Innovation and Pricing Frictions*

Xiaojie Liu† Northwestern

Sara Moreira‡ Northwestern, NBER

March 26, 2026

Abstract

We show that pricing frictions operate over much longer horizons than standard macroeconomic models assume, with important consequences for innovation and growth. Using product–firm longitudinal data, we document that incumbent nominal prices tend to remain relatively stable over the life cycle, while new products enter at sizable premia, a pattern we refer to as "price overshooting", and overshooting scales with sectoral stickiness and cost shocks. Motivated by these facts, we develop an endogenous growth model with long-run pricing frictions that rationalizes these patterns. Pricing frictions weaken innovation incentives, while price overshooting partially mitigates these losses and informs optimal R&D policy.

JEL Classification: E3, E4, E5, O4, O31, O34

Keywords: Innovation, growth, pricing friction, product life cycle, patents, competition, creative destruction.

*We would like to thank Weizhe Wang and Razvan Banescu for excellent research assistance. Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

†Email: xiaojie.liu@kellogg.northwestern.edu. Address: 2211 Campus Drive, Evanston, IL 60208.

‡Email: sara.moreira@kellogg.northwestern.edu. Address: 2211 Campus Drive, Evanston, IL 60208.

1 Introduction

How do pricing frictions shape firms' innovation decisions and, ultimately, economic growth? Standard macroeconomic models typically treat nominal rigidities as short-lived and therefore irrelevant for long-run outcomes.[1] Because firms are thought to reset prices frequently enough, such frictions are assumed not to influence forward-looking decisions such as innovation, product introduction, or creative destruction. In this paper, we revisit this view using rich product–firm longitudinal data and show that pricing frictions operate over much longer horizons, with meaningful consequences for the pricing of new products and firms' innovation incentives.

We start the paper with a set of motivation facts. First, we document that new products enter the market with substantial price premia relative to existing varieties. Using product–year scanner data for the consumer goods sector, we compute average log prices by cohort and follow their evolution over time. Each successive cohort enters at higher prices than earlier ones, and these differences are large and persistent. On average, new products are priced about 10-20 percent above the prevailing price level in their sector. This pattern is pervasive: it appears for products introduced by both incumbent and new firms, for items with high and low novelty, and across food and non-food categories. While some price premium for newer or higher-quality products is to be expected, the magnitude we observe is considerably larger than what existing estimates of quality improvements would predict (e.g., Cafarella, Ehrlich, Gao, Haltiwanger, Shapiro and Zhao, 2023; Argente, Lee and Moreira, 2024).

Second, we study how incumbent product prices evolve over the life cycle, a crucial dimension for understanding long-run pricing frictions. While standard analyses emphasize high-frequency price changes, what matters for firms' innovation decisions is the slow-moving, persistent component of pricing. Estimating age profiles of log prices while controlling for cohort and sector composition, we find that nominal prices show little systematic growth as products age, we find that nominal prices show little systematic growth over the life cycle, despite frequent high-frequency adjustments documented in prior work. Once sector-level trends are absorbed, relative prices decline steadily—by about 2–3 percent per year over the first five years. These results are robust across specifications, samples, and both national and store-level data. Although retail prices change frequently (Bils and Klenow, 2004; Klenow and Kryvtsov, 2008; Nakamura and Steinsson, 2008), these adjustments tend to be temporary and do not fully offset inflation, leaving the average prices of incumbent products relatively stable over the life cycle.

Having established that new products enter at sizable price premia and that prices of incumbent products are subject to long-run rigidity, we next examine whether these frictions help explain the pricing of new products. We present two pieces of evidence. First, in the cross section, we

[1] See for example Friedman (1968); Woodford (2003); Galí (2015).

construct two product-sector measures of price rigidity and show that sectors with more persistent incumbent-price stickiness exhibit systematically larger entry premia. This pattern suggests that when firms face greater difficulty adjusting the prices of existing varieties, they rely more heavily on new products as the margin through which to update their overall price level.

Second, we exploit variation in firms' exposure to import tariffs to provide complementary quasi-experimental evidence. Tariff changes generate plausibly exogenous cost shocks for firms that sell imported products, allowing us to compare how such shocks pass through to incumbent versus newly introduced products. We find a pronounced asymmetry: tariff exposure has no large detectable effect on the prices of incumbent products—consistent with strong life-cycle price rigidity—but significantly increases the relative prices of new products introduced in subsequent years, with especially large effects in sectors exhibiting greater rigidity. These results reinforce the interpretation that, when constrained in adjusting incumbent prices over multi-year horizons, firms pass through cost shocks primarily through the prices of newly launched products.

Motivated by these empirical facts, we build an endogenous growth model with price stickiness to examine their long-run implications. Our framework builds on the literature on endogenous technological change (Romer, 1990; Aghion and Howitt, 1992; Grossman and Helpman, 1991) and incorporates firm-level innovation dynamics in the spirit of Klette and Kortum (2004) with endogenous product varieties. We extend these canonical environments by embedding money demand on the household side, allowing the central bank to control trend inflation through the money supply. In the model, firms choose the introductory price of each newly launched product but thereafter face Calvo-style frictions: the opportunity to reset prices arrives only with a fixed probability. These pricing frictions capture a firm's inability to adjust prices in line with ongoing trend inflation, rather than the short-run response to business-cycle shocks emphasized in the New Keynesian literature.

Anticipating the gradual erosion of real prices caused by trend inflation and infrequent price adjustment, firms optimally choose an introductory markup above the static optimum. By setting a higher initial price, firms keep their effective markup closer to the static-optimal level over the expected duration of a price spell, thereby increasing average profits. We refer to this behavior as *price overshooting*. The mechanism is closely related to the role of inflation expectations in New Keynesian models: as Werning (2022) shows, forward-looking firms raise markups above the static optimum when they expect future marginal costs to rise. In our setting, similar forward-looking considerations lead firms to front-load markups at product introduction, consistent with the empirical evidence.

Innovation is tied to pricing frictions in two ways. First, pricing frictions depress product values by limiting firms' ability to maintain real markups, thereby weakening incentives to innovate. Second, the ability to reset prices for newly introduced products: price overshooting serves as a

hedge against the gradual erosion of real prices. This reset option can be strong enough that, in equilibrium, higher pricing frictions are associated with stronger innovation incentives. Absent this option to reset prices, pricing frictions would unambiguously weaken innovation.

An important implication of the model is that markup dispersion is endogenous and matters for aggregate outcomes (as in Peters, 2020). Price overshooting causes households to consume too little of products with elevated introductory markups and too much of products whose prices cannot adjust. Innovation mitigates these inefficiencies and expands the number of varieties: higher innovation rates increase product churning, concentrate products near the top of the markup distribution, and thereby reduce misallocation.

To estimate the model, we use product-level data to construct four moments: price overshooting, incumbent and entrant innovation rates, and the aggregate growth rate. These moments identify four key parameters—pricing frictions, the innovation capacities of incumbents and entrants, and the size of the quality step—and the estimated model matches the targeted moments closely despite its parsimony. Beyond the targeted moments, the model also performs well on untargeted dimensions: estimating price dynamics in the simulated data using the same empirical procedure yields life-cycle patterns that closely mirror those in the data.

The estimated Calvo parameter implies long price spells and therefore substantial life-cycle pricing frictions—far larger than the short-run adjustment frequencies documented by Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008). The key distinction is that most observed price changes are transitory and do not return prices to the inflation trend; only the persistent component of pricing matters for firms' innovation decisions, and this is what our estimates capture. The model also yields sizable gaps between reset and exit markups, generating meaningful efficiency losses due to misallocation.

To assess the importance of price overshooting for innovation and growth, we conduct a counterfactual that shuts it down by restricting firms to charge at most the static-optimal price. Without overshooting, innovation falls sharply because product values drop when firms cannot choose a higher reset price. Exit rises for the same reason, leading to substantially fewer varieties. As a result, aggregate growth and welfare decline meaningfully. For comparison, a benchmark model in the spirit of Klette and Kortum (2004) generates considerably smaller welfare losses.

These findings highlight the critical role of forward-looking price setting in mitigating the adverse effects of long-run pricing frictions. Without the ability to overshoot, firms underinvest in innovation and maintain fewer product lines, leading to markedly lower growth and welfare. Intuitively, the incentive to overshoot increases with the severity of pricing rigidities. This mechanism aligns with our tariff-based and cross-sectional evidence: firms and sectors that face greater frictions in adjusting the prices of incumbent products are precisely those that set higher introductory prices for new products.

We use the estimated model to study optimal R&D policy and find that subsidizing incumbent innovation generates substantially larger welfare gains than subsidizing entrants. Incumbents account for most product churning in retail markets, so boosting their innovation raises variety, improves efficiency, and delivers meaningful welfare improvements. Importantly, we also show that optimal policy depends crucially on the degree of pricing frictions. As frictions become more severe, the optimal incumbent subsidy rises because faster innovation introduces more new products, compresses the markup distribution, and mitigates misallocation. When frictions are mild, variety and efficiency gains diminish and the traditional growth margin dominates. Overall, the results imply that (i) incumbent R&D subsidies are far more effective than entrant subsidies, and (ii) optimal innovation policy should be stronger in sectors with greater pricing frictions to unlock gains from both variety expansion and improved efficiency.

Related Literature — This paper contributes to several strands of literature. First, it relates to models of firm innovation and dynamics in general equilibrium Klette and Kortum (2004); Acemoglu, Akcigit, Alp, Bloom and Kerr (2018); Akcigit and Kerr (2018). Our contribution is to incorporate empirically motivated pricing frictions into this framework. These frictions generate an endogenous markup distribution and price overshooting, which in turn shape product values, innovation incentives, and creative destruction.

Second, the paper connects to the macroeconomic literature on nominal rigidities and monetary non-neutrality. Traditional New Keynesian models emphasize short-lived price stickiness and its implications for business-cycle dynamics (e.g., Calvo, 1983; Woodford, 2003; Bils and Klenow, 2004; Klenow and Kryvtsov, 2008; Nakamura and Steinsson, 2008; Liu, 2025). A related contribution is Argente and Yeh (2022), who use weekly scanner-level data to show that young products adjust prices more frequently due to demand learning, highlighting short-run, age-dependent flexibility. More broadly, Adam and Weber (2023) use U.K. micro price data to document a negative trend in relative prices over a products' age and derive implications for optimal inflation targets under both Calvo and menu-cost frictions. Our findings complement this work by showing that life-cycle price dynamics matter not only for inflation targeting, but also for firms' innovation and growth decisions. In doing so, the paper also relates to empirical work on reference prices and infrequent adjustment (Eichenbaum, Jaimovich and Rebelo, 2011; Kehoe and Midrigan, 2015; Anderson, Jaimovich and Simester, 2015), but extends the focus to the long-run consequences for innovation and market dynamics. Importantly, our goal is not to explain the sources of these long-run frictions—such as the behavioral mechanisms proposed by Rebelo, Santana and Teles (2025)—but rather to quantify their implications for innovation and growth.

Third, this paper contributes to the literature on markups, misallocation, and the macroeconomic consequences of firm-level heterogeneity. Canonical endogenous growth models (Romer, 1990;

Aghion and Howitt, 1992; Grossman and Helpman, 1991) and recent quantitative frameworks linking innovation and market structure typically abstract from nominal frictions. A related literature studies how heterogeneous markups generate misallocation and affect aggregate outcomes (Edmond et al., 2022; Blanco et al., 2024; Peters, 2020). In particular, Peters (2020) shows that cross-sectional dispersion in markups across firms leads to endogenous misallocation with implications for growth. This paper departs from that literature by introducing a dynamic source of markup dispersion arising from pricing frictions and trend inflation. In our framework, nominal rigidities generate time-varying markups over the product life cycle, even in the absence of underlying productivity heterogeneity. This mechanism links pricing frictions to product turnover, innovation incentives, and long-run welfare, thereby providing a unified framework to study misallocation and growth in the presence of nominal rigidities.

Finally, we contribute to the literature exploiting tariff shocks to study cost pass-through (e.g., Fajgelbaum et al., 2020; Cavallo et al., 2021). Existing work examines how trade policy affects average prices, markups, and import patterns. We document a new asymmetry: tariff-induced cost increases are reflected in the prices of new products introduced in subsequent years. This provides novel evidence that long-run pricing frictions shape firms' adjustment margins and links trade shocks to the dynamics of product introduction.

The rest of paper is organized as follows. Section 2 presents the data and compute the inflation within the dataset. Section 3 shows main empirical findings. Section 4 presents the model. Section 5 estimates the model, present the fit of model, and presents the results and counterfactuals. Section 6 examines optimal R&D policies. Section 7 concludes. Appendix A contains detailed discussion of data and empirical robustness checks, and Appendix B contains additional quantitative results.

2 Data

In our analysis, we use product-level data to document the life-cycle behavior of prices for incumbent products and to quantify the relative pricing of newly introduced products. We then exploit tariff information to construct measures of firm- and module-level cost shocks, which we use to evaluate how these shocks transmit into the prices of incumbents and new products.

2.1 Product-level data

We primarily use scanner data from the NielsenIQ Retail Measurement Services (RMS), provided by the Kilts–Nielsen Data Center at the University of Chicago Booth School of Business. The RMS records weekly, store-level sales and quantities for every UPC sold through point-of-sale systems in participating retail outlets. The data cover roughly 30% of total U.S. expenditures on food and

beverages and about 2% of total household consumption (Beraja et al., 2019). The main strength of this dataset is its breadth and depth, both in the number of products and in the consistency of coverage over time (Argente et al., 2024). Our analysis spans 2006–2022, allowing us to track product prices over a very long horizon.

The dataset includes a wide variety of goods, ranging from food items to semi-durable non-food products. It contains over 1.7 million distinct UPCs, each classified into a hierarchical structure: 1,066 product modules, aggregated into 115 product groups, and then into 10 major departments.[2]

To link products to their producers, we use data from GS1 US, which provides information on the parent companies associated with each barcode prefix. Because GS1 US maintains the universe of company prefixes issued in the United States, we merge these prefixes with the UPCs observed in the RMS. This linkage allows us to map each product to its parent firm and, in turn, to characterize the full product portfolio of every firm represented in our sample.

Our baseline dataset aggregates information for each product across all stores in the sample at a quarterly frequency (or annually for product–year analyses). For every UPC, we observe prices, quantities, and sales, along with detailed product classifications (module, group, department), longitudinal characteristics (age and cohort), and firm-level attributes.

For each product u in period t, we define total sales Y_{ut} as the sum of sales across all stores and weeks in that period. Total quantity sold, y_{ut}, is defined analogously. The nominal price p_{ut} is constructed as the ratio of total sales to total quantity, which corresponds to the quantity-weighted average price.

To construct longitudinal variables, we infer product age and cohort from the timing of a product's first observed transaction. We define entry as the period in which a product first appears in the data and exit as the first period following its last observed sale. Products already present at the beginning of the sample are classified as left-censored, as their true entry dates cannot be identified. Similarly, products with transactions in the final period of the sample are classified as right-censored, since their exit cannot be observed.

At the firm level, we focus primarily on distinguishing new and incumbent firms and constructing measures of firm size, including the number of products offered and total firm sales.

