convergence Properties of the EM Algorithm,” *The Annals of Statistics*, 1983, *11* (1), 95 – 103.
- Wu, Jing Cynthia, Jin Xi, and Shihan Xie**, “LLM Survey Framework: Coverage, Reasoning, Dynamics, Identification,” Technical Report, National Bureau of Economic Research 2025.
- Yang, Yucheng, Yue Pang, Guanhua Huang et al.**, “The knowledge graph for macroeconomic analysis with alternative big data,” *arXiv preprint arXiv:2010.05172*, 2020.
- Zavodny, Madeline and Donna K Ginther**, “Does the Beige Book move financial markets?,” *Southern Economic Journal*, 2005, *72* (1), 138–151.

A Additional Empirical Results

A.1 Phrases labeled as different concepts

Figure A.1, A.2, and A.3 plot the wordclouds of the phrases in the Beige Book that are labeled as “price”, “input price”, and “demand”, respectively. They show the versatility of the language used to describe the same economic variables in different sectors.

Figure A.1: Price phrases



The most common economic variables in Beige Book. Figure A.4 plots the economic variables ranked by their total frequency in the sample. The most common ten variables are “sales”, “production”, “price”, “demand”, “market conditions”, “non-labor input cost”, “wage”, “employment”, “labor demand”, “inventory outcome”.

Word Clouds for Other Input Price for Production by Sector



Figure A.3: Demand phrases





Figure A.4: The common economic variables by sector



B Micro model derivations

This appendix collects the detailed derivations for the menu-cost model.

B.1 Household Problem

We take the household's problem as written in (1) subject to (2) and (3). The first order conditions give:

$$[c_{i,t} :] \quad \frac{\beta^t \omega_{i,t} u'(c_{i,t})}{C_t u'(C_t)} = \lambda_t P_{i,t} \quad [N_t :] \quad \beta^t e^n N_t^\nu = \lambda_t W_t \quad (29)$$

where λ_t is the lagrangian multiplier. Following the definition of aggregate demand factor in (5), the household's problem leads to:

$$\begin{cases} W_t &= \tilde{\lambda}_t C_t u'(C_t) e^n N_t^\nu \\ P_{i,t} &= \tilde{\lambda}_t \omega_{i,t} u'(c_{i,t}) \\ \int_0^1 P_{i,t} c_{i,t} di &= W_t N_t + e^{dt} + \Pi_t \end{cases} \quad (30)$$

B.2 Firm problem under flexible prices

The individual firm i observes $\{W_{i,t}, Q_{i,t}, \tilde{\lambda}_t, \omega_{i,t}\}$ and takes them and the demand curve as given to solve the cost minimization problem:

$$\min_{L_{it}, M_{it}} \quad W_{i,t} L_{it} + Q_{i,t} M_{it} \quad (31)$$

$$s.t. \quad Y_{it} = A_t L_{it}^\alpha M_{it}^{1-\alpha} \quad (32)$$

This gives optimal ratio of input:

$$\frac{M_{it}}{L_{it}} = \frac{1-\alpha}{\alpha} \frac{W_{i,t}}{Q_{i,t}} \quad (33)$$

Combine with production function:

$$L_{i,t} = \frac{Y_{it}}{A_t} \left(\frac{\alpha}{1-\alpha} \frac{Q_{i,t}}{W_{i,t}} \right)^{1-\alpha}, \quad M_{it} = \frac{Y_{it}}{A_t} \left(\frac{1-\alpha}{\alpha} \frac{W_{i,t}}{Q_{i,t}} \right)^\alpha \quad (34)$$

This gives total cost in terms of Y_{it} which leads to marginal cost:

$$MC_{it} = \frac{1}{A_t} \left(\frac{W_{i,t}}{\alpha} \right)^\alpha \left(\frac{Q_{i,t}}{1-\alpha} \right)^{1-\alpha}$$

Firm's optimal pricing: The individual firm i would choose $P_{i,t}$, given the demand curve (with demand $c_{i,t}$), taking into consideration that it will produce to match demand, i.e. $Y_{i,t} = c_{i,t}$. That means the firms are not holding inventories etc.

$$\begin{aligned} \max_{P_{i,t}} \quad & P_{i,t}Y_{i,t} - MC_{i,t}Y_{i,t} \\ \text{s.t.} \quad & P_{i,t} = \tilde{\lambda}_t \omega_{i,t} u'(Y_{i,t}) \end{aligned} \quad (35)$$

Before we solve the optimal pricing problem, note that from the demand curve, the elasticity of substitution is:

$$\varepsilon_{i,t} \equiv -\frac{d \ln Y_{i,t}}{d \ln P_{i,t}}$$

From the demand curve we have:

$$\frac{1}{\varepsilon_{i,t}} = -\frac{d \ln P_{i,t}}{d \ln Y_{i,t}} \quad (36)$$

$$= -\frac{dP_{i,t}}{dY_{i,t}} \frac{Y_{i,t}}{P_{i,t}} \quad (37)$$

$$= -\tilde{\lambda}_t \omega_{i,t} u''(Y_{i,t}) \frac{Y_{i,t}}{P_{i,t}} \quad (38)$$

$$= -\frac{u''(Y_{i,t})Y_{i,t}}{u'(Y_{i,t})} \quad (39)$$

Which apparently collapse to $\varepsilon_{i,t} = \varepsilon$ (constant) under CES. It is also important to point out that this ES is different from those using standard Kimball preference, where $\varepsilon_{i,t}$ is not constant but a function of **relative demand** (e.g. $\frac{c_{i,t}}{C_t}$) instead of level of individual demand here. Now take the F.O.C. of optimal pricing problem and firm's optimal *flexible* pricing is given by:

$$\begin{aligned} P_{i,t} &= MC_{i,t} - \frac{Y_{i,t}}{\frac{dY_{i,t}}{dP_{i,t}}} \\ &= MC_{i,t} - \frac{Y_{i,t}}{P_{i,t}} \frac{dP_{i,t}}{dY_{i,t}} P_{i,t} \\ &= MC_{i,t} + \frac{1}{\varepsilon_{i,t}(Y_{i,t})} P_{i,t} \end{aligned}$$

This then leads to **optimal flexible price**:

$$P_{i,t}^0 = \frac{\varepsilon(Y_{i,t})}{\varepsilon(Y_{i,t}) - 1} MC_{i,t} = \underbrace{\left(\frac{\varepsilon(Y_{i,t})}{\varepsilon(Y_{i,t}) - 1} \right)}_{\equiv \mu(Y_{i,t})} \frac{1}{A_t} \left(\frac{W_{i,t}}{\alpha} \right)^\alpha \left(\frac{Q_{i,t}}{1 - \alpha} \right)^{1-\alpha} \quad (40)$$

where $\mu(Y_{i,t})$ is the markup that changes according to individual firm's demand. In CES case, $\mu = \frac{\epsilon}{\epsilon-1}$ is constant because elasticity of substitution is constant. Linearizing (40) we get (7).