Our baseline data include all participating stores and exclude private-label products. Table 1 reports descriptive statistics for this sample. To address potential concerns regarding the measurement of product entry and exit, we construct robustness samples that restrict attention to a balanced set of stores and drop products that do not record at least one transaction per quarter after

[2]The structure of the source data changed in 2020. Appendix X documents how we harmonize the classification to ensure consistency over the full 2006–2022 period.

Table 1: Summary Statistics of Products by Censoring

	All	By Censoring Type			
		Not Censored	Right	Left	Right & Left
Total # of products	1,789,738	937,647	372,172	395,179	84,740
Duration (quarters)					
average	21	15	22	25	68
less than 4	22	28	18	14	0
above 20	37	26	39	48	100
above 28	26	15	30	35	100
Revenue (quarterly, $1,000)					
mean	51	23	105	23	207
25th percentile	0.2	0.1	0.7	0	2.1
median	1.6	1.2	5.6	0.5	14
75th percentile	14	9.4	45	5	95
90th percentile	79	44	217	35	444
95th percentile	199	98	478	92	930
Price (quarterly)					
mean	0.097	0.14	0.17	-0.059	0.022
25th percentile	-0.36	-0.34	-0.25	-0.51	-0.36
median	0.13	0.18	0.18	-0.0014	0.015
75th percentile	0.62	0.69	0.64	0.43	0.42
90th percentile	1.1	1.2	1.1	0.9	0.85
95th percentile	1.4	1.5	1.4	1.2	1.1

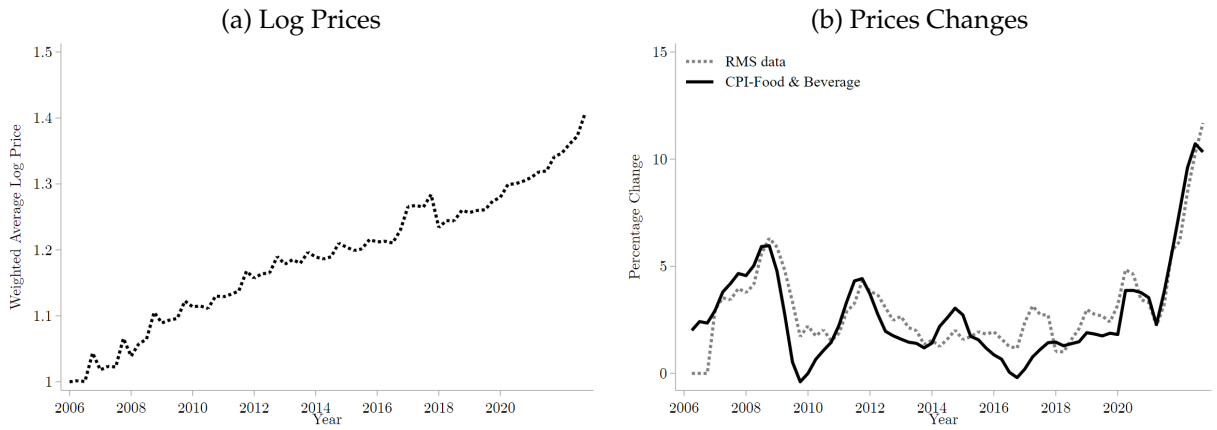
Notes: The table presents the summary statistics for the products included in the baseline sample for 2006q1–2022q4. Products already active in 2006q1–2006q2 are left-censored; products with sales in 2022q3–2022q4 are right-censored. "Not Censored" includes products observed from entry to exit. "Right Censored" indicates products still sold at the end of the sample. "Left Censored" refers to products already on the market at the start of the sample. "Right & Left" includes products with both entry and exit censored. Duration is in quarters. Revenue and price are quarterly. Revenue is in thousands of dollars. Prices are differenced from the sector×quarter median price.

entry. These samples also exclude private-label items and departments that are less representative. Appendix B reports the results based on these robustness specifications.

Our analysis further draws on product×store–level datasets, following Campbell and Eden (2014). Using the weekly RMS microdata, we construct measures of the frequency and size of price changes, aggregated to the product–store–year level. Appendix C provides details on these constructions.

Matching Aggregate Price Indexes– We evaluate how the average prices in our dataset relate to the Consumer Price Index (CPI) for Urban Consumers, specifically the Food and Beverages category. Using the product-quarter dataset, we compute the average prices and price changes

Figure 1: Price Indexes



Notes: Figure (a) shows the weighted average log price using the Nielsen RMS data at the quarterly-level, computed as $\log\bar{p}_t = \sum_m^M \omega_m \log\bar{p}_{m,t}$, where $\log\bar{p}_{m,t}$ is the average (log) prices of products in product module m at time t and ω_m represents the average sales weights of module m. The weighted average Log price is normalized to 1 at 2006q1. Figure (b) shows the quarterly price changes computed as in equation (1) and the one implied from CPI for all urban consumers: food-at-home.

at the sector-quarter level (here defined as RMS's product modules). To obtain a measure of the aggregate price index, we use the following formula:

$$\Delta_t = \sum_j^J \omega_j \Delta_{j,t}, \quad \Delta_{j,t} = \frac{P_{j,t} - P_{j,t-4}}{P_{j,t-4}} \times 100 \tag{1}$$

where ω_m represents the average sales weights, and $\Delta_{j,t}$ denotes the sector-specific price changes. Figure1(a) presents the sales-weighted log prices, while Figure1(b) illustrates Δ_t alongside the CPI index. The plot indicates that the average prices in our dataset broadly approximate the official inflation patterns, that was on average 3% annual in the period 2006-2022. The ability of RMS to approximate official inflation indexes is consistent with Argente and Lee (2020), who compare the Nielsen data using a Laspeyres price index and demonstrate that it closely matches the overall patterns of the CPI index.

2.2 Tariff-exposure

We construct a firm's measure of exposure to international tariffs and estimate its relationship with firm outcomes. Our approach integrates detailed tariff schedules, barcode-level country of origin information, and annual product-level data from NielsenIQ.

We begin by assembling a comprehensive dataset on product-level tariffs. Our primary source is the UNCTAD TRAINS database accessed through WITS, which provides tariff information at the exporter-country $\times$ HS code level. Specifically, we use the *Effectively Applied Tariff*—a weighted-average tariff that incorporates both MFN rates and preferential tariffs arising from

trade agreements. Because tariffs are defined at the HS8 level, we aggregate to HS6 using import-weighted averages. This measure excludes additive tariffs (e.g., fixed dollar amounts per kilogram). Using Bai and Stumpner (2019), we employ concordance tables to map HS codes into NielsenIQ module codes. For each country of origin, we generate country $\times$ module $\times$ year tariff series that can be matched to firm products.

The second component of the exposure measure is product-level country of origin (COO). We rely on LabelInsights, which provides country of origin for consumer goods products based on products' package. The data does not cover all products and some products do not report contry of origin in its package.

Product information comes from the NielsenIQ annual UPC-level data from 2010–2022. For each barcode, we construct a panel of firm identifiers, quantities sold, and revenues. This dataset forms the base structure onto which country of origin and tariff information are merged. We merge the UPC-level product data with COO information and subsequently with tariffs using the country-of-origin $\times$ module combination. Products that cannot be matched to country of origin or tariffs are retained to measure coverage and selection patterns.

For each firm i, module j, and year t, we construct as measures of tariff exposure:

$$T_{ijt}^{(1)} = \sum_{u \in j} w_{uijt} T_{uijt}, \qquad w_{uijt} = \frac{q_{uijt}}{\sum_{u \in j} q_{uijt}}. \tag{2}$$

All weights are computed using only products with tariff information. Alternative time-invariant weights will be estimated as robustness checks. Appendix X provides more details and summary statistics of the measures.

3 Prices of New and Incumbent Products: Stylized Facts

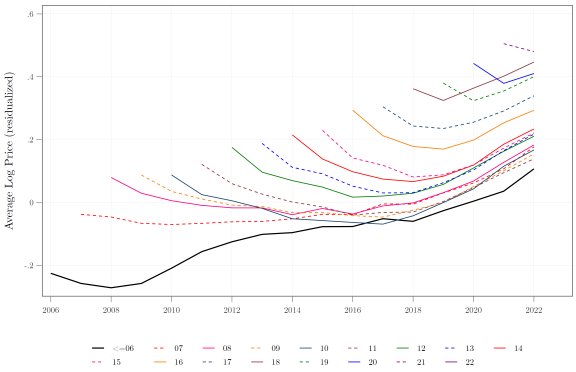
We start by providing stylized facts about the relative prices of new products and the evolution of prices after entry, followed by an analysis of the impact of cost costs (as proxied by a firms' tariff exposure) on the prices of new and incumbent products.

Fact 1: New Products Enter with Price Premia

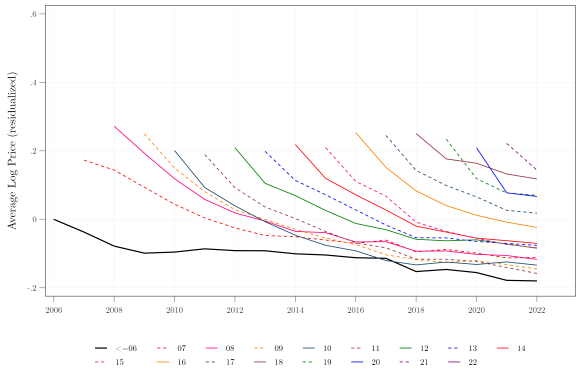
We begin by evaluating the relationship between the prices of new and incumbent products. We compute the average (log) prices by product generation to compare price levels. Using the product-year dataset, we address compositional effects by first regressing log prices on sector fixed effects and then using the residuals to compute the averages by cohort and time period. Figure 2 (a) presents the baseline patterns. The black line represents products created before or in 2006. Note that for products already in existence at the beginning of the period, cohort information cannot be

Figure 2: Prices by Cohort of Products

(a) Without Time Controls



(b) With Time Controls



Notes: The figures display average log prices by cohort over time using the baseline annual-level sample. In panel (a), log prices are residualized using sector fixed effects. In panel (b), log prices are residualized using sector-by-year fixed effects.

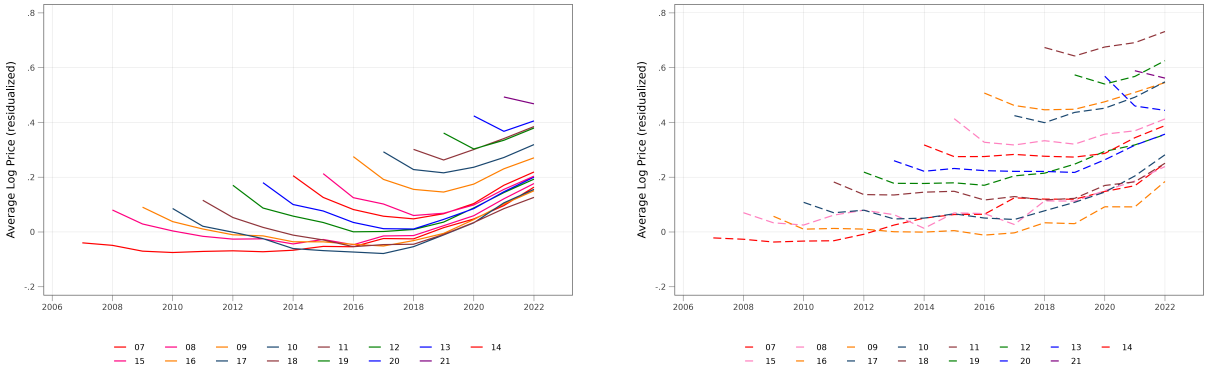
determined, so we bundle them together. All other lines show the average log prices (residualized by sector fixed effects) over time for each product cohort.

Figure 2 (a) shows that each new generation of products, on average, has higher prices than older generations during their initial years of activity (and, for most cohorts, prices remain higher than previous cohorts' prices throughout the entire period). We complement the analysis by evaluating the patterns when we also control for time controls. Figure 2 (b) presents equivalent plot when we residualized log prices with sector×year fixed effects. Relative to Figure 2 (a), the inclusion of time effects allows us to observe the average log prices of each generation of products relative to the average log prices of all products in the same sector at that time period. The results show that the prices of new products are on average 20% above the average price of products in the same sector.

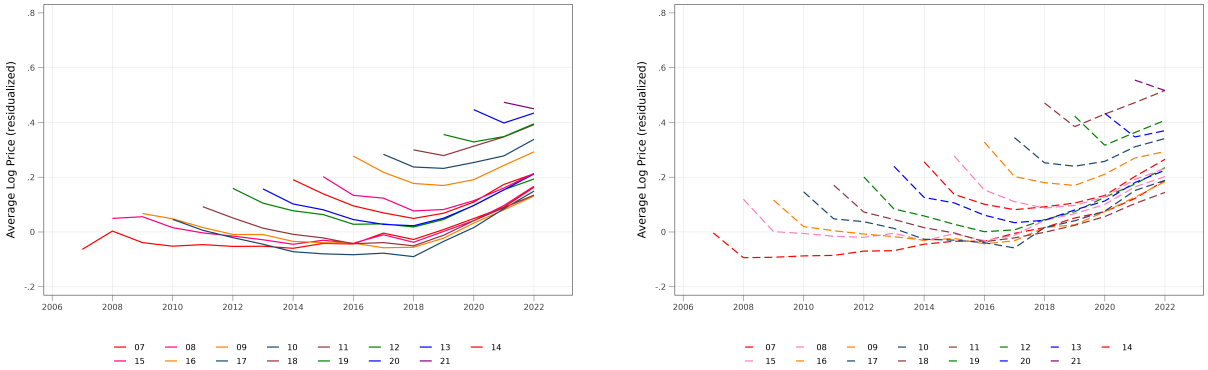
To further understand the importance of these results, we provide evidence that they are pervasive across different types of products. Figure 3 (a) explores heterogeneity by computing average log prices by cohort × incumbent/startups over time. Our results show that both new products introduced by new firms and by incumbent firms exhibit higher price levels than existing products. Figure 3 (b) explores heterogeneity by computing average log prices by cohort × novelty over time. We compute novelty for each product (as in Argente and Lee (2020)) and divide products in two groups, below and above average novelty. Our results show that higher prices of new product come from both high and low novelty products. Figure 3 (c) explores heterogeneity by computing average log prices by cohort × Non-food/food products over time. While the price premium of new products is higher for non-food product categories, they are also evident among food product categories.