Firm's decision rules: From firms perspective, the decision rules are choices of $\{P_{it}, Y_{it}, L_{it}, M_{it}\}$ as functions of observed exogenous terms: $\{P_t, W_{i,t}, \tilde{\lambda}_t, \omega_{i,t}\}$:

- The pricing and production decisions are:

$$P_{i,t} = \frac{\varepsilon(Y_{i,t})}{\varepsilon(Y_{i,t}) - 1} MC_{i,t} = \underbrace{\left(\frac{\varepsilon(Y_{i,t})}{\varepsilon(Y_{i,t}) - 1} \right)}_{\equiv \mu(Y_{i,t})} \frac{1}{A_t} \left(\frac{W_{i,t}}{\alpha} \right)^\alpha \left(\frac{Q_{i,t}}{1 - \alpha} \right)^{1-\alpha} \quad (\text{Pricing decision})$$

$$P_{i,t} = \tilde{\lambda}_t \omega_{i,t} u'(Y_{i,t}) \quad (\text{Production/Demand})$$

That is, the monopolistic firms solve a “dual problem” of jointly choosing $P_{i,t}$ and $Y_{i,t}$, taking into consideration their demand gives them market power.

- **How does flexible price depend on demand $Y_{i,t}$ (or $c_{i,t}$)?** It is pretty clear from (Pricing decision) that we need $\mu'(Y_{i,t}) > 0$ or equivalently $\varepsilon'(Y_{i,t}) < 0$ for optimal $P_{i,t}^0$ to be increasing in $Y_{i,t}$. In CES, demand or any aggregate demand side factors will not put pressures on desired flexible prices because markup is constant. In Kimball, the two conditions are a bit different from here (all written in relative demands due to homotheticity) but will yield similar results – aggregate demand has no pricing pressure. In all these cases, demand curve will shift in response to demand factors anyway. We can show this more clearly in an analytical example later.
- **How demand factors d_t and $\omega_{i,t}$ affect flexible price?** Let's consider a simpler problem w.l.o.g. As the problem above is actually static and the F.O.C for labor will only determine labor supply after solving $\tilde{\lambda}_t$ and $c_{i,t}$, let's consider for any given N_t we can write the system of equations in (30) as:

$$c_{i,t} = v \left(\frac{P_{i,t}}{\tilde{\lambda}_t \omega_{i,t}} \right)$$

$$H_t = \int_0^1 P_{i,t} c_{i,t} di$$

where $v(\cdot) = u'^{-1}(\cdot)$ is the demand function and $H_t \equiv W_t N_t + e^{d_t} + \Pi_t$ so an aggregate demand shock can come in form of aggregate income shock $d_t > 0$ that will increase H_t . Note that $v'(x) = 1/u''(v(x)) < 0$. Now plug the demand function into the budget

constraint, we have implicit function that defines $\tilde{\lambda}_t(H_t, \{\omega_{i,t}\}_{i=0}^1)$:¹⁵

$$\int_0^1 P_{i,t} v \left(\frac{P_{i,t}}{\tilde{\lambda}_t \omega_{i,t}} \right) di = H_t$$

We show:

1. $\frac{\partial \tilde{\lambda}_t}{\partial H_t} > 0$: Take derivative of the implicit function of $\tilde{\lambda}_t$ above:

$$\frac{\partial \tilde{\lambda}_t}{\partial H_t} \int_0^1 P_{i,t} v' \left(\frac{P_{i,t}}{\tilde{\lambda}_t \omega_{i,t}} \right) \left(-\frac{P_{i,t}}{\tilde{\lambda}_t^2 \omega_{i,t}} \right) di = 1$$

This leads immediately to:

$$\frac{\partial \tilde{\lambda}_t}{\partial H_t} > 0 \quad \text{and} \quad \frac{\partial c_{i,t}}{\partial H_t} = v' \left(\frac{P_{i,t}}{\tilde{\lambda}_t \omega_{i,t}} \right) \left(-\frac{P_{i,t}}{\tilde{\lambda}_t^2 \omega_{i,t}} \right) \frac{\partial \tilde{\lambda}_t}{\partial H_t} > 0$$

because $v'(\cdot) < 0$.

2. $\frac{\partial c_{i,t}}{\partial \omega_{i,t}} > 0$: When thinking about the effect of change for single $\omega_{i,t}$ in a continuum of goods $i \in [0, 1]$, recognize that each variety is measure zero, changing the wait on one measure zero variety has no impact on total expenditure – $\frac{\partial \tilde{\lambda}_t}{\partial \omega_{i,t}} = 0$ for any i .¹⁶¹⁷ This means that:

$$\frac{\partial c_{i,t}}{\partial \omega_{i,t}} = -v'(\cdot) \frac{P_{i,t}}{\tilde{\lambda}_t \omega_{i,t}^2} > 0$$

B.3 Non-homothetic Preference

A non-homothetic demand system implies that changes in aggregate demand conditions can change the elasticity of demand faced by firms and therefore desired markups.¹⁸ This breaks

¹⁵It is also a function of $P_{i,t}$ but as we are examining shift of demand curves, it is partial movements fixing $P_{i,t}$ so we drop it to save notations.

¹⁶It might be useful to think of finite variety version with:

$$\begin{aligned} \max_{c_k} \quad & \ln C_t \\ \text{s.t.} \quad & \frac{1}{N} \sum_k p_k c_k = H \\ & u(C_t) = \frac{1}{N} \sum_{k=1}^N w_k u(c_k) \end{aligned}$$

Then take the limit $N \rightarrow \infty$

¹⁷Note that in this continuum setup, we implicitly eliminated the possibility that idiosyncratic demand shock for one variety serves like negative demand shock for the other. Because $\frac{\partial \tilde{\lambda}_t}{\partial \omega_{j,t}} = 0$.

¹⁸See Bertoletti and Etro (2017); Cavallari and Etro (2020); Matsuyama (2023), etc.

the CES benchmark in which markups are constant and aggregate demand shifts do not affect desired prices absent marginal-cost movements. Moreover, with Kimball-style homothetic preference, aggregate demand factor shifts demand curve but only proportionally. As a result, optimal pricing will not respond to aggregate demand factors as well. In this Appendix we show how changes to aggregate income d_t , idiosyncratic taste $\omega_{i,t}$ and cost shock (through $MC_{i,t}$) change the desired flexible price of individual firm, under (1) our non-homothetic preference; (2) CES; (3) Kimball.