Figure 3: Prices by Cohort of Products: Heterogeneity

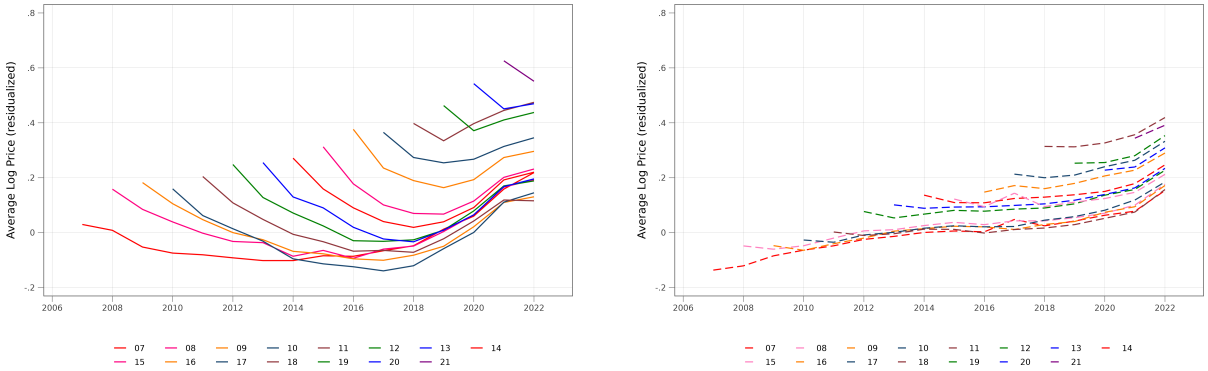
(a) Introduced by Incumbent (left) and New Firms (right)



(b) Products with High Novelty (left) and Low Novelty (right)



(c) Non-Food Products (left) and Food Products (right)



Notes: The figure shows the average residuals of (log) price on sector fixed-effects.

While it may not be surprising that new products are priced above existing ones—particularly if newer products offer higher quality or greater appeal—the magnitude of the observed price premium exceeds what is typically implied by standard models of quality ladders and endogenous

growth. In fact, existing estimates in the literature suggest that quality improvements alone would predict more modest increases in price. Furthermore, our analysis of heterogeneity in product novelty reveals no systematic relationship between higher prices and measures of product novelty. This suggests that the observed price premium for new products is not fully explained by differences in quality or degree of novelty relative to existing products.

Fact 2: Prices of Incumbent Products Remain Largely Stable Over Life Cycle

The patterns above allow us to trace the unconditional average prices upon entry. We now account for selection forces and examine how the prices of incumbent products evolve over their life cycle.

We characterize the life cycle of products by estimating the evolution of the prices of a product as a function of its age. We consider two types of specifications. First, we consider a specification that does not control for the fact that we observe products in different time periods. The estimation of log prices Y of product u at time t is expressed as function of age (a), sector(j) and cohort (c) effects:

$$\ln Y_{u,t} = \alpha + \sum_{a=2}^{A} \beta_a D_a D_c + \eta(1 - D_c) + \lambda_j + \theta_c + u_{u,t} \tag{3}$$

We are interested in the series of coefficients, β_a, that capture the average aging process of the products relative to the level of the outcome in the first full quarter of activity. We estimate age fixed effects for different sets of products (dummy D_c equal to 1), but still include other products in the regression to help estimate the other controls. In this specification we control for the fact that we have products from different sectors (sector fixed-effects λ_j) and that otherwise comparable products might behave differently depending on the timing of their entry (cohort fixed-effects θ_c).[3]

In the second type of specifications, we isolate the effect of age by accounting for the fact that we observe products in different time periods. In particular, we use sector$\times$time to account for the evolution of average prices in a particular product category:

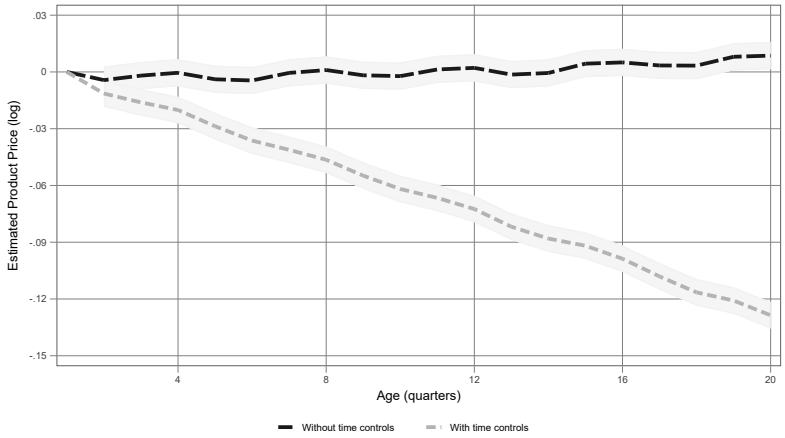
$$\ln Y_{u,t} = \alpha + \sum_{a=2}^{A} \beta_a D_a D_c + \eta(1 - D_c) + \lambda_{jt} + \theta_c + u_{u,t} \tag{4}$$

The sector$\times$time effects enable us to account for price changes that are common across products within the same sector, such as inflation and other calendar-related effects. By capturing age effects relative to the average price trajectory within that category, this specification allows us to approximate the evolution of real prices over a product's life cycle.

The baseline estimation focuses on age effects and is computed for products that remained active for at least 20 quarters. Figure 4 presents the estimated age fixed effects, both without and

[3]In practice we normalize the cohort effect as suggested in Deaton (1997). Later, when including time effects, there are problems with colinearity (there is an exact linear relation between the age-cohort-time effects).

Figure 4: Prices over the Product Life Cycle



Notes: The figure shows the estimated age fixed effects of (log) prices of products computed using equations (3) and (4). The estimation uses all products and the age fixed effects are computed for products from cohorts 2006q3-2016q3 and duration of more than 20 quarters. The gray area indicates the 95% confidence interval.

with time controls. The results show that, on average, the nominal prices of incumbent products remain largely unchanged over their life cycle (black dashed line). However, their prices relative to the average market product decline with age (gray dashed line). Between the first and fifteenth year of activity, relative product prices decline by an average of 3% per year.[4]

We evaluate the consistency of this result using several robustness exercises. The first set of exercises examines robustness across different sets of products. The baseline results are based on a balanced sample of long-lasting products, defined as products that survive for at least 20 quarters (i.e., balanced cohorts without exit). We also consider a semi-balanced sample that includes products from cohorts surviving at least five years, even if individual products exit earlier (i.e., balanced cohorts with exit), as well as an unbalanced sample including any product regardless of cohort survival or individual longevity (i.e., unbalanced cohorts with exit). Figure A1, panels (a) and (b), in the appendix presents the results. Exit-based selection plays a limited role, as results remain quantitatively similar across specifications—particularly when time controls are included. The only notable difference is that, without time controls, unconditional nominal prices are more likely to decline.[5] Another issue concerns whether to include products not used to estimate age fixed effects in the main sample as controls (i.e., using a dummy for products excluded from age

[4]This is consistent with Figure 1, which shows that average product prices increase over time.

[5]To further understand the nature of selection in our data set, we study the life cycle patterns conditional on the product's ex-post duration by estimating age fixed effects for products with durations of between 2 and 28 quarters. Figure A2, panels (a) and (b) in appendix confirm that prices conditional of time-effects of both short- and long lived products decline throughout their life cycles.

fixed effects). Figure A3 in the appendix presents results for the baseline sample (balanced cohorts without exit), estimated using only those products included in the regression. Sector, cohort, and time effects are thus identified using this restricted set. The results show qualitatively similar patterns, with differences in magnitude arising primarily with time controls, due to variation in the estimated fixed effects when comparing all products versus only those in their first 20 years of activity.

A second set of robustness exercises examines the sensitivity of the results to alternative specifications of control variables, with particular emphasis on the treatment of cohort effects. The baseline specification follows the normalization approach proposed by Deaton (1997). As an alternative, we estimate models that include cohort fixed effects, allowing the normalization of time and cohort components to be determined by the estimation procedure. In addition, we assess the robustness of the results to the inclusion of firm-level controls. Specifically, we incorporate firm$\times$sector$\times$time fixed effects to account for unobserved heterogeneity at the firm-sector level over time. Figures A4 and A5 in the appendix report the corresponding estimates. Across specifications, the results remain qualitatively similar, indicating that neither the method of controlling for cohort effects nor the inclusion of firm-specific controls materially affects the main findings.

The third set of robustness exercises assesses the potential role of store selection in shaping the results. Our baseline measure of prices is constructed as a weighted average of product-level prices across all stores in the United States. As documented by Kaplan and Menzio (2015), there exists substantial price dispersion across retail outlets. Consequently, variation in the composition of stores selling a product at different points in its life cycle may influence the observed price dynamics. To address this concern, we utilize our product-store-year dataset to examine how prices evolve at the store level as a function of product age, defined as the time elapsed since the product was first sold in a given store. Figure C.1 in the Appendix presents the evolution of (log) prices over the store-level product life cycle. The results indicate that store-level prices, like the national averages, do not exhibit systematic growth over time. This consistency suggests that the national-level findings are not driven by changes in store composition and hold even at the most disaggregated level of analysis.

Overall, our findings indicate that the nominal prices of incumbent products remain largely unchanged over the course of the product life cycle, while real prices decline slowly and steadily. This pattern is particularly striking given the positive average inflation rate in these sectors, which might lead one to expect upward trends in nominal prices over time.

At first glance, this result may seem inconsistent with the findings of Bils and Klenow (2004); Klenow and Kryvtsov (2008); Nakamura and Steinsson (2008), who report that prices typically change every four to twelve months. Our data likewise show that products undergo frequent price adjustments (Table X in Appendix C). However, these changes tend to be *not* persistent

and *do not* vary systematically with product age (Appendix C). This pattern suggests that firms often engage in temporary pricing actions—such as sales, promotions, or short-run responses to business-cycle fluctuations—while maintaining relatively stable average prices over a product's life cycle. In other words, even though nominal prices change frequently, these adjustments do not bring prices back in line with the aggregate inflation trend. For our purposes of studying the relationship between pricing frictions and innovation, it is essential to distinguish short-lived deviations associated with temporary price changes from trend-restoring adjustments that occur over the long run. Accordingly, the relevant component of price stickiness for our analysis is the one that evolves with the product life cycle rather than with high-frequency, temporary price movements.

Fact 3: Price Premium of New Products Depend on Price Frictions of Incumbent Products

Does price rigidity affect the initial price premium of new products relative to incumbent products—referred to as price overshooting? While we do not take a position on the underlying causes of our finding that nominal prices of incumbent products remain largely unchanged over the product life cycle, we seek to understand whether such frictions influence innovation. To this end, we first present cross-sectional evidence showing a strong association between the degree of price rigidity and the extent of price overshooting among new products. We then complement this evidence by showing that cost shocks—captured through firms' exposure to import tariffs—have substantial effects on the relative prices of new products, but not on the prices of incumbent products.

Cross-Sectional Evidence

We estimate equation 4 separately for each product category and construct category-specific measures of price frictions. Our data reveal substantial heterogeneity across product categories in the degree of price rigidity, as reflected in the flatness of the price life-cycle profile. Roughly 75% of categories display declining life-cycle price paths, with an average annual slope of -0.015. Table 2, columns (1)–(3), shows a significant negative association between these life-cycle estimates and the price premium of new products relative to incumbents. This implies that product categories with greater price rigidity—that is, flatter or more negative life-cycle profiles—exhibit larger price overshooting. In other words, the inability of firms to adjust prices of incumbent products is systematically associated with higher entry premia for new products.

To further assess this cross-sectional relationship, we consider alternative measures of price rigidity. In particular, we use product-category measures of the average frequency of price changes,

Table 2: Price Overshooting of New Products and Proxies for Price Frictions

	Price Overshooting New Products					
	(1)	(2)	(3)	(4)	(5)	(6)
Frequency Price Changes	-0.596***	-0.591***	-0.269***			
	(0.036)	(0.036)	(0.057)			
Life Cycle Estimated				-2.438***	-2.437***	-1.299***
				(0.094)	(0.094)	(0.113)
Observations	58,207	58,207	58,207	58,207	58,207	58,207
R-squared	0.011	0.018	0.037	0.005	0.011	0.035
Time FE	N	Y	Y	N	Y	Y
Sector FE	N	N	Y	N	N	Y

Notes: The table reports the association between two proxies for the frequency of price changes and the price overshooting of new products, using module–quarter variation, covering 2007q1-2022q4. "Frequency Price Changes" is the average module-level frequency of price adjustments, as defined in Appendix X. "Life Cycle Estimated" is the annualized average price growth over the product life cycle, estimated separately for each module using equation 4. Price overshooting is defined as $\tilde{P}_{m,t} = \left(P^{\text{new}}_{m,t} / P^{\text{inc}}_{m,t}\right) - 1$. Sectors correspond to NielsenIQ product groups.

following Nakamura and Steinsson (2008).[6] The underlying assumption is that short-run frictions—such as sales, promotions, or responses to business-cycle conditions—are correlated with long-run frictions that shape the evolution of average prices over the product life cycle. Table 2, columns (4)–(6), shows that categories with more frequent price changes exhibit lower price overshooting, consistent with the interpretation that stronger price rigidities (i.e., lower adjustment frequencies) are associated with larger new-product entry price premia.

Although our evidence does not speak directly to why some products exhibit greater price rigidity, it suggests that when firms face difficulties adjusting prices of incumbent products, they are more likely to use new products as a margin through which to adjust the prices of their overall portfolio.

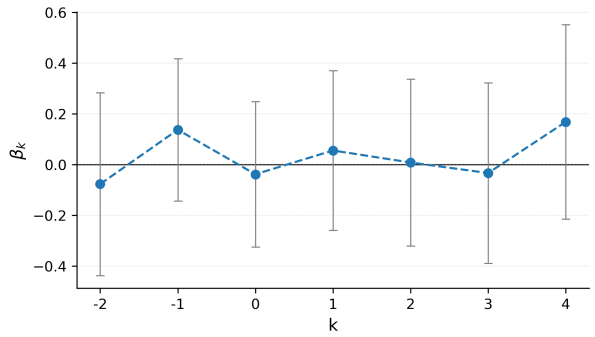
Evidence from Tariffs

To evaluate whether price frictions in incumbent products affect the pricing decisions of new products, we exploit variation in firms' exposure to import tariffs as a proxy for cost shocks to existing products. A substantial share of consumer goods is imported from multiple countries. When tariffs change, firms that sell imported products in the U.S. experience corresponding changes

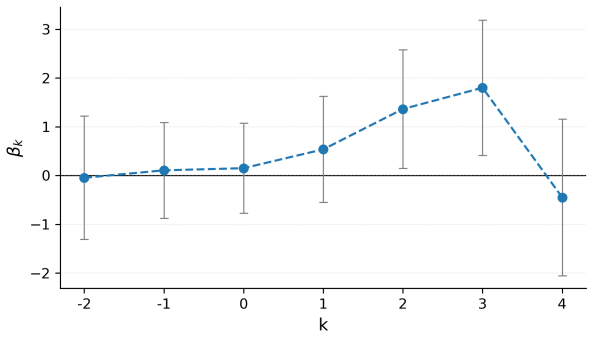
[6]We construct measures of the frequency of price changes using weekly store–product scanner data and then average them at the category level. See Appendix C for details.

Figure 5: Tariff Exposure and Prices

(a) Price of Incumbent Products



(b) Relative Price New Products



Notes: The figure plots the estimated coefficients after estimating equation $y_{i,j,t+k} = \beta T_{i,j,t} + \omega_{j,t+k} + \phi_{i,j} + \varepsilon_{i,j,t+k}$, $k = -2, ..., 0, ..., 4$, where $T_{i,j,t}$ is the level of tariff exposure. The outcome in panel (a) is the log prices of incumbent products $P^{\text{inc}}_{i,j,t+k}$, and in panel (b) is the relative price of new products $\tilde{P}_{i,j,t+k} = \left(P^{\text{new}}_{i,j,t+k} / P^{\text{inc}}_{i,j,t+k}\right) - 1$. The vertical bands represent +/- 1.96 $\times$ st. error of each point estimate.

in their cost of supplying these goods. We use this variation to study how tariff exposure affects the prices of incumbent products and the price premia of newly introduced products.