Non-homothetic and CES Our non-homothetic framework nests CES as a special case, so it's easy to put them together. Follow [Cavallari and Etro \(2020\)](#), we consider:

$$u(c) = \gamma c + \frac{\varepsilon}{\varepsilon - 1} c^{\frac{\varepsilon - 1}{\varepsilon}} \quad (41)$$

when $\gamma = 0$ this collapse to CES. When $\gamma \neq 0$, this is a non-homothetic preference. The firm's optimal pricing and production $\{P_{i,t}, Y_{i,t}\}$ are given by ([Pricing decision](#)) and ([Production/Demand](#)):

$$\begin{aligned} P_{i,t} &= \frac{\varepsilon(Y_{i,t})}{\varepsilon(Y_{i,t}) - 1} MC_{i,t} = \mu(Y_{i,t}) \frac{1}{A_t} \left(\frac{W_t}{\alpha} \right)^\alpha \left(\frac{Q_t}{1 - \alpha} \right)^{1 - \alpha} \\ P_{i,t} &= \tilde{\lambda}_t \omega_{i,t} u'(Y_{i,t}) \end{aligned}$$

with markup:

$$\mu(Y_{i,t}) = \left(\frac{\varepsilon(Y_{i,t})}{\varepsilon(Y_{i,t}) - 1} \right), \quad \varepsilon(Y_{i,t}) = - \frac{u'(Y_{i,t})}{u''(Y_{i,t}) Y_{i,t}}$$

Note that for aggregate income shock, we showed that $\tilde{\lambda}_t$ increases in H_t , so that a positive aggregate income shock is like $\tilde{\lambda}_t$ increases. From the demand curve, it shifts outward, then for price to respond, we need the pricing curve to be upward sloping. That is $\mu'(Y_{i,t}) > 0$ and $\varepsilon'(Y_{i,t}) < 0$. Given the analytical form, this translates into $\gamma < 0$. The following figure is when $\varepsilon = 2$, $\gamma = -0.2$, $\omega_{i,t} = 0.5$ and we consider exogenous increases of (1) $\tilde{\lambda}_t$ (aggregate demand/income); (2) $MC_{i,t}$ (marginal cost); (3) $\omega_{i,t}$ (idiosyncratic demand).




Figure A.5: Non-homothetic preference

Then we do the same exercise with CES where $\gamma = 0$:




Figure A.6: CES preference

It is clear that given the shift on demand curve of both aggregate and idiosyncratic demand factors are (almost) identical across these two preferences, whether flexible price responds depends on how markup responds to demand. In CES, markup is constant so shifts of demand only affect production but not desired prices.

Kimball Preference We now turn to the case for Kimball preference, consider the simple static problem (but yield essentially the same intra-period solutions):

$$\max_{C_t, c_{i,t}} \ln C_t \quad (42)$$

$$s.t. \quad \int_0^1 P_{i,t} c_{i,t} di = H_t \quad (43)$$

$$\int_0^1 \omega_{i,t} \underbrace{\Gamma\left(\frac{c_{i,t}}{C_t}\right)}_{\equiv s_{i,t}} di = 1 \quad (44)$$

where (44) is the general Kimball aggregator. Attach λ_t to B.C. and ϕ_t to the Kimball aggregator, we get the demand curve from F.O.C of $c_{i,t}$:

$$P_{i,t} = \frac{\phi_t}{C_t \lambda_t} \omega_{i,t} \Gamma'(s_{i,t}) \quad (45)$$

We can again define $\tilde{\lambda}_t = \frac{\phi_t}{C_t \lambda_t}$ so we get similar form of demand function like before:

$$P_{i,t} = \tilde{\lambda}_t \omega_{i,t} \Gamma'(s_{i,t})$$

The major difference is that $P_{i,t}$ is a function of **relative demand** $s_{i,t}$ instead of level of demand $c_{i,t}$ like in the general setup before. Now again we need to consider how H_t shifts the above demand curve. To do that, we need to know how H_t affects $\tilde{\lambda}_t$. Note that the HH maximization problem gives solution for $\{s_{i,t}, C_t, \tilde{\lambda}_t\}$ (the separation of λ_t and ϕ_t is irrelevant for our purpose). These endogenous variables are given by the system of equations:

$$\begin{cases} P_{i,t} &= \tilde{\lambda}_t \omega_{i,t} \Gamma'(s_{i,t}) \\ \int_0^1 \Gamma(s_{i,t}) di &= 1 \\ \int_0^1 P_{i,t} s_{i,t} di &= H_t / C_t \end{cases} \quad (46)$$

The first equation defines (relative) demand function of $s_{i,t}$:

$$s_{i,t}(\tilde{\lambda}_t, \omega_{i,t}) = v\left(\frac{P_{i,t}}{\tilde{\lambda}_t \omega_{i,t}}\right) \quad (47)$$

The second and third equations implicitly define $\tilde{\lambda}_t$ and C_t functions of $\{p, w, H_t\}$. Now take derivatives of these two equations w.r.t. H_t :

$$\begin{cases} -\frac{\partial \tilde{\lambda}_t}{\partial H_t} \left(\int_0^1 \Gamma'(s_{i,t}) v' \left(\frac{P_{i,t}}{\tilde{\lambda}_t \omega_{i,t}} \right) \frac{P_{i,t}}{\tilde{\lambda}_t^2 \omega_{i,t}} di \right) &= 0 \\ -\frac{\partial \tilde{\lambda}_t}{\partial H_t} \left(\int_0^1 P_{i,t} v' \left(\frac{P_{i,t}}{\tilde{\lambda}_t \omega_{i,t}} \right) \frac{P_{i,t}}{\tilde{\lambda}_t^2 \omega_{i,t}} di \right) &= 1/C_t - H_t/C_t^2 \frac{\partial C_t}{\partial H_t} \end{cases} \quad (48)$$

These give you that:

$$\frac{\partial \tilde{\lambda}_t}{\partial H_t} = 0, \quad \frac{\partial C_t}{\partial H_t} = \frac{C_t}{H_t}$$

Basically: increase of income H_t will not shift the relative demand curve $s_{i,t} = v(P_{i,t}/\tilde{\lambda}_t)$, but it indeed shifts the level demand curve: $c_{i,t} = C_t v(P_{i,t}/\tilde{\lambda}_t)$ by shifting C_t one by one. Intuitively, because of the homothetic preference like in Kimball, increase of income leads to increase in aggregate consumption, which scales up demand for all varieties.

Now the individual firm's problem is the same and gives:

$$P_{i,t} = \frac{\varepsilon(s_{i,t})}{\varepsilon(s_{i,t}) - 1} MC_{i,t} \quad (49)$$

But note that

$$\varepsilon_{i,t} = -\frac{\Gamma'(s_{i,t})}{\Gamma''(s_{i,t}) s_{i,t}}$$

That is, the elasticity of substitution is conveniently defined as $\varepsilon_{i,t} = -\frac{\partial \ln \frac{c_{i,t}}{C_t}}{\partial \ln \frac{P_{i,t}}{P_t}}$ where $P_t \equiv \tilde{\lambda}_t$, as a result, it depends on the **relative demand** $s_{i,t}$. Again, note that if $\Gamma(s) = s^{\frac{\varepsilon-1}{\varepsilon}}$ it collapses into CES.

Now it is straightforward that the aggregate income shock H_t won't affect the desired price as: for any given solution $s_{i,t}^*, P_{i,t}^*$, increase of H_t will only shift C_t one-for-one, which yields solution that $c_{i,t} = C_t s_{i,t}^*$ with the same $P_{i,t}^*$ – production for all varieties scales up, but desired price doesn't change. Consider the analytical example from Dorsey and King (2005):

$$\Gamma(s) = \frac{1}{(1+\eta)\gamma} [(1+\eta)s - \eta]^\gamma - [1 + \frac{1}{(1+\eta)\gamma}]$$

Where $\gamma > 0$ (essentially it's roughly $\gamma = \frac{\varepsilon-1}{\varepsilon}$ as in CES). Take parameters $\gamma = 1.02$, $\eta = -6$ (this is from their baseline Kimball, if $\eta = 0$ we back to CES).




Figure A.7: Kimball preference

It is quite clear that (1) aggregate price shifts both pricing and demand curve but proportionally, as a result, only production increase but not desired prices; (2) marginal cost and idiosyncratic demand will both increase desired prices.

B.4 Quadratic approximation and the (S,s) policy

Following [Auclert et al. \(2024\)](#) and the canonical menu-cost literature [Sheshinski and Weiss \(1977\)](#); [Caplin and Spulber \(1987\)](#), we approximate the static profit function around the flexible optimum and show that the state variable is the price gap. Under stationarity and symmetry, the optimal policy admits a constant, approximately symmetric inaction region.

We will do quadratic approximation of static profit function (net of adjustment costs) so we will drop t for now. From ([Production/Demand](#)) equation we write the implied demand function:

$$y = v \left(\frac{P_i}{\tilde{\lambda} w_i} \right) \quad (50)$$

with $v(\cdot) \equiv u'^{-1}(\cdot)$. We can write the per-period profit function for individual firm i , net of adjustment cost, as:

$$\pi(P_i; S_i) = P_i v \left(\frac{P_i}{\tilde{\lambda} w_i} \right) - MC_i v \left(\frac{P_i}{\tilde{\lambda} w_i} \right) \quad (51)$$

where $S_i \equiv \{A_i, W_i, Q_i, \tilde{\lambda}, w_i\}$, note that these are collections of all our “factor” variables that affect pricing decision, and we allow idiosyncratic factors of productivity, wage and input price, which are captured by MC_i .

Quadratic Approximation of profit function Recall the flexible optimal price as P_i^0 ,¹⁹ define the **log** of it as $p_i^0 \equiv \ln P_i^0$ and the log price gap:

$$g_i \equiv p_i - p_i^0 \quad (52)$$

We can write the profit function in terms of g_i , and do Taylor expansion at $g_i = 0$:

$$\begin{aligned} \pi(P_i; S_i) &\equiv \pi(\exp(p_i^0 + g_i); S_i) = \pi(\exp(p_i^0); S_i) + \pi_P(\exp(p_i^0); S_i) \exp(p_i^0) g_i + \frac{1}{2} \pi_{PP}(\exp(p_i^0); S_i) (P_i^0)^2 g_i^2 + o(g_i^2) \\ &= \pi(P_i^0; S_i) + \frac{1}{2} \pi_{PP}(P_i^0; S_i) (P_i^0)^2 g_i^2 + o(g_i^2) \end{aligned} \quad (53)$$

where the second line follows from the fact $\pi_P(P_i^0; S_i) = 0$. Now define $\theta(S_i) \equiv -\pi_{PP}(P_i^0; S_i) (P_i^0)^2$, i.e. the second order derivative of profit function evaluated at flexible optimal price, which makes it a function of state S_i .²⁰ We get the profit loss function:

$$L(g_i; S_i) \equiv \pi(P_i^0; S_i) - \pi(P_i; S_i) \approx \frac{1}{2} \theta(S_i) g_i^2 \quad (54)$$

As a result, the per-period profit maximization is approximately (upto second order) equivalent to loss minimization that depends on g_i , i.e the distance between log of price and the static flexible price.

Linearized flexible optimal price Take log of the pricing equation ([Pricing decision](#)) and define $x_{i,t} \equiv \ln X_{i,t}$, and $m(Y) = \ln \mu(Y)$, the log flexible price is given by:

$$p_{i,t}^0 = m(Y_{i,t}) - a_{i,t} + \alpha w_{i,t} + (1 - \alpha) q_{i,t} + \text{cons} \quad (55)$$

where $\mu(Y)$ is the markup that is implied by demand function (50). We then linearize the demand function and $m(Y)$:

$$Y_{i,t} - \bar{Y} \approx v'(\bar{X}) \bar{X} (p_{i,t}^0 - \bar{p}^0) - v'(\bar{X}) \bar{X} (\ln \tilde{\lambda}_t - \ln \bar{\lambda}) - v(\bar{X}) \bar{X} (\ln \omega_{i,t} - \ln \bar{\omega}) \quad (56)$$

$$= \psi_p (p_{i,t}^0 - \bar{p}^0) + \psi_\lambda (\ln \tilde{\lambda}_t - \ln \bar{\lambda}) + \psi_\omega (\ln \omega_{i,t} - \ln \bar{\omega}) \quad (57)$$

where $X \equiv \frac{P^0}{\lambda \omega}$ for shorthand notations. And $\psi_p < 0$ and $\psi_\lambda, \psi_\omega > 0$ because $v'(\cdot) < 0$. Consistent with our later estimation procedure, we assume for each linearized shock $x \in \{a, q, w, \hat{\lambda}\}$, it is a linear combination of an aggregate component that follows AR(1) and an idiosyncratic shock ($e^a, e^q, e^w, \hat{\omega}$) that is i.i.d normal with mean zeros. To save notations, we denoted $\hat{\lambda} \equiv \ln \tilde{\lambda}$

¹⁹The solution is given by equations ([Pricing decision](#)) and ([Production/Demand](#)), note the flexible price is a static problem, so p^0 only depends on current factors s_i .

²⁰Note that profit function is strictly concave to guarantee unique optimal price so $\pi_{PP} < 0$, so $\theta(S_i) > 0$

and $\hat{\omega} \equiv \ln \omega$:

$$\begin{aligned} x_{i,t} &= x_t + e_{i,t}^x, & x_t &= \rho_x x_{t-1} + \epsilon_t^x, & x &\in \{a, w, q, \hat{\lambda}\} \\ e^x &\sim N(0, \sigma_x^2), & \epsilon_t^x &\sim N(0, \sigma_{\epsilon,x}^2), & \hat{\omega}_{i,t} &\sim N(0, \sigma_\omega^2) \end{aligned}$$

Now we linearize $m(Y_{i,t})$:

$$m(Y_{i,t}) = m(\bar{Y}) + m_Y \left[\psi_p(p_{i,t}^0 - \bar{p}^0) + \psi_\lambda \hat{\lambda}_t + \psi_\omega \hat{\omega}_{i,t} \right] \quad (58)$$

where²¹

$$m_Y \equiv \frac{\bar{Y}}{\mu(\bar{Y})} \mu'(\bar{Y})$$

Finally we reach the following linearized flexible price rule that can be simplified into (8):

$$p_{i,t}^0 = \frac{m_Y \psi_\lambda}{\Theta} \hat{\lambda}_t + \frac{m_Y \psi_\omega}{\Theta} \hat{\omega}_{i,t} - \frac{1}{\Theta} a_{i,t} + \frac{\alpha}{\Theta} w_{i,t} + \frac{1-\alpha}{\Theta} q_{i,t} \quad (59)$$

with $\Theta \equiv 1 - m_Y \psi_p$.

B.5 Optimal pricing with menu-cost

Now given the per-period profit loss function, the menu cost ξ , the individual firm solves dynamic loss minimization problem.²²

$$\min_{\{p_{i,t}\}} E_0 \sum_{t=0}^{\infty} \left[\frac{1}{2} (p_{i,t} - p_{i,t}^0)^2 + \xi \mathbb{1}(p_{i,t} \neq p_{i,t-1}) \right] \quad (60)$$

where $p_{i,t}^0$ follows from equation derived before:

$$p_{i,t}^0 = \beta_\lambda \hat{\lambda}_t + \beta_\omega \hat{\omega}_{i,t} - \beta_a a_{i,t} + \beta_w w_{i,t} + \beta_q q_{i,t} \quad (61)$$

We can now separate the aggregate factors with mean-zero, idiosyncratic factors as in [Auclert et al. \(2024\)](#) and define $p_{i,t}^*$ as the static optimal price from idiosyncratic factors, Φ_t as the

²¹Note we required $\mu'(\bar{Y}) > 0$ so that market power increases in demand, so $m_Y > 0$.