We construct measures of tariff exposure for firm i in product category j and year t, as defined in equation 2.2, and examine how these measures correlate with prices of both old and new products. In particular, we estimate specifications of the form

$$y_{i,j,t+k} = \beta\, T_{i,j,t} + \omega_{j,t+k} + \phi_{i,j} + \varepsilon_{i,j,t+k}, \tag{5}$$

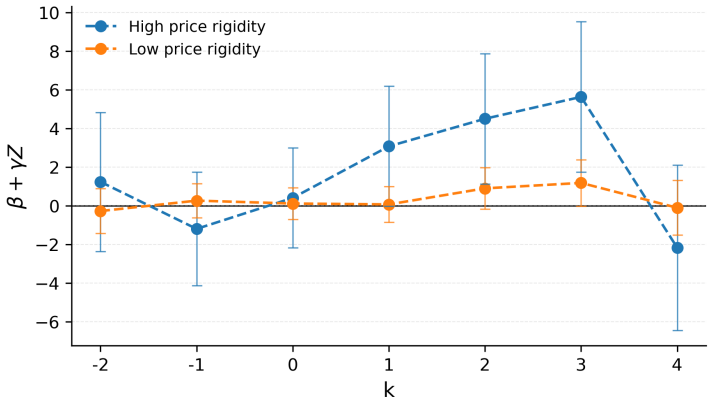
where $y_{i,j,t+k}$ is the outcome of interest (e.g., log prices of incumbent products or the relative price of new products) and $T_{i,j,t}$ is the tariff-exposure measure. We include year–category fixed effects to absorb changes common to all firms within a category, such as aggregate shifts in production costs, demand, or market power. We also include firm–category fixed effects to control for time-invariant differences in price levels and baseline tariff exposure across firms.[7]

Figure A2 presents the estimated effect of tariff exposure on the average prices of incumbent products and on the relative price of new products. The results show no significant effect of a firm's tariff exposure on the prices of incumbent products (panel (a)). By contrast, tariff exposure has a significant positive effect on the future relative prices of new products: a 0.10 percentage point increase in tariff exposure is associated with an increase of approximately 0.14 percentage points in the relative price of new products two years later.[8]

[7] In some specifications, we additionally control for the share of a firm's quantity covered by non-missing tariff information.

[8] The figure shows that it takes about two years for tariff changes to pass through to the prices of new products. This is consistent with the fact that the introduction of new products is relatively infrequent for many firms (CITE).

Figure 6: Tariff Exposure and Relative Price of New Products



Notes: The figure plots the estimated impact of tariffs for product categories with high and low frequencies of price change. The coefficients are estimated using equation $y_{i,j,t+k} = \beta T_{i,j,t} + \gamma T_{i,j,t} \times Z_j + \omega_{j,t+k} + \phi_{i,j} + \varepsilon_{i,j,t+k}$, $k = -2, ..., 0, ..., 4$, and evaluating the estimated impact when $Z = -1$ (high price rigidity) and $Z = 1$ (low price rigidity). The variable $T_{i,j,t}$ is the level of tariff exposure. The outcome is the relative price of new products $\tilde{P}_{i,j,t+k} = \left(P^{\text{new}}_{i,j,t+k} / P^{\text{inc}}_{i,j,t+k}\right) - 1$. The variable Z_j is the standardized frequency of price changes at the product category level. The vertical bands represent +/- 1.96 $\times$ st. error of each point estimate.

To further assess the robustness of the positive association between tariff exposure and the future price premium of new products, we estimate a series of alternative specifications. Table A2 in Appendix reports the estimated coefficients from equation 6 with $k = 2$, including specifications that add additional controls. In columns (3)–(6), we replace the level of tariff exposure with its change. Across all specifications, the results consistently show a significant positive impact of a firm's tariff exposure on the future relative price of new products.[9] In all cases, tariff exposure has no significant effect on the log price of incumbent products (Appendix Table X).

Finally, we examine how the impact of tariff exposure on the relative price of new products varies across product categories with different degrees of price rigidity. The underlying hypothesis is that if incumbent prices adjust sluggishly in some product categories, tariffs should induce larger price overshooting among the new products introduced in those sectors. To test this idea, we estimate the specification:

$$y_{i,j,t+k} = \beta\, T_{i,j,t} + \gamma\, T_{i,j,t} \times Z_j + \omega_{j,t+k} + \phi_{i,j} + \varepsilon_{i,j,t+k}, \tag{6}$$

where Z_j is a sector-specific measure of price frictions. As a proxy for these frictions, we use the standardized frequency of price changes (defined in Appendix C).

Figure 6 plots the estimated effects of tariffs on the relative price of new products separately for product categories with high and low price rigidity, defined as one standard deviation above

[9]Table X in the Appendix reports analogous results using the log difference between new and incumbent product prices.

and below the mean, respectively. The results are striking: the positive impact of tariffs is several orders of magnitude larger in high-rigidity product categories. Consistent with our identifying assumptions, there are no significant pre-trends, and by three years after the tariff shock the differential effect between high and low rigidity categories dissipates. Table A3 in the Appendix reports the corresponding coefficient estimates.

Taken together, our findings indicate that when firms face greater frictions in adjusting prices of incumbent products, they set higher prices for subsequently introduced products. We interpret these relationships as causal in episodes where tariff changes are unexpected or not fully anticipated by firms. A growing literature shows that many tariff changes occur under substantial policy uncertainty or are implemented with limited predictability, particularly during episodes of trade conflict, retaliatory measures, or policy renegotiations.

4 Model

This section develops a Schumpeterian model of creative destruction with endogenous innovation, embedded in a monetary general equilibrium environment with trend money growth and nominal pricing frictions. The economy is in continuous time and admits a balanced growth path (BGP) on which real allocations are stationary after appropriate normalization. The innovation block follows the multi-product framework in Klette and Kortum (2004) and the quality-ladder structure in Aghion and Howitt (1992); Acemoglu et al. (2018); the monetary side adopts a money-in-the-utility formulation commonly used in menu-cost models (e.g., Golosov and Lucas, 2007).

4.1 Monetary Economy

Time is continuous. A representative household consumes the final good C_t, holds nominal money balances M_t, and owns a portfolio of nominal assets A_t (claims to firms and government bonds) that yields the nominal return R_t. The preference is:

$$\int_0^\infty e^{-\rho t} \left(\frac{C_t^{1-\vartheta} - 1}{1-\vartheta} + \ln\left(\frac{M_t}{\hat{P}_t}\right) \right) dt, \tag{7}$$

where $\rho > 0$ is the discount rate, $\vartheta > 0$ is the inverse intertemporal elasticity of substitution, and $\hat{P}_t$ is the quality-adjusted price index defined below. The household supplies labor inelastically, $L_t = L$. The nominal budget constraint is:

$$\dot{A}_t + \dot{M}_t = W_t L + R_t A_t - \hat{P}_t C_t, \tag{8}$$

where W_t is the nominal wage. The first-order conditions imply (i) the Euler equation and (ii) a standard money-demand equation:

$$R_t - \pi_t = \rho + \vartheta g_t, \tag{9}$$

$$R_t M_t = \hat{P}_t C_t^{\vartheta}, \tag{10}$$

where $\pi_t \equiv \dot{\hat{P}}_t / \hat{P}_t$ is inflation and $g_t \equiv \dot{C}_t / C_t$ is the consumption growth rate.

The monetary authority supplies nominal balances M_t^s that grow at the constant rate μ:

$$\frac{\dot{M}_t^s}{M_t^s} = \mu, \tag{11}$$

and the money market clears, $M_t = M_t^s$. Along a BGP, growth rates are constant. Combining (10) with $M_t = M_t^s$ yields the BGP restriction:

$$\pi + (\vartheta - 1)g = \mu, \tag{12}$$

so trend money growth maps into constant inflation on the BGP.

4.2 Production, Product Lines, and Demand

There is a unit measure of product lines. At any time t, a measure $N_t \in (0, 1]$ of lines is active and produces intermediate goods; the remaining $1 - N_t$ lines are inactive. The final good is a CES aggregate of quality-adjusted intermediate inputs:

$$Y_t = \left(\int_{i \in \mathcal{N}_t} (q_{it} y_{it})^{\frac{\sigma-1}{\sigma}} \, di \right)^{\frac{\sigma}{\sigma-1}}, \tag{13}$$

where $\sigma > 1$ is the elasticity of substitution, q_{it} denotes product quality, and y_{it} is the quantity of intermediate input i. Intermediate goods are produced with labor under a linear technology,

$$y_{it} = A_t \, \ell_{it}, \tag{14}$$

where ℓ_{it} is labor used in line i and A_t is an aggregate productivity shifter (normalized to one if desired). The demand for intermediate good i is:

$$y_{it} = q_{it}^{\sigma-1} \left(\frac{p_{it}}{\hat{P}_t} \right)^{-\sigma} Y_t, \tag{15}$$

where p_{it} is the nominal price of good i and the quality-adjusted price index is

$$\hat{P}_t = \left(\int_{i \in \mathcal{N}_t} \left(\frac{p_{it}}{q_{it}} \right)^{1-\sigma} di \right)^{\frac{1}{1-\sigma}}. \tag{16}$$

For later use, define the (normalized) quality index

$$Q_t \equiv \left(\frac{1}{N_t} \int_{i \in \mathcal{N}_t} q_{it}^{\sigma-1} \, di \right)^{\frac{1}{\sigma-1}}, \tag{17}$$

and relative quality $\hat{q}_{it} \equiv q_{it}/Q_t$, which satisfies

$$\frac{1}{N_t} \int_{i \in \mathcal{N}_t} \hat{q}_{it}^{\sigma-1} \, di = 1. \tag{18}$$

4.3 Pricing, Markups, and Profits

The nominal marginal cost is same across firms, $\frac{W_t}{A_t}$. We define the product-level markup μ_{it} by

$$p_{it} = \mu_{it} \frac{W_t}{A_t}. \tag{19}$$

Given (101), markup dispersion affects aggregate allocations through the quality-adjusted price index. Define the quality-weighted aggregate markup

$$m_t \equiv \left(\frac{1}{N_t} \int_{i \in \mathcal{N}_t} \mu_{it}^{1-\sigma} \hat{q}_{it}^{\sigma-1} \, di \right)^{\frac{1}{1-\sigma}} = \left(\mathbb{E}_{\hat{q}} \left[\mu_{it}^{1-\sigma} \right] \right)^{\frac{1}{1-\sigma}} = \frac{\hat{P}_t Q_t A_t}{N_t^{\frac{1}{1-\sigma}} W_t}, \tag{20}$$

where $\mathbb{E}_{\hat{q}}[\cdot]$ denotes expectation under the quality-weighted cross-sectional distribution over active lines. On a BGP, both $\hat{P}_t Q_t$ and W_t grow at the inflation rate, implying that m_t is constant.

Combining (101) with labor-market clearing yields

$$Y_t = A_t N_t^{\frac{1}{\sigma-1}} \, v_t \, Q_t \, L, \qquad v_t \equiv \frac{m_t^{-\sigma}}{\mathbb{E}_{\hat{q}} \left[\mu_{it}^{-\sigma} \right]}. \tag{21}$$

where v_t measures the aggregate efficiency loss from the dispersion of markups across products. Compared to the efficient case where all firms charge the same markup, greater dispersion in markups implies that the household over-consumes low-markup goods and under-consumes high-markup goods, leading to misallocation and lower aggregate efficiency. In addition, the term $N_t^{1/(\sigma-1)}$ captures the standard love-for-variety effect. The labor share in quality-adjusted terms satisfies

$$\frac{W_t L}{\hat{P}_t Y_t} = (v_t m_t)^{-1} = \frac{\mathbb{E}_{\hat{q}} \left[\mu_{it}^{1-\sigma} \right]}{\mathbb{E}_{\hat{q}} \left[\mu_{it}^{-\sigma} \right]}, \tag{22}$$

which is a sales-weighted harmonic mean of markups, consistent with the result in Edmond et al. (2022). Finally, the real profit flow of product i (in units of the quality-adjusted final good) is

$$\frac{\pi_{it}}{\hat{P}_t} = m_t^{\sigma-1} N_t^{-1} Y_t \, \hat{q}_{it}^{\sigma-1} \left[(\mu_{it} - 1) \mu_{it}^{-\sigma} \right] = \hat{\pi}(\hat{q}_{it}, \mu_{it}) Y_t. \tag{23}$$

Profits are proportional to market size Y_t and increasing in aggregate markups. Holding aggregate conditions fixed, higher relative quality raises profits, while profits are hump-shaped in the product's own markup with the static optimum at $\mu^* = \sigma/(\sigma-1)$. With nominal pricing frictions (introduced below), firms face a dynamic trade-off between current and future profits.

4.4 Pricing Frictions and Exit Decision

Firms face nominal rigidities. A product's price is set optimally upon introduction. Thereafter, the firm receives opportunities to reset the nominal price at Poisson rate γ, as in Calvo pricing models. Between reset opportunities, the nominal price remains fixed while the aggregate price level grows at the constant inflation rate π along the BGP. Consequently, the product-level markup evolves deterministically according to

$$\frac{\dot{\mu}_{it}}{\mu_{it}} = -\pi, \tag{24}$$

so markups erode over time in the absence of adjustment.

We assume that a product is *active* in production only while its markup exceeds an exit threshold. Specifically, for each product i there exists a lower cutoff μ_{ib} such that continued operation is optimal if and only if $\mu_{it} > \mu_{ib}$; once μ_{it} falls to μ_{ib}, the firm shuts down production and the product is no longer supplied, so consumers can no longer purchase or consume that variety. We refer to this event as *deactivation* (or exit) of the product line. Importantly, deactivation is a production decision rather than a loss of technological capability: while the product disappears from the consumption set, the incumbent retains the know-how associated with the line and can continue to conduct R&D that builds on it. As a result, innovation remains feasible on inactive lines.

After deactivation, the incumbent cannot simply resume producing the same vintage via a price reset. Reactivation of the product line occurs only through innovation: the incumbent may upgrade the line by introducing a higher-quality product, or another firm may innovate on the line and displace it from the current owner. We describe the innovation technology in detail below.

When a reset opportunity arrives for an *active* product, the firm chooses a reset markup μ_{iu} to maximize continuation value. A product's life cycle is therefore characterized by an initial markup μ_{iu}, deterministic erosion at rate π between adjustment opportunities, stochastic reset opportunities arriving at rate γ, and deactivation when μ_{it} reaches μ_{ib}.

4.5 Innovation

Innovation improves product quality and reallocates ownership of product lines through creative destruction. Incumbent firms conduct two types of R&D. *Internal innovation* upgrades product lines already owned by the firm and, in this environment, plays an additional central role: it is a primary way for firms to circumvent pricing frictions on their existing products. Because nominal prices are otherwise slow to adjust, markups on continuing products erode mechanically with inflation; by innovating internally, the firm can effectively "relaunch" the line and re-optimize its price, restoring its markup and resetting the product's pricing cycle.

External innovation improves upon product lines not currently owned by the firm; successful innovators acquire the improved line, displacing the previous owner if the line was active.[10] New entrants innovate in the same manner and enter upon success. Crucially, *external and entrant innovation are untargeted*: an innovation attempt draws a product line uniformly at random from the unit measure of lines, which may be active or inactive and may have any quality. A firm's size is the number of product lines it operates; when a firm loses a line its portfolio shrinks, and if the displaced line is its last remaining line the firm exits. As with internal innovation, a successful external innovation introduces a new higher-quality vintage that is priced optimally upon introduction. Importantly, external innovation provides an additional channel for firms to offset the erosion of markups caused by nominal rigidity at the *portfolio level*, not just within existing products. In this sense, internal innovation relaxes price rigidity on continuing lines, while external innovation refreshes the firm's product mix with newly priced lines.

Following Klette and Kortum (2004), innovation is based on the firm's knowledge capital, proxied by the number of product lines in its portfolio. Formally, a firm operating n lines runs n research programs in parallel, so the arrival rate of innovation opportunities scales proportionally with n. Accordingly, we work with per-line innovation intensities and aggregate innovation costs by summing across the firm's portfolio.

Consider a product line with relative quality $\hat{q}$. If the incumbent chooses internal innovation intensity $x^I > 0$, it incurs a flow cost (in units of final goods)

$$c^I(x^I, \hat{q}, \mu) = \xi^I \left(x^I\right)^{\frac{1}{1-\alpha}} G(\mu, \hat{q}), \tag{25}$$

where $\xi^I > 0$ is an R&D cost shifter and $\alpha \in (0, 1)$ governs the elasticity of R&D costs with respect to innovation intensity. G is the cost function that may depend on markup and relative quality, which we specify in detail below. Conditional on success, quality improves multiplicatively,

$$q' = \lambda^I q, \qquad \lambda^I > 1, \tag{26}$$

and the firm resets the product's price at the time the upgraded version is introduced.

External innovation is untargeted in the sense described above. An incumbent chooses per-line external innovation intensity $x^E > 0$ and pays the cost for the random quality it draws,

$$c^E(x^E, \hat{q}) = \xi^E \left(x^E\right)^{\frac{1}{1-\alpha}} \hat{q}^{\sigma-1}, \tag{27}$$

where $\xi^E > 0$. Conditional on success, the innovator acquires the drawn line (displacing the current owner if the line was active), upgrades its quality by

$$q' = \lambda^E q, \qquad \lambda^E > 1, \tag{28}$$

[10] If the targeted line is inactive, external innovation reactivates the line and assigns ownership to the innovator.

and resets the price upon takeover. Entrants behave analogously: an entrant chooses innovation intensity $x^N > 0$ at the cost

$$c^N(x^N, \hat{q}) = \xi^N \left(x^N\right)^{\frac{1}{1-\alpha}} \hat{q}^{\sigma-1}, \tag{29}$$

and, upon success, acquires the randomly drawn line, upgrades quality by step size $\lambda^N > 1$, and sets an optimal reset price.

The economy-wide rate of creative destruction due to takeovers is

$$\tau \equiv x^E + x^N. \tag{30}$$

Since each product line undertakes untargeted external innovation at rate x^E and successful entrant innovations arrive at rate x^N, each successful external or entrant innovation transfers ownership of the affected line to the innovator. Hence τ summarizes the arrival rate at which product lines experience ownership turnover in the economy.

Because successive vintages within a line may differ in both quality and price, a standard CES demand system would in principle allow multiple vintages to coexist within a line. To preserve the canonical structure that at most one product is operated per line, we impose a simple selection device. Each product line has a single "slot" (shelf space) controlled by a retailer. The retailer receives an ε share of the operating profit generated by the stocked vintage and therefore chooses the vintage that maximizes the line's profit flow. We take $\varepsilon \to 0$, so the device affects selection but not equilibrium payoffs. Under this device, only one vintage is operated on each line in equilibrium, and when a new (improved) vintage arrives, the previously operated vintage is immediately deactivated.

4.6 Obsolescence and Creation of New Product Lines

To match the average product life observed in the data, we assume that product lines become obsolete exogenously at rate ψ. An equal measure of new product lines is created so that the total mass of lines remains constant and normalized to one. These newly created lines are operated by new producers. Empirically, large firms innovate much more intensively (see Akcigit and Kerr 2018), so it is natural to interpret these new producers as new establishments of existing large firms that expand into new industries or newly created lines (Aghion et al., 2023). Importantly, such producers may operate multiple new lines, distinguishing them from a traditional "entrant" that is modeled as a brand-new firm with a single product. To preserve a balanced growth path, we assume that products introduced on newly created lines start at the aggregate quality.

At each instant, events occur in sequence. (i) Internal innovation upgrades a measure x^I of lines, reactivating $x^I(1-N)$ inactive lines. (ii) Incumbent and entrant external innovation induces takeovers at rate τ, reactivating $\tau(1-N)$ inactive lines. (iii) A fraction ψ of lines becomes obsolete

and is replaced one-for-one by newly created lines which are produced by new firms, implying an additional inflow $\psi(1-N)$ into active status. Finally, lines with markups below μ_b are deactivated.

4.7 Value Functions

We normalize all the growing variables in the value function by Y_t to keep the equilibrium values constant on the BGP, i.e., $V = \hat{V}Y_t$. Let $\Omega_i = \{\hat{q}_j, \mu_j\}_{j=1}^{M_i}$ be the set of product qualities and markups that an incumbent firm i holds, including both active and inactive. Since each firm's innovation cost is independent across products, the firm's value function can be expressed as the sum of individual product values. This result is standard in the literature (Klette and Kortum, 2004). Formally,

$$\hat{V}(\Omega_i) = \sum_{j=1}^{M_i} \max\Big\{\Gamma_1(\hat{q}_j, \mu_j),\, \Gamma_0(\hat{q}_j)\Big\}, \tag{31}$$

where $\Gamma_1(\hat{q}_j, \mu_j)$ is the value of product line j when it is *active* (the firm produces and sells the current vintage), and $\Gamma_0(\hat{q}_j)$ is the value when the line is *inactive* (the firm shuts down production so the flow profit is zero). In both regimes, the line remains subject to exogenous firm exit at rate ψ and creative destruction at rate τ, and the firm can conduct internal and external innovation.

Active line. For an active line with state $(\hat{q}, \mu)$, the line-level HJB is

$$\begin{aligned} r\Gamma_1(\hat{q},\mu) = &-\pi\mu\frac{\partial\Gamma_1(\hat{q},\mu)}{\partial\mu} - g\hat{q}\frac{\partial\Gamma_1(\hat{q},\mu)}{\partial\hat{q}} + \hat{\pi}(\hat{q},\mu) - F - (\tau+\psi)\Gamma_1(\hat{q},\mu) \\ &+ \max_{x^I\geq 0}\Big\{x^I\big[\Gamma_1(\lambda^I\hat{q},\mu_u(\lambda^I\hat{q})) - \Gamma_1(\hat{q},\mu)\big] - \xi^I(x^I)^{\frac{1}{1-\alpha}}G(\hat{q},\mu)\Big\} \\ &+ \max_{x^E\geq 0}\Big\{x^E\mathbb{E}\big[\Gamma_1(\lambda^E\hat{q},\mu_u(\lambda^E\hat{q}))\big] - \mathbb{E}\xi^E(x^E)^{\frac{1}{1-\alpha}}\hat{q}^{\sigma-1}\Big\} \qquad (32) \\ &+ \gamma\Big[\Gamma_1(\hat{q},\mu_u(\hat{q})) - \Gamma_1(\hat{q},\mu)\Big]. \qquad (33) \end{aligned}$$

where the expectation in external innovation is taken over the cross-sectional distribution of *all* product lines, including both active and inactive lines.

Inactive line. For an inactive line with relative quality $\hat{q}$, the line-level HJB is

$$\begin{aligned} r\Gamma_0(\hat{q}) = &-g\hat{q}\frac{\partial\Gamma_0(\hat{q})}{\partial\hat{q}} - F - (\tau+\psi)\Gamma_0(\hat{q}) \qquad (34) \\ &+ \max_{x^I\geq 0}\Big\{x^I\big[\Gamma_1(\lambda^I\hat{q},\mu_u(\lambda^I\hat{q})) - \Gamma_0(\hat{q})\big] - \xi^I(x^I)^{\frac{1}{1-\alpha}}G(\hat{q})\Big\} \\ &+ \max_{x^E\geq 0}\Big\{x^E\mathbb{E}\big[\Gamma_1(\lambda^E\hat{q},\mu_u(\lambda^E\hat{q}))\big] - \mathbb{E}\xi^E(x^E)^{\frac{1}{1-\alpha}}\hat{q}^{\sigma-1}\Big\}. \qquad (35) \end{aligned}$$

The two line-level HJBs define the continuation value of a product line under two regimes: *active* (produced) and *inactive* (not produced), while allowing innovation in either state. In both equations, the term $(r+\psi)\Gamma_s$ (with $s \in \{1,0\}$) is the required return on the line's value, reflecting

discounting at rate r and an exogenous obsolescence hazard ψ of product lines. The right-hand side collects the sources of value common to both regimes: relative quality declines at rate g, so $-g\hat{q}\,\partial_{\hat{q}}\Gamma_s$ captures the capital loss from falling behind the frontier; the fixed cost F is paid to maintain the line's know-how and R&D capacity even when production is shut down; and the firm chooses internal and external innovation intensities (x^I, x^E) by trading off the expected value gain from a successful upgrade, which yields an active line at new price and higher quality, against convex R&D costs. Hence, deactivation does not drive value to zero: it eliminates operating profits but preserves the real option to restore production through innovation.

The regimes differ only in the components tied to production and pricing. When the line is active, it earns flow profits $\hat{\pi}(\hat{q}, \mu)$ and its markup is a state variable that erodes deterministically with inflation, generating the drift term $-\pi\mu\,\partial_\mu\Gamma_1$. Active lines also receive Calvo reset opportunities at rate γ, producing the jump term $\gamma\big[\Gamma_1(\hat{q}, \mu_u(\hat{q})) - \Gamma_1(\hat{q}, \mu)\big]$. When the line is inactive, flow profits are zero and the markup is irrelevant, so there is no μ-drift and no reset term; the line's value is pinned down solely by quality drift, fixed costs, and the innovation options that may reactivate it.

To maintain the tractability while preserving the key economic mechanism—that internal innovation both upgrades quality and resets the line's price—we impose a disciplined restriction on the internal-innovation cost shifter $G(\cdot)$ and fixed cost of operation.

Assumption 1 (Fixed cost and internal-innovation cost shifter)**.** 1. ***Fixed cost equals the external-innovation option value.*** *The per-period fixed operating cost equals the equilibrium per-line option value of external innovation:*

$$F = \Lambda^E. \tag{36}$$

2. ***Internal-innovation cost shifter.*** *The internal-innovation cost shifter satisfies*

$$G(\hat{q}, \mu) = \hat{q}^{\sigma-1}\,\frac{\Gamma_1(\lambda^I\hat{q}, \mu_u(\lambda^I\hat{q})) - \Gamma_1(\hat{q}, \mu)}{\Gamma_1(\lambda^I\hat{q}, \mu_u(\lambda^I\hat{q}))}, \qquad G(\hat{q}) = \hat{q}^{\sigma-1}\,\frac{\Gamma_1(\lambda^I\hat{q}, \mu_u(\lambda^I\hat{q})) - \Gamma_0(\hat{q})}{\Gamma_1(\lambda^I\hat{q}, \mu_u(\lambda^I\hat{q}))}. \tag{37}$$

The first part of Assumption 2 ensures the separability between markup and quality in the value function, which renders a clean ODE on markup dynamics as shown below. The second part is a targeted restriction on the internal-innovation cost shifter. It delivers two properties: (i) it preserves *separability* by maintaining the $\hat{q}^{\sigma-1}$ scaling; and (ii) it removes the nonlinearity in value function introduced by the maximization over x^I by making the effective gain of internal innovation independent of the current markup. Economically, the restriction preserves the idea that a successful internal innovation both upgrades quality and embeds an option value of resetting prices. In addition, the option value of internal innovation still increases with the distance between the current markup and optimal reset markup. The cost of this tractability is that the model

no longer allows internal innovation rate to depend on the current markup; instead, innovation intensity is constant with respect to μ and across regimes.

Under Assumption 2, the model admits a tractable characterization. The following lemma shows the direct result of these assumptions.

Lemma 1 (Separability and common markup thresholds)**.** *Under Assumption 2, the following properties hold:*

- ***Separation.*** *The value of a product line factorizes as*

$$\Gamma_1(\hat{q}, \mu) = \hat{q}^{\sigma-1} A(\mu), \quad \Gamma_0(\hat{q}) = \hat{q}^{\sigma-1} \bar{A} \tag{38}$$

  where $A(\mu)$ depends only on the markup μ.

- ***Common reset and exit thresholds.*** *The optimal pricing policies are quality-invariant: there exist constants (μ_u, μ_b) such that the optimal reset markup and the exit cutoff are the same for all lines.*

- ***Independence of quality and markups.*** *The relative-quality process (and hence its stationary distribution) is independent of the markup process and exit decisions.*

- ***External innovation option value.*** *The per-line option value of external innovation equals*

$$\Lambda^E = \xi^E \frac{\alpha}{1-\alpha} \left(x^{E*}\right)^{\frac{1}{1-\alpha}}, \qquad x^{E*} = \left(\frac{1-\alpha}{\xi^E} A(\mu_u)(\lambda^E)^{\sigma-1}\right)^{\frac{1-\alpha}{\alpha}}. \tag{39}$$

Proof. See Appendix D. ■

Lemma 2 yields two tractability gains. First, since both operating profits and R&D costs scale with $\hat{q}^{\sigma-1}$, the line value is homogeneous in relative quality and factorizes as $\Gamma(\hat{q}, \mu) = \hat{q}^{\sigma-1} A(\mu)$. The dynamic pricing problem therefore reduces to the one-dimensional function $A(\mu)$, and the associated policies are *quality-invariant*: the optimal reset markup μ_u and exit cutoff μ_b are common across lines, because quality scales payoffs without altering the trade-offs governing exit.

Second, because all types of innovations are feasible on both active and inactive lines and are untargeted with respect to quality and ownership, the quality ladder evolves independently of the markup process and exit decisions. Exit shuts down production and affects profits, but does not impede innovation on the line. Hence the stationary distribution of $\hat{q}$ is pinned down by the innovation block alone, while the stationary distribution of markups is determined by inflation-driven erosion, Calvo resets, and the common cutoff μ_b. In particular, an untargeted innovation draw has expected quality equal to the aggregated quality index defined on active lines, Q_t, so the expected relative quality is equal to one. Finally, the external-innovation problem admits closed-form expressions for Λ^E and x^{E*}.