²²Note that to get (60) is complicated. The key difference here with the static loss is that $\theta(S_{i,t})$ is dropped, also in writing the explicit non-linear profit maximizing problem it involves a stochastic discounting factor Λ_t :

$$\min_{p_{i,t}} E_0 \sum_{t=0}^{\infty} \Lambda_t \left[\frac{1}{2} \theta(S_{i,t}) (p_{i,t} - p_{i,t}^0)^2 + \sigma^2 \tilde{\xi} \mathbb{1}(p_{i,t} \neq p_{i,t-1}) \right]$$

where σ^2 is the variance of all idiosyncratic shocks, we follow [Auclert et al. \(2024\)](#) to normalize the cost by this constant factor to guarantee its limit when $\sigma^2 \rightarrow 0$ the profit function is well behaved. Following Auclert et al. (2024, App. E.1, E.3, E.4), under small idiosyncratic shocks and small aggregate fluctuations, our micro-founded profit function with menu costs is first-order equivalent to the canonical quadratic problem (60), up to a constant scaling of prices and the menu cost.

aggregate factors:

$$p_{i,t}^* = \beta_\omega \hat{\omega}_{i,t} - \beta_a e_{i,t}^a + \beta_w e_{i,t}^w + \beta_q e_{i,t}^q \quad (62)$$

$$\Phi_t = \beta_\lambda \hat{\lambda}_t - \beta_a a_t + \beta_w w_t + \beta_q q_t \quad (63)$$

So that $p_{i,t}^0 \equiv p_{i,t}^* + \Phi_t$. Define the price gap as $g_{i,t} = p_{i,t} - p_{i,t}^0$, we can then rewrite (60) as:²³

$$\min_{\{g_{i,t}\}} E_0 \sum_{t=0}^{\infty} \left[\frac{1}{2} g_{i,t}^2 + \xi \mathbb{1}(g_{i,t} \neq g_{i,t}^N) \right] \quad (65)$$

where $g_{i,t}^N$ denotes the price gap at time t **without adjusting**:

$$g_{i,t}^N = p_{i,t-1} - p_{i,t}^0 = g_{i,t-1} - \Delta \Phi_t - \Delta p_{i,t}^* \quad (66)$$

Following the standard literature on the canonical menu-cost model, as Φ_t is stationary, $\Delta p_{i,t}^*$ is mean zero and symmetrically distributed, the problem in (64) gives an s-S band policy that takes the following form:

$$g_{i,t} = \begin{cases} g_t^*(\Phi_t) & \text{if } g_{i,t}^N \notin [g_t^L(\Phi_t), g_t^U(\Phi_t)] \\ g_{i,t}^N & \text{otherwise} \end{cases} \quad (67)$$

Note that given the definition of the gap: $g_{i,t} = p_{i,t} - p_{i,t}^0$, when desired price $p_{i,t}^0$ goes up, the gap is negative, the price without adjustment is “too low”, so that the $g_{i,t}^N$ is more likely to cross the $g_t^L(\Phi_t)$ to be lower. When $g_{i,t}^N$ crosses, the reset price is increasing as $p_{i,t}^0$ is higher. Now notice that in general, the sS bands and the optimal reset gap depends on aggregate states Φ_t but not idiosyncratic factors $p_{i,t}^*$ as shown in Auclert et al. (2024). However, idiosyncratic shocks do move the likelihood for individual firm to adjust by shifting $g_{i,t}^N$ and the reset price for individual firm by shifting $p_{i,t}^0$. Moreover, it is well-known in the literature (Sheshinski and Weiss (1977); Caplin and Spulber (1987); Auclert et al. (2024) etc) that around the steady state with $\Delta \Phi = 0$ and $\Delta p_i^* = 0$, the band is constant and symmetric $g^U = -g^L$ ($g^L < 0$, $g^U > 0$) and optimal reset gap is zero: $g^* = 0$.

For purpose of our analysis, we focus on the small aggregate fluctuations around the steady state, so we follow Auclert et al. (2024) to use fixed bands and $g^* = 0$. One can think of this

²³Or as in Auclert et al. (2024) set up, define $x_{i,t} \equiv p_{i,t} - p_{i,t}^*$ we write:

$$\min_{\{x_{i,t}\}} E_0 \sum_{t=0}^{\infty} \left[\frac{1}{2} (x_{i,t} - \Phi_t)^2 + \xi \mathbb{1}(x_{i,t} \neq x_{i,t-1} - \Delta p_{i,t}^*) \right] \quad (64)$$

exercise as first order approximation of $g^*(\Phi)$, $g^U(\Phi)$ and $g^L(\Phi)$ functions around steady state level of $\Phi = 0$. This then gives our optimal price policy of individual firm:²⁴

$$p_{i,t} = \begin{cases} p_{i,t}^0 = \Phi_t + p_{i,t}^* & \text{if } p_{i,t-1} \notin [g^L + \Phi_t + p_{i,t}^*, g^U + \Phi_t + p_{i,t}^*] \\ p_{i,t-1} & \text{if } p_{i,t-1} \in [g^L + \Phi_t + p_{i,t}^*, g^U + \Phi_t + p_{i,t}^*] \end{cases} \quad (69)$$

C Estimation model derivations

This appendix collects the probability formulas, likelihood blocks, and weighting formulas used in the estimation procedure.

C.1 Phrase probabilities

Under (15) and (16), the noisy phrase index satisfies

$$s_{i,m,t}^R \mid z_{m,t} \sim \mathcal{N}(z_{m,t}, \sigma_{m,R}^2), \quad \sigma_{m,R}^2 \equiv \sigma_{m,v}^2 + \sigma_{m,\xi}^2. \quad (70)$$

Therefore,

$$p_{m,t}^{R,+}(z_{m,t}) \equiv \mathbb{P}(R_{i,m,t} = + \mid z_{m,t}) = \Phi\left(\frac{z_{m,t} - d_m}{\sigma_{m,R}}\right), \quad (71)$$

$$p_{m,t}^{R,-}(z_{m,t}) \equiv \mathbb{P}(R_{i,m,t} = - \mid z_{m,t}) = \Phi\left(\frac{-z_{m,t} - d_m}{\sigma_{m,R}}\right), \quad (72)$$

$$p_{m,t}^{R,0}(z_{m,t}) = 1 - p_{m,t}^{R,+}(z_{m,t}) - p_{m,t}^{R,-}(z_{m,t}). \quad (73)$$

²⁴Alternatively, one can allow $g^*(\Phi)$, $g^L(\Phi)$ and $g^U(\Phi)$ to depend on Φ up to first order by doing Taylor expansion around steady state Φ (we can do this as the Ss policy solves a smooth Bellman Equation with respect to the aggregate state Φ , which is continuously differentiable in Φ):

$$g^*(\Phi) \approx a\Phi, \quad g^L(\Phi) \approx g^L + L\Phi, \quad g^U(\Phi) \approx g^U + U\Phi$$

The resulted optimal reset price is then:

$$p_{i,t} = \begin{cases} p_{i,t}^0 + a\Phi_t = (1 + a)\Phi_t + p_{i,t}^* & \text{if } p_{i,t-1} \notin [g^L + (1 + L)\Phi_t + p_{i,t}^*, g^U + (1 + H)\Phi_t + p_{i,t}^*] \\ p_{i,t-1} & \text{if } p_{i,t-1} \in [g^L + (1 + L)\Phi_t + p_{i,t}^*, g^U + (1 + H)\Phi_t + p_{i,t}^*] \end{cases} \quad (68)$$

That is, the cutoff value and reset prices still depend linearly on Φ_t and $p_{i,t}^*$. In general, it is reasonable to think of price setters “overshoot” in proportion to Φ_t when Φ_t is persistent, which explains why $g^*(\Phi)$ depends on Φ .