Proposition 1. *The value function $A(\mu)$ is characterized by*

$$(\varphi+\gamma)A(\mu) = -\pi\mu\frac{\partial A(\mu)}{\partial\mu} + m^{\sigma-1}N^{-1}(\mu-1)\mu^{-\sigma} + \alpha x^{I*}\Big((\lambda^I)^{\sigma-1}A(\mu_u) - A(\mu)\Big) + \gamma A(\mu_u) \qquad \text{if } \mu > \mu_b, \tag{40}$$

$$\varphi\bar{A} = \alpha x^{I*}\Big((\lambda^I)^{\sigma-1}A(\mu_u) - \bar{A}\Big) \qquad \text{otherwise}$$

where $\varphi = r+\psi+\tau+g(\sigma-1)$. The optimal internal-innovation intensity x^{I*} is,

$$x^{I*} = \left(\frac{1-\alpha}{\xi^I}(\lambda^I)^{\sigma-1}A(\mu_u)\right)^{\frac{1-\alpha}{\alpha}}. \tag{41}$$

The exit (inactivity) markup μ_b is pinned down by value matching and smooth pasting,

$$m^{\sigma-1}N^{-1}(\mu_b-1)\mu_b^{-\sigma} + \gamma(A(\mu_u)-\bar{A}) = 0. \tag{42}$$

Proof. See Appendix D. ■

The product line has two modes. When active ($\mu > \mu_b$), it earns operating profits and its markup erodes deterministically with inflation; it may be displaced by creative destruction and exit due to exogenous shock and occasionally receives Calvo opportunities to reset to the optimal markup μ_u. When inactive ($\mu \leq \mu_b$), the line no longer produces, so operating profits vanish and Calvo resetting is irrelevant. The only channel that remains is innovation, which is feasible regardless of active status and, upon success, upgrades quality and reintroduces the line at the reset markup. This generates the piecewise HJB: in the inactive region the line's value is pinned down solely by the option value of innovation, while in the active region it equals profits plus the innovation option plus the value of potential reset. The exit cutoff μ_b is characterized by smooth pasting: at the point of exit, the firm is indifferent, and the marginal value of continuing production is zero, yielding the boundary condition (42). Since $A(\mu_u)$ is larger than $\bar{A}$ by definition, μ_b must be smaller than one, implying negative profit at the threshold. Intuitively, a firm may still be willing to retain a product with negative flow profit because, first, keeping the product active enables the firm to engage in innovation based on its knowledge about this product, and second, there is a possibility that the firm may be able to adjust its price. In general, μ_b decreases as γ increases.

The following proposition shows the solution to the above differential equation in closed form.

Proposition 2. *Given μ_b and μ_u, let $\kappa = \varphi+\gamma+\alpha x^{I*}$. Then the value function $A(\mu)$ is given by,*

$$A(\mu) = \Big[\xi_1\mu^{1-\sigma}\left(1-(\frac{\mu_b}{\mu})^{\frac{\kappa}{\pi}+1-\sigma}\right) - \xi_2\mu^{-\sigma}\left(1-(\frac{\mu_b}{\mu})^{\frac{\kappa}{\pi}-\sigma}\right) + \xi_3\left(1-(\frac{\mu_b}{\mu})^{\frac{\kappa}{\pi}}\right)\Big] + \bar{A} \tag{43}$$

where

$$\xi_1 = \frac{m^{\sigma-1}N^{-1}}{\pi}(\frac{\kappa}{\pi}-\sigma+1)^{-1}; \quad \xi_2 = \frac{m^{\sigma-1}N^{-1}}{\pi}(\frac{\kappa}{\pi}-\sigma)^{-1}; \quad \xi_3 = \frac{(\gamma+\alpha x^{I*}(\lambda^I)^{\sigma-1})A(\mu_u)}{\kappa} - \bar{A}$$

Proof. See Appendix D. ■

Expression (43) decomposes the line value into three components. The first two terms are the present value of operating profits along the deterministic markup path $\dot{\mu}/\mu = -\pi$; they arise from integrating the flow-profit kernel $(\mu - 1)\mu^{-\sigma}$ with discounting at the effective rate κ. The third term captures the value of future price resets: with Poisson arrival rate γ, the firm can jump back to the reset markup μ_u, generating continuation value $A(\mu_u)$. The boundary condition at μ_b truncates all components through the terms $(\mu_b/\mu)^{\kappa/\pi}$, reflecting the fact that when the markup falls to μ_b the product stops generating operating profits. The magnitude of this truncation is governed by the ratio κ/π: higher inflation speeds up markup erosion and makes the exit boundary matter more for any given μ, while a higher effective discount rate shortens the product life cycle and reduces the effect of pricing frictions.

4.8 Optimal Reset Markup

Until now, we have characterized the value function given the optimal reset markup μ_u. The following proposition characterizes the optimal reset markup and its property.

Proposition 3. *The optimal reset markup μ_u satisfies,*

$$\Big[\varphi - \alpha x^{I*}\big((\lambda^I)^{\sigma-1} - 1\big)\Big] A(\mu_u) = m^{\sigma-1} N^{-1} (\mu_u - 1)\mu_u^{-\sigma}. \tag{44}$$

*In particular, the optimal reset markup exhibits **price-overshooting**, i.e., it is strictly greater than the frictionless markup $\mu^* = \frac{\sigma}{\sigma-1}$. In the limiting case where $\pi = 0$ or $\gamma \to \infty$, we recover the frictionless benchmark, i.e., $\mu_u = \mu^*$.*

Proof. See Appendix D. ■

The optimal reset markup μ_u is determined by the firm's problem of maximizing the value of the product conditional on the ability to reset its price. Specifically, the firm chooses μ_u to maximize $A(\mu_u)$, the value function evaluated at the reset markup. Although the value of innovation depends on μ_u, its elasticity with respect to μ_u does not directly enter the solution. This is because the innovation value is maximized *whenever* the product value at μ_u is maximized following equation (41).

The optimal reset markup is strictly larger than the static optimal markup μ^*—a phenomenon we refer to as *price-overshooting*. To see the intuition, take the derivative of both sides of equation (123) and rearrange:

$$m^{\sigma-1} N^{-1} (\mu_u - \frac{\sigma}{\sigma-1}) \mu_u^{-\sigma-1} (\sigma - 1) = -\pi \mu_u A''(\mu_u) \tag{45}$$

The key step in the proof is to observe that the value function $A(\mu)$ is increasing and concave around μ^*, which implies a higher optimal reset markup μ_u. Intuitively, when firms face the

possibility of being unable to reset prices in the future as the price level rises, the optimal price should reflect a forward-looking average of future desired prices. Overshooting allows firm's profit to be centered around static optimal markup over average duration of a price cell, leading to higher average profit. This mechanism mirrors the role of expectations of inflation in driving inflation dynamics in New Keynesian models (Werning, 2022). Combine equations 44 and 42, we can jointly solve μ_u and μ_b.

In the limits, as $\pi \to 0$ or $\gamma \to \infty$, the value function (123) evaluated at μ_u becomes,

$$\varphi A(\mu_u) = m^{\sigma-1} N^{-1} (\mu_u - 1) \mu_u^{-\sigma} + \Lambda(\mu_u) \tag{46}$$

Since the value of innovation is maximized when the profit flow is maximized, the optimal reset markup is equal to $\frac{\sigma}{\sigma-1}$. Intuitively, in both cases, the pricing friction is effectively zero. Consequently, there is no incentive to overshoot the optimal markup, as doing so would reduce the firm's profit flow without offering any additional dynamic advantage.

4.9 Markup Distributions and Aggregate Markup

Let H denote (unnormalized) markup distribution for products. The following proposition characterizes the markup distribution, the measure of products, and derived aggregate markup and aggregate efficiency.

Proposition 4. *The (unnormalized) markup distribution $H(\mu)$ for products satisfy,*

$$H(\mu) = N \frac{(\frac{\mu}{\mu_b})^{\phi} - 1}{(\frac{\mu_u}{\mu_b})^{\phi} - 1}, \tag{47}$$

where $\phi = \frac{\tau + x^I + \gamma + \psi}{\pi}$. The measure of products is given by,

$$N = (\tau + x^I + \psi) \Big[\phi\pi \big((\frac{\mu_u}{\mu_b})^{\phi} - 1 \big)^{-1} + \tau + x^I + \psi \Big]^{-1}, \tag{48}$$

The normalized aggregate markup m and the aggregate efficiency v are given by,

$$m = \left[\frac{\phi}{\phi - \sigma + 1} \frac{\mu_u^{\phi-\sigma+1} - \mu_b^{\phi-\sigma+1}}{\mu_u^{\phi} - \mu_b^{\phi}} \right]^{\frac{1}{1-\sigma}}, \tag{49}$$

$$v = m^{-\sigma} \frac{\mu_u^{\phi} - \mu_b^{\phi}}{\frac{\phi}{\phi-\sigma} \left(\mu_u^{\phi-\sigma} - \mu_b^{\phi-\sigma} \right)}, \tag{50}$$

When $\phi \to \infty$, they become,

$$\lim_{\phi \to \infty} m = \frac{\sigma}{\sigma - 1}, \quad \lim_{\phi \to \infty} v = 1, \quad \lim_{\phi \to \infty} N = 1. \tag{51}$$

Proof. See Appendix D. ■

The markup distribution is defined on the support $[\mu_b, \mu_u]$ and has total mass N. The parameter ϕ measures how concentrated the distribution is around μ_u. Intuitively, a lower rate of inflation implies that, once prices are reset to μ_u, they drift away more slowly. Similarly, a higher probability of price adjustment and innovation rate means that a larger share of products maintain markups near μ_u. Anticipating that markups will remain close to μ_u for an extended period, firms choose μ_u closer to the static optimal markup μ^*, which further reduces the dispersion of markups, thereby mitigates misallocation and leads to higher v. In the limit, as $\phi \to \infty$, the normalized aggregate markup converges to the static optimum, and the misallocation component vanishes. To roughly see how v moves with ϕ, we can ignore terms with μ_b. Then, we can show

$$v \approx \left(\frac{\phi}{\phi - \sigma + 1}\right)^{\frac{1}{\sigma-1}} \frac{\phi - \sigma}{\phi - \sigma + 1} \tag{52}$$

It is easy to show that v increases with ϕ, consistent with the above intuition. In general, the aggregate markup and efficiency depend on distance between μ_u and μ_b as well.

The measure of active products is determined by endogenous exit: firms endogenously shut down products when the markup falls below the threshold μ_b. Lower ϕ increases the likelihood of exit through this channel. Exogenous obsolescence shock does not affect the measure directly because there are new firms carry new product lines immediately back to the economy which are active. In the limit, when inflation or pricing friction is absent, the measure of products is always one. The measure of products has a direct impact on welfare through love-for-variety channel.

4.10 Entrants, Expected Lifetime, Growth Rate and Consumption Share

There is a unit measure of potential entrants, which are brand-new firms in the data. Each entrant can innovate by introducing a new variety. They first randomly draw a quality from the quality distribution of existing products. Once they innovate successfully, they improve such quality by λ^N. The optimization problem for entrants is given by,

$$\max_{x^N} \left\{ x^N \mathbb{E}\left[\Gamma(\lambda^N \hat{q}, \mu_u) - \xi^N x^{N\frac{1}{1-\alpha}} \hat{q}^{\sigma-1}\right] \right\} \tag{53}$$

where $E\Gamma$ is the expected value of entry. The arrival rate of innovation by entrants is,

$$x^{N*} = \left(\frac{1-\alpha}{\xi^N} A(\mu_u)(\lambda^N)^{\sigma-1}\right)^{\frac{1-\alpha}{\alpha}} \tag{54}$$

To derive the growth rate, notice that equation (106) shows that aggregate output is proportional to the quality index Q_t. Therefore the growth rate of aggregate output is given by $g = \dot{Q}_t / Q_t$. Also, the goods market clearing implies $Y_t = C_t + D_t + FN_t$, where D_t is the aggregate R&D cost and FN_t is the aggregate fixed cost. The expected product lifetime must account for multiple exit

and renewal channels: (i) displacement through endogenous innovation (creative destruction), (ii) exogenous obsolescence shocks, (iii) endogenous deactivation when the markup falls below the cutoff μ_b, and (iv) stochastic price resets that restart the product's pricing cycle. The following proposition characterizes the growth rate, consumption share and expected life of a product.

Proposition 5. *The growth rate of the economy is equal to*

$$g = x^{E*}\frac{(1+\lambda^E)^{\sigma-1}-1}{\sigma-1} + x^{I*}\frac{(1+\lambda^I)^{\sigma-1}-1}{\sigma-1} + x^{N*}\frac{(1+\lambda^N)^{\sigma-1}-1}{\sigma-1} \tag{55}$$

The share of consumption of total final goods,

$$\chi_c = 1 - \left(\xi^I x^{I*\frac{1}{1-\alpha}}\mathbb{E}_\mu\Big[\frac{(\lambda^I)^{\sigma-1}A(\mu_0)-A(\mu)}{A(\mu_0)}\Big] + \xi^E x^{E*\frac{1}{1-\alpha}} + \xi^N x^{N*\frac{1}{1-\alpha}}\right) - FN \tag{56}$$

The expected lifetime of a product,

$$\mathbb{E}T = \frac{1-(\mu_b/\mu_u)^{(\tau+\psi+x^I+\gamma)/\pi}}{\tau+\psi+x^I+\gamma(\mu_b/\mu_u)^{(\tau+\psi+x^I+\gamma)/\pi}} \tag{57}$$

Proof. See Appendix D. ■

The replacement of products with higher-quality varieties implies an increase in the aggregate quality index and thus allows us to derive the growth rate of the economy. Aggregate consumption equals total output net of the total R&D costs incurred by incumbents and entrants.

4.11 Equilibrium and Welfare

The equilibrium of the above economy is as follows:

Definition 1 (Stationary Equilibrium)**.** *A stationary equilibrium of this economy is a tuple*

$$\{y_i, p_i, x^I, x^E, \mu_b, \mu_u, A(\mu), \Gamma(\hat{q},\mu), m, v, \bar{H}(\mu), H(\mu), g, r, \chi_c\}$$

such that (i) y_i and p_i maximize profits as in (101); (ii) $A(\mu), \Gamma(\hat{q},\mu)$ solve value functions in (123) and (μ_b, μ_u) satisfies (44) and (42); (iii) x^I, x^E, x^N are given by the R&D policy functions in (41), (127) and (132); (iv) markup distributions are given by (47). The measure of active lines satisfies (48); (v) the aggregate consumption share is given by (56); (vi) the growth rate and real interest rate solve (55) and (9).

Now, we derive the welfare of the representative household in this economy. Recall that consumption is a constant share of aggregate output, denoted by χ_c. Since the quality index is already normalized, it facilitates the comparison of different economies. To this purpose, we set $Q_0 = 1$. The initial consumption is therefore $C_0 = \chi_c Y_0 = \chi_c v N^{\frac{1}{\sigma-1}}$. We only consider the welfare

derived from consumption and ignore the utility from money. Then the welfare of the above economy can be obtained as,

$$U_0(C_0, g) = \int_0^\infty e^{-\rho t} \cdot \frac{C_t^{1-\vartheta} - 1}{1-\vartheta} \, dt = \frac{1}{1-\vartheta} \left[\underbrace{\frac{\chi_c^{1-\vartheta}}{\rho - (1-\vartheta)g}}_{\text{growth-related}} (vN^{\frac{1}{\sigma-1}})^{1-\vartheta} - \frac{1}{\rho} \right] \tag{58}$$

Welfare in this economy is shaped by three key components. First, there is a growth-related term that depends on the economy's growth rate and the consumption share. A higher growth rate reflects a faster pace of innovation, given the step size, but it also implies greater R&D investment and hence fewer resources available for current consumption, introducing a tradeoff. Second, welfare is negatively affected by misallocation arising from markup dispersion. Greater dispersion distorts the allocation of resources across firms and reduces aggregate efficiency. Third, the number of products contributes to welfare through love of variety. This term is weighted by the elasticity of substitution: when goods are more substitutable, the marginal value of additional varieties is lower. Importantly, the aggregate markup does not directly enter the welfare expression, as labor is supplied inelastically and profits are fully redistributed. However, the markup still plays a critical role indirectly—it affects innovation incentives by shaping the profitability of products and the effective cost of innovation.[11]

In comparing welfare in two economies, say economy 1 and 2, and growth rates $g(1)$ and $g(2)$ and initial consumption levels $C_0(1)$ and $C_0(2)$, we compute consumption-equivalent changes in welfare by considering the fraction of initial consumption κ that will ensure the same discounted utility. Formally, the consumption-equivalent change κ is given such that,

$$U_0(\kappa C_0(2), g(2)) = U_0(C_0(1), g(1))$$

Or equivalently, we have $\kappa = (\frac{\rho-(1-\vartheta)g_1}{\rho-(1-\vartheta)g_2})^{\vartheta-1} \frac{C_1}{C_2}$. Then, we can define the log-change of consumption-equivalent welfare,

$$\log \kappa = \underbrace{\frac{1}{(\vartheta-1)} \log(\frac{\rho-(1-\vartheta)g_1}{\rho-(1-\vartheta)g_2}) + \log(\frac{\chi_{c1}}{\chi_{c2}})}_{\text{growth-related}} + \underbrace{\log(\frac{v_1}{v_2})}_{\text{misallocation}} + \underbrace{\frac{1}{\sigma-1} \log(\frac{N_1}{N_2})}_{\text{variety}} \tag{59}$$

[11] It is straightforward to eliminate the wedge on the labor supply condition due to aggregate markup using a production subsidy Edmond et al. (2022). Here, we are interested in how such production subsidy can also affect markup dispersion and available varieties.

5 Model Estimation and Results

In this section, we begin by estimating the model and assessing its fit using both targeted and untargeted moments. We then present equilibrium outcomes and quantify the role of price overshooting via counterfactual experiments.

5.1 Estimation

Our model has 12 structural parameters: $\{\vartheta, \sigma, r, \alpha, \pi, \xi^I, \xi^E, \xi^N, \lambda^I, \lambda^E, \lambda^N, \gamma, \psi\}$. First, we set the discount rate to $\rho = 2\%$. The median elasticity of substitution across products in the NielsenIQ data is set to 6 as in Argente et al. (2024). Following Acemoglu et al. (2018) and Akcigit and Kerr (2018), we set the elasticity of innovation with respect to R&D, α, to 0.5. The average inflation rate in our dataset is $\pi = 2.65\%$. The elasticity of intertemporal substitution is chosen to be 3, which provides reasonable welfare implications when considering optimal policies, as discussed in Jones (2024). Finally, we set the quality improvement steps to be the same across incumbents and entrants, i.e., $\lambda^I = \lambda^E = \lambda^N = \lambda$. We plan to relax this assumption in future iterations.

The remaining five parameters $\{\xi^I, \xi^E, \xi^N, \lambda, \gamma\}$ are calibrated to four product-level moments. We choose five moments that have closed-forms expressions in theory. In particular, we minimize the following distance,

$$\min \sum_{i=1}^{5} \omega(i) \frac{|\text{Model}(i) - \text{Data}(i)|}{|\text{Data}(i)|}$$

We assign more weights to growth rate and price-overshooting.

Table 3 summarizes the empirical moments used in the calibration. We use product-level data to construct innovation rates that correspond closely to their model counterparts. We measure the incumbent internal innovation rate as follows. For each firm, we compute the ratio of newly introduced products in the product categories in which the firm is already active to the firm's total product stock. This captures innovation that expands the firm's existing product lines. The incumbent external innovation rate is defined as the ratio of newly introduced products in product categories that are new to the firm to the firm's total product stock. This measure captures expansion into new product categories. By construction, the sum of internal and external innovation rates equals the firm's overall innovation rate, defined as the total number of newly introduced products relative to its total product stock. Both incumbent innovation measures are computed at the firm level and then averaged across firms within each sector. Finally, we measure the entrant innovation rate as the share of newly introduced products by entrant firms within a sector, relative to the total number of products in that sector.

Besides the entry rates, we used moments for price overshooting, aggregate growth rate and expected life cycle of a product. For price overshooting we compute the ratio of average price of

Table 3: Internal calibration of the model

Moments	Equation	Data	Model
Incumbents' internal innovation rate	x^I/N	8.92%	8.92%
Incumbents' external innovation rate	x^E/N	1.10%	1.10%
New entrant's entry rate	x^N/N	0.25%	0.25%
Aggregate growth rate	g	1.71%	1.71%
Price over-shooting (in log)	$\log(\mu_u) - \mathbb{E}_\mu \log(\mu)$	8.76%	8.76%
Expected life of product	$\mathbb{E}T$	6.5	6.5

Notes: The table presents the moments used in the calibration algorithm. The moments are estimated using data across all product groups, weighting sectors by their sales share. The moments are computed using data at the firm $\times$ sector $\times$ quarter level for the period 2006–22.

new products relative to old products within sector. The aggregate growth rate correspond to the average growth of sales by firms. The expected life cycle is computed as discussed in Table 1.

Table 3 shows that our model is able to match the four empirical moments closely, despite the model parsimony. Importantly, the data suggests that the internal innovation rate is significantly higher than the external innovation rate. This finding is consistent with (Garcia-Macia et al., 2019), who argue that improvements to a firm's own products contribute more to growth than creative destruction using manufacturing firms dataset. A high internal innovation rate indicates that firms primarily rely on quality improvements to their existing products to overcome pricing frictions, rather than using external innovation to acquire product lines from other firms and reset the overall price level of their product portfolio. Additionally, the entry rate of new firms is relatively low (about 0.25%), while incumbents account for much higher rates of product turnover. This suggests that most product churning, quality improvement, and price resetting are driven by incumbents engaged in internal innovation.

Table 4 presents the baseline calibration, with all parameters expressed at annual frequency. Notably, the Calvo parameter γ is estimated to be low, implying that it takes a firm about 9 years on average to adjust its price back to trend. Given that the average product life span is about 7 years[12], this suggests that a vast fraction of products do not adjust their prices at all during their lifetime following initial introduction.

As discussed in Section 3 this result may seem inconsistent with the findings of Nakamura and Steinsson (2008); Bils and Klenow (2004); Eichenbaum et al. (2011), who report the average duration of a price cell is about 3-4 quarters. However, this apparent discrepancy arises from a crucial distinction. While prices are indeed adjusted frequently, these adjustments do not necessarily bring prices back in line with the aggregate inflation trend. In fact, they tend to be temporary sales,

[12]We calculate the average product lifetime non-parametrically using Kaplan–Meier estimator.

Table 4: Model Parameters

Parameter	Description	Identification	Value
ϑ	Inverse EIS	External calibration	3
σ	Elasticity of substitution	External calibration	6
α	Innovation elasticity	External calibration	0.5
ρ	Discount rate	External calibration	0.02
π	Inflation rate	External calibration	2.65%
ξ^I	Internal innovation cost of incumbents	Internal calibration	11.02
ξ^E	External innovation cost of entrants	Internal calibration	89.52
ξ^N	Innovation cost of entrants	Internal calibration	393.18
γ	Frequency of price adjustment	Internal calibration	0.110
λ	Step size in innovation	Internal calibration	1.141
ψ	Exogenous obsolescence rate	Internal calibration	4.58%

Notes: The table reports the baseline parameter calibration. Parameters labeled "External calibration" are informed by the literature; those labeled "Internal calibration" are chosen to match key moments in the data.

promotions, and price responses to shocks on the business-cycle frequency. In our framework, the Calvo parameter γ does not measure the overall frequency of nominal price changes; rather, it captures the slow-moving life-cycle component of pricing frictions, which is relevant for innovation decisions.

5.2 Results

In this section, we present values of key equilibrium variables and evaluate the fit of our model to different non-targeted moments in the data, and conduct counterfactuals.

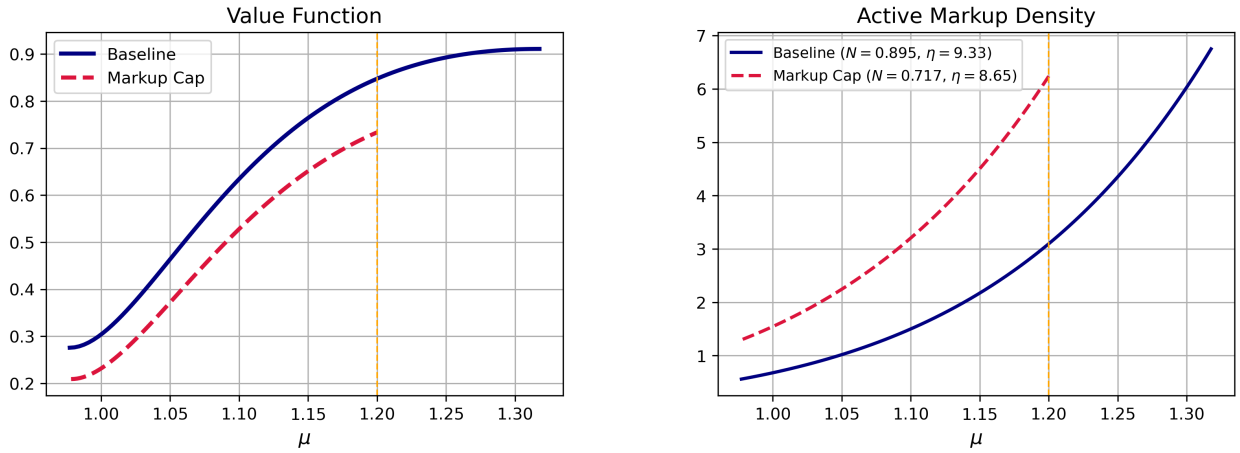
Equilibrium Values. Table 5 reports the key equilibrium outcomes in our baseline economy. The average creative destruction rate τ is about 1%, indicating a mall product turnover caused by entrants and external innovation each year. The optimal reset markup, μ_0, is 10% higher than the static optimal markup $\mu^* = 1.2$, reflecting the impact of infrequent price adjustments as established in Proposition 9. The exit threshold μ_b falls below one due to the option value of price-resetting, as shown in Proposition 1. Given the inflation rate $\pi = 2.65\%$, firms are willing to tolerate negative profits for about one year.

The distance between μ_0 and μ_b corresponds to roughly thirteen years of price drift due to inflation. Given price adjustment probability of 11% and the total innovation rate of 9.2% ($x^I + x^E + x^N$), the share of products that ultimately reach μ_b and exit endogenously is small. As a result, the measure of available varieties remains high, at roughly $N = 0.9$. The left panel of Figure

Table 5: Baseline Economy

x^I	x^E	x^N	μ_u	μ_b	N	χ_c	v	m	g	Life	Wel
7.980%	0.982%	0.224%	1.318	0.977	0.895	0.969	98.08%	1.190	1.71%	6.5	100

Figure 7: Value Function and Markup Density for Baseline and Markup Cap Cases



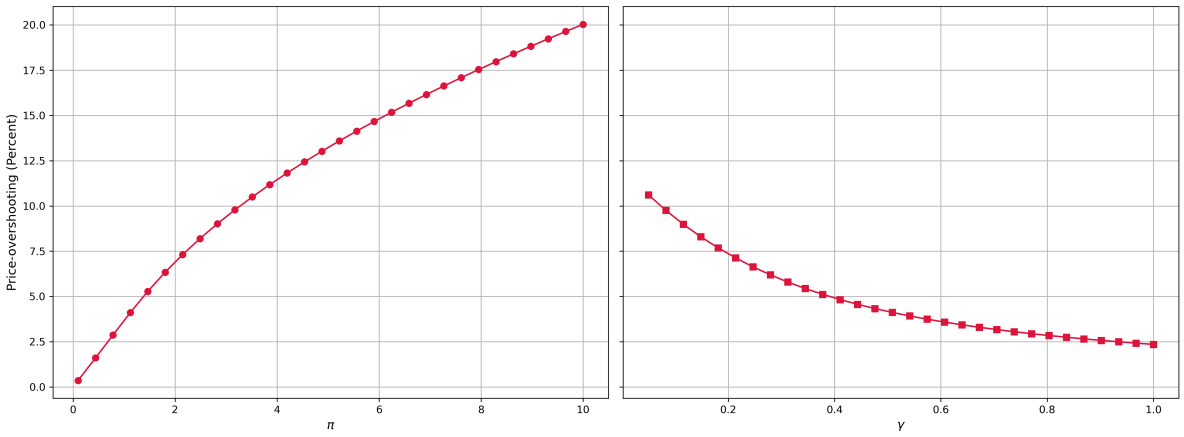
Notes: The figure plots value functions and markup density for both baseline economy and the case with markup cap at the static optimum. The vertical line represents the static optimal markup.

7 depicts the value function in the baseline economy. It is monotonically increasing and exhibits local concavity at the static optimal markup, indicated by the vertical line. The value function attains its maximum at μ_0 and smoothly approaches zero as μ nears μ_b.

The right panel of Figure 7 illustrates the distribution of markup densities. The density is highest around the reset markup μ_0 and declines as the markup decreases. Near the exit threshold μ_b, the density is low. However, because the markup distribution is not tightly concentrated around μ_0, the resulting misallocation is non-negligible. We estimate that dynamic pricing frictions alone generate a total factor productivity (TFP) loss of 2%. This estimate is comparable in magnitude to that of Edmond et al. (2022).[13] For the same reason, despite large price premium, the aggregate markup in the baseline economy is about 1.19—close to the static optimum. Additionally, aggregate R&D expenditures plus fixed cost of keeping research capacity represent around 3% of nominal GDP, which is roughly consistent with the R&D expenditure share of GDP in the US before 2020.

[13]Edmond et al. (2022) connect the markup distribution to underlying productivity heterogeneity. In their framework, high-productivity firms charge higher markups and command larger market shares. In contrast, our model assumes same productivity across firms. Hence, the markup distribution arises solely from dynamic pricing frictions, and firms with higher markups have lower market shares in our model.

Figure 8: Price-Overshooting over Inflation and Pricing Friction



For ease of comparison with the counterfactual experiments presented in the next section, we normalize welfare in the baseline to 100.

Non-targeted Moments. We evaluate the model's performance by comparing its implications to a set of non-targeted moments—economically relevant quantities *not* used in estimation. This serves as an out-of-sample test of the model implications and estimated parameters.