Conditional on $z_{m,t}$, the observed phrase counts satisfy

$$(N_{+,m,t}, N_{0,m,t}, N_{-,m,t}) \mid z_{m,t} \sim \text{Multinomial}\left(N_{m,t}; p_{m,t}^{R,+}(z_{m,t}), p_{m,t}^{R,0}(z_{m,t}), p_{m,t}^{R,-}(z_{m,t})\right), \quad (74)$$

where $N_{m,t} = N_{+,m,t} + N_{0,m,t} + N_{-,m,t}$.

C.2 Logic probabilities within phrase branches

Because phrase selection is determined by $s_{i,m,t}^R = z_{m,t} + v_{i,m,t} + \xi_{i,m,t}$ while logic status is determined by $s_{i,m,t}^L = z_{m,t} + v_{i,m,t}$, the conditional logic probabilities are bivariate-normal selection probabilities.

Define

$$\rho_m \equiv \frac{\sigma_{m,v}}{\sigma_{m,R}}, \quad \sigma_{m,R}^2 = \sigma_{m,v}^2 + \sigma_{m,\xi}^2. \quad (75)$$

Positive phrase branch. Let

$$a_{m,t}^+ \equiv \frac{\kappa_{m,+} - z_{m,t}}{\sigma_{m,v}}, \quad b_{m,t}^+ \equiv \frac{d_m - z_{m,t}}{\sigma_{m,R}}. \quad (76)$$

Then the conditional active-logic probability in the positive branch is

$$q_{A|R+,m}(z_{m,t}) \equiv \mathbb{P}(A \mid R = +, z_{m,t}) = \frac{1 - \Phi_2(a_{m,t}^+, b_{m,t}^+; \rho_m)}{1 - \Phi(b_{m,t}^+)}, \quad (77)$$

where $\Phi_2(\cdot, \cdot; \rho)$ denotes the standard bivariate normal cdf with correlation ρ .

Negative phrase branch. Let

$$a_{m,t}^- \equiv \frac{-\kappa_{m,-} - z_{m,t}}{\sigma_{m,v}}, \quad b_{m,t}^- \equiv \frac{-d_m - z_{m,t}}{\sigma_{m,R}}. \quad (78)$$

Then the conditional active-logic probability in the negative branch is

$$q_{A|R-,m}(z_{m,t}) \equiv \mathbb{P}(A \mid R = -, z_{m,t}) = \frac{\Phi_2(a_{m,t}^-, b_{m,t}^-; \rho_m)}{\Phi(b_{m,t}^-)}. \quad (79)$$

Hence the observed logic counts satisfy

$$M_{A,+,m,t} \mid N_{+,m,t}, z_{m,t} \sim \text{Binomial}(N_{+,m,t}, q_{A|R+,m}(z_{m,t})) , \quad (80)$$

$$M_{A,-,m,t} \mid N_{-,m,t}, z_{m,t} \sim \text{Binomial}(N_{-,m,t}, q_{A|R-,m}(z_{m,t})) . \quad (81)$$

C.3 Likelihood blocks

Let θ collect all model parameters. Let $X_t = [z_{d,t}, z_{w,t}, z_{c,t}]'$ denote the narrative-factor state vector. In the inflation block, the residual component u_t follows (20). The full state vector in the final stage is therefore $(X'_t, u_t)'$.

State transition block. The latent factor process contributes

$$\ell_X(\theta) = \sum_{m \in \{d,w,c\}} \sum_{t=1}^T \log f(z_{m,t} \mid z_{m,t-1}; \rho_m) , \quad (82)$$

with $z_{m,t} = \rho_m z_{m,t-1} + \eta_{m,t}$ and $\eta_{m,t} \sim \mathcal{N}(0, 1)$. When the inflation residual is included, its state block is

$$\ell_u(\theta) = \sum_{t=1}^T \log f(u_t \mid u_{t-1}; \rho_u, \sigma_u^2) , \quad (83)$$

with $u_t = \rho_u u_{t-1} + \eta_{u,t}$ and $\eta_{u,t} \sim \mathcal{N}(0, \sigma_u^2)$.

Optional Gaussian signal block. If Y_t is included,

$$\ell_Y(\theta) = -\frac{1}{2} \sum_{t=1}^T \left[\log |R_Y| + (Y_t - CX_t)' R_Y^{-1} (Y_t - CX_t) \right] + \text{const.} \quad (84)$$

Phrase gate block. Using (74), the phrase log likelihood is

$$\ell_R(\theta) = \sum_{m \in \{d,w,c\}} \sum_{t=1}^T \left[N_{+,m,t} \log p_{m,t}^{R,+}(z_{m,t}) + N_{0,m,t} \log p_{m,t}^{R,0}(z_{m,t}) + N_{-,m,t} \log p_{m,t}^{R,-}(z_{m,t}) \right] + \text{const.}, \quad (85)$$

and we denote the factor-specific contribution by

$$\ell_{R,m}(\theta) = \sum_{t=1}^T \left[N_{+,m,t} \log p_{m,t}^{R,+}(z_{m,t}) + N_{0,m,t} \log p_{m,t}^{R,0}(z_{m,t}) + N_{-,m,t} \log p_{m,t}^{R,-}(z_{m,t}) \right]. \quad (86)$$

Logic block. Using (80)–(81), the conditional logic log likelihood is

$$\begin{aligned} \ell_L(\theta) = & \sum_{m \in \{d,w,c\}} \sum_{t=1}^T \left\{ M_{A,+,m,t} \log q_{A|R+,m}(z_{m,t}) + (N_{+,m,t} - M_{A,+,m,t}) \log(1 - q_{A|R+,m}(z_{m,t})) \right\} \\ & + \sum_{m \in \{d,w,c\}} \sum_{t=1}^T \left\{ M_{A,-,m,t} \log q_{A|R-,m}(z_{m,t}) + (N_{-,m,t} - M_{A,-,m,t}) \log(1 - q_{A|R-,m}(z_{m,t})) \right\}. \end{aligned} \quad (87)$$

Inflation block. Conditional on (z_t, u_t) ,

$$P_t = \theta_0 + \sum_{m \in \{d,w,c\}} \psi_m [M_{m,t}^+(z_{m,t}) + M_{m,t}^-(z_{m,t})] + u_t + \varepsilon_t^P, \quad \varepsilon_t^P \sim \mathcal{N}(0, R_P). \quad (88)$$

Therefore,

$$\ell_P(\theta) = -\frac{1}{2} \sum_{t=1}^T \left[\log R_P + \frac{\left(P_t - \theta_0 - \sum_{m \in \{d,w,c\}} \psi_m [M_{m,t}^+(z_{m,t}) + M_{m,t}^-(z_{m,t})] - u_t \right)^2}{R_P} \right] + \text{const.} \quad (89)$$

C.4 Three-stage estimation objective

Stage 1: conditional likelihood using the logic block. The first-stage objective uses only the latent factor state law of motion and the logic block:

$$\ell_{cond}(\theta) \equiv \ell_X(\theta) + \ell_L(\theta). \quad (90)$$

Equivalently, the first stage conditions on the observed branch exposures $N_{+,m,t}$ and $N_{-,m,t}$ and fits the conditional likelihood of $M_{A,+,m,t}$ and $M_{A,-,m,t}$ given those exposures. The phrase multinomial block, inflation block, Gaussian signal block, and the residual inflation state u_t are excluded at this stage.

Stage 2: Dirichlet-multinomial effective sample size calibration. For each factor m , let

$$n_{m,t} \equiv (N_{+,m,t}, N_{0,m,t}, N_{-,m,t}), \quad \hat{p}_{m,t}^R \equiv (\hat{p}_{m,t}^{R,+}, \hat{p}_{m,t}^{R,0}, \hat{p}_{m,t}^{R,-}), \quad (91)$$

where $\hat{p}_{m,t}^R$ is computed from the first-stage fitted states and parameters. We then fit the auxiliary Dirichlet-multinomial calibration

$$n_{m,t} \mid \hat{p}_{m,t}^R \sim \text{Dirichlet-Multinomial}(N_{m,t}; \alpha_{m,t}), \quad \alpha_{m,t} = \beta_m \hat{p}_{m,t}^R, \quad (92)$$

with factor-specific concentration parameter $\beta_m > 0$.

The implied effective sample size at (m, t) is

$$N_{m,t}^{eff} = \frac{N_{m,t}(1 + \beta_m)}{N_{m,t} + \beta_m}, \quad (93)$$

and the corresponding gate weight is

$$\lambda_{gate,m} = \frac{\sum_{t=1}^T N_{m,t}^{eff}}{\sum_{t=1}^T N_{m,t}}. \quad (94)$$

This weight interprets the phrase block as contributing only its effective information content, rather than its raw multinomial sample size, to the final estimation criterion.

Stage 3: weighted composite likelihood. The final-stage objective is

$$\ell_{CL}(\theta) = \ell_X(\theta) + \ell_u(\theta) + \ell_Y(\theta) + \ell_L(\theta) + \ell_P(\theta) + \sum_{m \in \{d, w, c\}} \lambda_{gate,m} \ell_{R,m}(\theta). \quad (95)$$

Equivalently, one may write

$$\ell_{CL}(\theta) = \ell_{core}(\theta) + \ell_{gate}^w(\theta), \quad (96)$$

where

$$\ell_{core}(\theta) = \ell_X(\theta) + \ell_u(\theta) + \ell_Y(\theta) + \ell_L(\theta) + \ell_P(\theta), \quad (97)$$

$$\ell_{gate}^w(\theta) = \sum_{m \in \{d, w, c\}} \lambda_{gate,m} \ell_{R,m}(\theta). \quad (98)$$

The final estimator is

$$\hat{\theta}_{CL} = \arg \max_{\theta} \ell_{CL}(\theta). \quad (99)$$

Because the final criterion is a weighted composite likelihood rather than a literal full likelihood, standard Hessian-based uncertainty formulas need not be exact. In principle, inference should be based on the Godambe information; in practice, a bootstrap is often the most transparent route when the filtering and smoothing problem is computationally feasible.

C.5 From the micro pricing rule to the inflation measurement equation

This appendix derives the inflation measurement equation from the micro pricing rule and makes explicit where the approximation enters.

Under the fixed-band pricing rule in (10), firm i 's realized price change is

$$\Delta p_{i,t} \equiv p_{i,t} - p_{i,t-1} = (p_{i,t}^0 - p_{i,t-1}) \mathbf{1}\{p_{i,t-1} \notin [g^L + p_{i,t}^0, g^U + p_{i,t}^0]\}. \quad (100)$$

Define the exact desired-price gap

$$\delta_{i,t} \equiv p_{i,t}^0 - p_{i,t-1}. \quad (101)$$

Then the micro pricing rule can be written as

$$\Delta p_{i,t} = h(\delta_{i,t}), \quad (102)$$

where

$$h(\delta) \equiv \delta \mathbf{1}\{\delta > -g^L\} + \delta \mathbf{1}\{\delta < -g^U\}. \quad (103)$$

Using the linearized flexible-price rule in the micro model, the exact desired-price gap is

$$\delta_{i,t} = \beta_\lambda s_{i,d,t}^L + \beta_w s_{i,w,t}^L + \beta_q s_{i,c,t}^L + r_{i,t}, \quad (104)$$

where $r_{i,t}$ collects residual forces outside the three-factor decomposition, and the factor-specific latent indices are exactly those implied by the mapping in (12). Therefore the exact aggregate inflation map implied by the micro pricing rule is

$$P_t^* = G(z_{d,t}, z_{w,t}, z_{c,t}) \equiv \mathbb{E}_i[h(\beta_\lambda s_{i,d,t}^L + \beta_w s_{i,w,t}^L + \beta_q s_{i,c,t}^L + r_{i,t}) | z_{d,t}, z_{w,t}, z_{c,t}]. \quad (105)$$

Equation (105) is exact, but it is not additive because the reset rule is applied to the joint desired-price gap. To obtain a tractable decomposition, we use the anchored decomposition of G around the baseline $(0, 0, 0)$, also known as cut-HDMR (Kuo, Sloan, Wasilkowski, and

Woźniakowski, 2010; Li, Wang, Rosenthal, and Rabitz, 2001). This gives

$$G(z_{d,t}, z_{w,t}, z_{c,t}) = \theta_0 + \sum_{m \in \{d,w,c\}} F_m(z_{m,t}) + R^{int}(z_{d,t}, z_{w,t}, z_{c,t}), \quad (106)$$

where

$$\theta_0 \equiv G(0, 0, 0), \quad F_m(z_{m,t}) \equiv G(0, \dots, z_{m,t}, \dots, 0) - G(0, 0, 0), \quad (107)$$

and R^{int} collects the interaction terms omitted by the additive representation. The approximation used in the estimation model is therefore

$$G(z_{d,t}, z_{w,t}, z_{c,t}) \approx \theta_0 + \sum_{m \in \{d,w,c\}} F_m(z_{m,t}), \quad (108)$$

with the interaction remainder R^{int} absorbed by the residual inflation component u_t .

It remains to characterize each one-factor main effect. Fix a factor m . By definition,

$$F_m(z_{m,t}) = G(0, \dots, z_{m,t}, \dots, 0) - G(0, 0, 0). \quad (109)$$

Under this factor-alone counterfactual, all other classified factors are held at baseline, so the desired-price gap reduces to

$$\delta_{i,m,t}^{FA} = \beta_m s_{i,m,t}^L, \quad (110)$$

where, following (12), $(\beta_m, \kappa_{m,+}, \kappa_{m,-})$ corresponds to $(\beta_\lambda, \kappa_{d,+}, \kappa_{d,-})$ for demand, $(\beta_w, \kappa_{w,+}, \kappa_{w,-})$ for wage, and $(\beta_q, \kappa_{c,+}, \kappa_{c,-})$ for non-labor input cost. Applying the same reset map $h(\cdot)$ to this factor-alone gap gives

$$\Delta p_{i,m,t}^{FA} = h(\beta_m s_{i,m,t}^L). \quad (111)$$

Using the cutoff mapping in (12), this becomes

$$\Delta p_{i,m,t}^{FA} = \beta_m s_{i,m,t}^L \mathbf{1}\{s_{i,m,t}^L \geq \kappa_{m,+}\} + \beta_m s_{i,m,t}^L \mathbf{1}\{s_{i,m,t}^L \leq -\kappa_{m,-}\}. \quad (112)$$

Aggregating across firms gives the factor-alone contribution

$$P_{m,t}^{FA} \equiv \mathbb{E}_i[\Delta p_{i,m,t}^{FA} \mid z_{m,t}] = \beta_m \mathbb{E}_i[s_{i,m,t}^L \mathbf{1}\{s_{i,m,t}^L \geq \kappa_{m,+}\} \mid z_{m,t}] + \beta_m \mathbb{E}_i[s_{i,m,t}^L \mathbf{1}\{s_{i,m,t}^L \leq -\kappa_{m,-}\} \mid z_{m,t}]. \quad (113)$$

Conditional on $z_{m,t}$, the firm-level pricing index satisfies

$$s_{i,m,t}^L = z_{m,t} + v_{i,m,t}, \quad v_{i,m,t} \sim \mathcal{N}(0, \sigma_{m,v}^2), \quad (114)$$

so that $s_{i,m,t}^L \mid z_{m,t} \sim \mathcal{N}(z_{m,t}, \sigma_{m,v}^2)$. Define the upward and downward truncated moments

$$M_{m,t}^+(z_{m,t}) \equiv \mathbb{E}[s_{i,m,t}^L \mathbf{1}\{s_{i,m,t}^L \geq \kappa_{m,+}\} \mid z_{m,t}], \quad (115)$$

$$M_{m,t}^-(z_{m,t}) \equiv \mathbb{E}[s_{i,m,t}^L \mathbf{1}\{s_{i,m,t}^L \leq -\kappa_{m,-}\} \mid z_{m,t}]. \quad (116)$$

Then

$$P_{m,t}^{FA} = \beta_m M_{m,t}^+(z_{m,t}) + \beta_m M_{m,t}^-(z_{m,t}). \quad (117)$$

Equation (117) is the exact factor-alone benchmark implied by the micro pricing rule. It shows that the same coefficient multiplies the upward and downward truncated moments for a given factor, which is why the benchmark inflation loading is symmetric. Motivated by this result, the main-text inflation equation is written as

$$P_t^{\text{nar}} = \theta_0 + \sum_{m \in \{d,w,c\}} \psi_m [M_{m,t}^+(z_{m,t}) + M_{m,t}^-(z_{m,t})]. \quad (118)$$

Here ψ_m coincides with β_m up to normalization in the exact factor-alone benchmark, and in the empirical specification also absorbs the approximation error from replacing the exact joint map in (105) by the additive approximation in (108).

Closed-form expressions. Let

$$\alpha_{m,t} \equiv \frac{\kappa_{m,+} - z_{m,t}}{\sigma_{m,v}}, \quad \beta_{m,t} \equiv \frac{-\kappa_{m,-} - z_{m,t}}{\sigma_{m,v}}. \quad (119)$$

For the upward truncated moment, using the change of variables $s = z_{m,t} + \sigma_{m,v}u$ with $u \sim \mathcal{N}(0, 1)$,

$$\begin{aligned} M_{m,t}^+(z_{m,t}) &= \int_{\kappa_{m,+}}^{\infty} s \frac{1}{\sigma_{m,v}} \varphi\left(\frac{s - z_{m,t}}{\sigma_{m,v}}\right) ds \\ &= \int_{\alpha_{m,t}}^{\infty} (z_{m,t} + \sigma_{m,v}u) \varphi(u) du \\ &= z_{m,t} \int_{\alpha_{m,t}}^{\infty} \varphi(u) du + \sigma_{m,v} \int_{\alpha_{m,t}}^{\infty} u \varphi(u) du. \end{aligned} \quad (120)$$

The first integral equals $1 - \Phi(\alpha_{m,t}) = \Phi(-\alpha_{m,t}) = \Phi\left(\frac{z_{m,t} - \kappa_{m,+}}{\sigma_{m,v}}\right)$. For the second, note that $\frac{d}{du}\varphi(u) = -u\varphi(u)$, hence

$$\int_{\alpha_{m,t}}^{\infty} u \varphi(u) du = [-\varphi(u)]_{\alpha_{m,t}}^{\infty} = \varphi(\alpha_{m,t}). \quad (121)$$

Substituting into (120) yields

$$M_{m,t}^+(z_{m,t}) = z_{m,t} \Phi\left(\frac{z_{m,t} - \kappa_{m,+}}{\sigma_{m,v}}\right) + \sigma_{m,v} \varphi\left(\frac{\kappa_{m,+} - z_{m,t}}{\sigma_{m,v}}\right). \quad (122)$$

For the downward truncated moment,

$$\begin{aligned} M_{m,t}^-(z_{m,t}) &= \int_{-\infty}^{-\kappa_{m,-}} s \frac{1}{\sigma_{m,v}} \varphi\left(\frac{s - z_{m,t}}{\sigma_{m,v}}\right) ds \\ &= \int_{-\infty}^{\beta_{m,t}} (z_{m,t} + \sigma_{m,v} u) \varphi(u) du \\ &= z_{m,t} \int_{-\infty}^{\beta_{m,t}} \varphi(u) du + \sigma_{m,v} \int_{-\infty}^{\beta_{m,t}} u \varphi(u) du. \end{aligned} \quad (123)$$

The first integral equals $\Phi(\beta_{m,t}) = \Phi\left(\frac{-\kappa_{m,-} - z_{m,t}}{\sigma_{m,v}}\right)$. For the second, using again $\frac{d}{du}\varphi(u) = -u\varphi(u)$,

$$\int_{-\infty}^{\beta_{m,t}} u \varphi(u) du = [-\varphi(u)]_{-\infty}^{\beta_{m,t}} = -\varphi(\beta_{m,t}). \quad (124)$$

Substituting into (123) yields

$$M_{m,t}^-(z_{m,t}) = z_{m,t} \Phi\left(\frac{-\kappa_{m,-} - z_{m,t}}{\sigma_{m,v}}\right) - \sigma_{m,v} \varphi\left(\frac{-\kappa_{m,-} - z_{m,t}}{\sigma_{m,v}}\right). \quad (125)$$

Equations (122)–(125) are the closed-form objects used in the inflation measurement equation in Section 5.2.