First, we simulate thousands of firms and use the same empirical methods as in Section 3 and compare the empirical results in both datasets. Note that in the real data, products may exit due to various reasons that our model does not capture, which introduces different patterns in the dynamics. In particular, the price-overshooting in the real data will be 15% but this is not comparing the prices of the same product over the life cycle. Our estimation instead uses price dynamics of the same product, which results in the difference of the initial responses of prices. Nevertheless, Figure A6 in Appendix shows that the price dynamics is very similar to both datasets, despite the model parsimony.

Second, Figure 8 shows the comparative statics of price-overshooting over inflation and pricing friction. Consistent with the model, Table A1 in Appendix shows that the relationship between inflation and price-overshooting is positive and significant in the data. Moreover, consistent with Table 2, the model exhibits a negative association between pricing friction and the magnitude of price-overshooting.

The Role of Price-Overshooting. In this section, we show that price-overshooting significantly increases number of available varieties and growth rate. We keep the calibrated parameters fixed at baseline and only allow firms to reset prices to the static optimal markup.

Table 6: The Role of Price-Overshooting

x^I	x^E	x^N	μ_u	μ_b	N	χ_c	v	m	g	Life	Wel
Panel A: Frictionless Benchmark											
8.086%	0.995%	0.227%	1.200	–	1.000	0.913	100.00%	1.200	1.733%	7.2	105.8
Panel B: Baseline											
7.980%	0.982%	0.224%	1.318	0.977	0.895	0.904	98.08%	1.190	1.710%	6.5	100.0
Panel C: No Overshooting											
6.427%	0.791%	0.180%	1.200	0.978	0.717	0.929	99.02%	1.106	1.377%	6.0	93.0

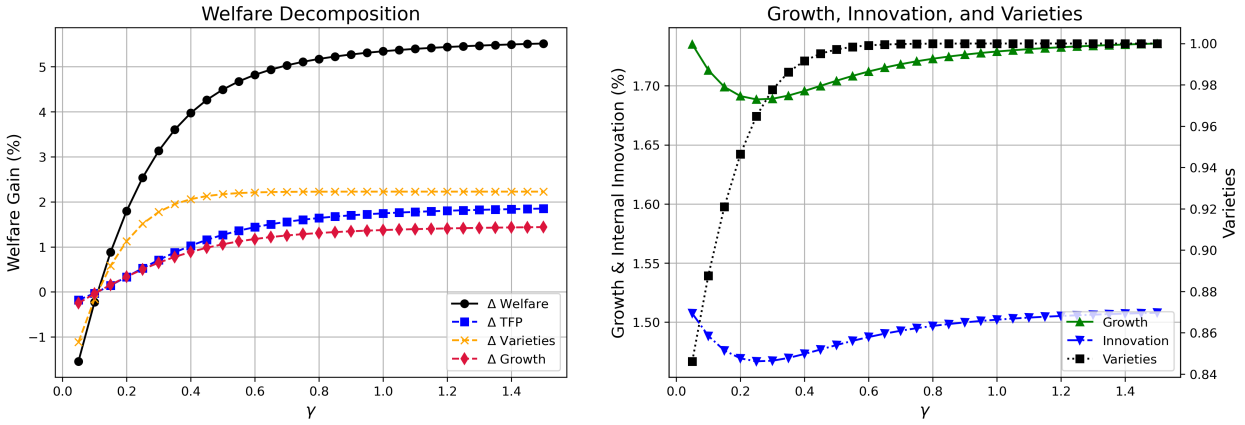
The red dashed lines in Figure 7 illustrate the value function and markup density in a counterfactual scenario where prices are not allowed to overshoot. In this environment, the exit threshold μ_b remains largely unchanged, as firms are still unwilling to tolerate more than one year of negative profits. However, the elimination of price overshooting also leads to a reduction in the number of available product varieties. Specifically, the distance between μ_0 and μ_b shrinks, so a product without price resetting and churning by innovation exits after approximately 8.4 years instead of 12.9 years, as in the baseline. This shortened horizon reduces the steady-state mass of varieties and expected product lifetime. Consistent with this, the markup density rises sharply near μ_b, indicating a higher frequency of endogenous product exit.

The maximal value attainable through innovation or price resetting declines significantly relative to the baseline, leading to lower innovation incentives. Table 6 shows that both the incumbent and entrant innovation rates, x^I, x^E and x^N, fall, which leads to a significant decline of growth rate by 33 bpts. Two opposing forces shape the value of a product and incumbent innovation rate in this setting. On one hand, banning overshooting limits firms' ability to smooth prices around the static optimal markup, thereby reducing the continuation value of products, and thus increasing innovation rate. On the other hand, a smaller number of available varieties could in principle raise the return to innovation due to weaker competition on the product market as hinted in firm's profit shown in (108). In this case, however, the former effect dominates. In general, the innovation rates are non-monotone in pricing frictions γ, as we will show in the next section.

The consumption-equivalent welfare declines by 7 ppts (7.24 log points) in this counterfactual, implying that the household would be willing to give up 7% of current consumption to avoid the no-overshooting scenario. We decompose the welfare change into three components as in (59):

$$\Delta \log \text{Welfare} = \underbrace{-3.77}_{\text{growth-related}} + \underbrace{0.95}_{\text{misallocation}} + \underbrace{-4.41}_{\text{variety}}. \tag{60}$$

Figure 9: Welfare and Growth vs. Pricing Friction



Notes: The left panel plots the welfare change compared to the baseline welfare in the range of $\gamma = 0.1$ to $\gamma = 1.5$, and decomposition into misallocation, varieties and growth. The right panel plots the growth rate and its decomposition into innovation rate and number of varieties.

Roughly 52% of the welfare loss is attributable to slower growth, while about 61% is driven by the decline in available varieties. By contrast, aggregate efficiency increases by about 13% due to reduced misallocation associated with compressed markup dispersion.

Panel A of Table 6 reports the frictionless benchmark. Relative to it, pricing frictions in the baseline model reduce welfare by 5.8 ppts (5.6 log points) and lower long-run growth by 2.3 bpts. Price overshooting partially offsets these losses by raising overall profits over the product life time and therefore product values and innovation incentives. Without overshooting, welfare falls by 13 ppts and growth drops by 36 bpts (about 2 times and 15 times larger declines, respectively). This highlights the key role of overshooting in mitigating the welfare and growth costs of pricing frictions.

Comparative Statics on Pricing Friction. Finally, we study comparative statics of welfare, growth, varieties, and innovation with respect to the pricing-friction parameter.

Figure 9 plots consumption-equivalent welfare as the price adjustment probability γ varies, relative to the baseline. Welfare rises through three channels: higher variety, higher aggregate efficiency (TFP), and faster growth.

Near the baseline, the variety channel dominates. The growth contribution is muted (and even flat) because stronger competition from additional varieties reduces incumbent innovation (see the right panel). As γ increases further, the variety margin quickly saturates. Lower pricing frictions then raise product values by increasing expected lifetime profits, which strengthens innovation incentives and raises long-run growth. Higher overall innovation also compresses the markup distribution, reducing misallocation and increasing TFP; correspondingly, the growth and TFP

Table 7: Optimal R&D Subsidies

x^I	x^E	x^N	μ_u	μ_b	N	χ_c	v	m	g	Life	Wel
Panel B: Baseline											
7.980%	0.982%	0.224%	1.318	0.977	0.895	0.904	98.08%	1.190	1.710%	6.5	100.0
Panel B: Incumbent R&D Subsidy Only (41%)											
10.881%	1.339%	0.181%	1.298	0.981	0.917	0.833	98.37%	1.184	2.309%	5.4	102.6
Panel C: Entrant R&D Subsidy Only (23%)											
7.934%	0.977%	0.289%	1.317	0.977	0.895	0.904	98.08%	1.190	1.713%	6.5	100.02

curves move together. Accordingly, the right panel exhibits a non-monotone pattern: growth and internal innovation rates first decline with stronger competition, then increases with higher product value.

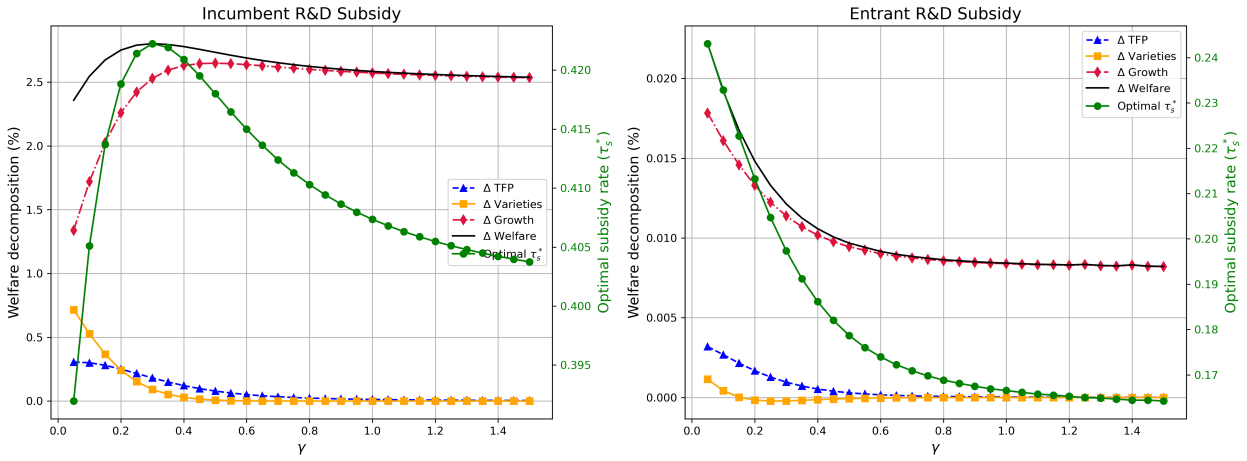
6 Optimal R&D Policy

In this section, we explore how different policy instruments, for example incumbent R&D subsidies and entry subsidies, affect the equilibrium outcomes given the parameters fixed at our baseline values.

We begin by analyzing incumbent R&D subsidies. Panel B of Table 7 presents the equilibrium outcomes when only incumbent R&D is subsidized. The optimal subsidy rate is 41%. As expected, the incumbent innovation rate rises from 12.73% to 17.17%, a 5 percentage point increase. While the reset and exit markups remain largely unchanged, the number of varieties increases by 2%, and aggregate efficiency improves by 0.34%. This improvement is driven by a higher innovation rate, which increases the fraction of products eligible for price resetting in each period, thereby concentrating the markup distribution more tightly around μ_0. Overall, welfare rises by 2.56 ppts, with higher growth contributing 1.69 ppts, an expansion in product variety contributing 0.53 ppts, and the remainder stemming from improved aggregate efficiency.

Our findings contrast with those of Acemoglu et al. (2018), who conclude that incumbent R&D subsidies have only limited welfare effects. The key distinction again lies in the underlying data. Acemoglu et al. (2018) focus on manufacturing firms, whereas we examine retail industries. In the manufacturing sector during the 1990s, entry rates were substantially higher—Acemoglu et al. (2018) report a five-year entrant share of 40%—compared to an annual entry rate of just 2.23% in our data. This stark difference implies that the increase in the creative destruction rate τ, induced

Figure 10: Optimal Incumbent and Entrant R&D Policies over γ



Notes: The left panel plots the welfare change after optimal incumbent R&D compared to the welfare in non-subsidized equilibrium in the range of $\gamma = 0.1$ to $\gamma = 1.5$, and decomposition into misallocation, varieties and growth. The right panel plots the growth rate and its decomposition into innovation rate and number of varieties. The right panel plots the welfare change after R&D subsidy on entrant R&D and its decomposition.

by higher incumbent innovation, reduces entry by only 0.44% in our context, with minimal impact on overall growth. However, in their context, higher incumbent innovation discourages entry.

We next consider entrant R&D subsidies. Panel C of Table 7 shows the outcomes when only entrant R&D is subsidized. The optimal subsidy rate in this case is 19%. The entrant innovation rate rises from 2.17% to 2.61%, but this crowds out incumbent innovation, resulting in a mere 0.02% increase in the growth rate and a modest welfare gain of 0.06 ppts. Again, since entrants represent a relatively small share of economic activity in the baseline, the impact on equilibrium outcomes is limited.

Figure 10 illustrates how optimal incumbent and entrant R&D subsidies, along with the resulting welfare gains, vary with the degree of pricing friction. The left panel presents the welfare improvement and its decomposition under optimal incumbent R&D subsidies, benchmarked against the non-subsidized equilibrium at each value of γ. Two key insights emerge. First, both the optimal incumbent subsidy rate and the resulting welfare gain are larger when pricing friction is high and decline as it decreases. This implies that R&D subsidies are dependent on pricing frictions and the policy help alleviate negative impact of pricing frictions. Second, under severe pricing friction, the welfare gains from increased variety and improved aggregate efficiency are comparable to those from higher growth. In particular, the welfare contribution from growth exhibits an inverted-U pattern.

The intuition is as follows. When pricing friction is large, the number of available varieties is low. An increase in the incumbent innovation rate leads to smaller changes in the creative destruction

rate and hence in growth, explaining the limited welfare contribution from the growth margin. However, the same increase in innovation helps overcome pricing frictions by introducing more new products. This compresses the markup distribution toward μ_0, thereby improving aggregate efficiency and also raises the number of varieties. These gains in turn justify a larger optimal R&D (45% when $\gamma = 0.05$) subsidy even when growth gain is lower when pricing friction is high. As pricing frictions decline, the gains from variety and efficiency diminish significantly. In the frictionless benchmark, only the gain from growth remains, yielding an optimal subsidy rate of 38% and a welfare gain of 1.25 ppts.

The right panel of Figure 10 displays the case for optimal entrant R&D subsidies. Here, both the optimal subsidy rate follows an inverted-U shape. When pricing frictions are high and variety is limited, an increase in entry crowds out incumbent innovation, thereby reducing growth gains. Interestingly, the induced improvement in aggregate efficiency due to the increase in creative destruction rate contributes more to welfare than the growth effect. A key distinction from the incumbent subsidy case is that the number of varieties declines slightly. This is because a higher creative destruction rate raises the probability of product exit, lowering the value of price overshooting and the incentive to wait through periods of negative profit. This narrows the gap between the reset and exit markups. In this case, this effect dominates the effect of slightly more concentrated markup distribution, leading to shrinking varieties. Therefore, the subsidy rate is lower when pricing friction is extremely large.

Taken together, these counterfactual R&D subsidy experiments yield two primary implications: (i) in the retail sector, subsidizing incumbent R&D generates substantially greater welfare gains than subsidizing entrants; and (ii) in sectors with greater pricing frictions, optimal R&D subsidies should be higher to stimulate gains through both variety expansion and improved efficiency.

7 Conclusion

This paper revisits the conventional wisdom that pricing frictions are irrelevant for long-run economic outcomes. Using new product–firm evidence, we show that pricing frictions operate over the entire product life cycle rather than at business-cycle frequencies. Incumbent prices remain nearly flat in nominal terms and decline steadily in real terms, while new products enter with sizable price premia. These patterns reflect a slow-moving form of price rigidity distinct from temporary sales or high-frequency adjustments emphasized in prior work. We further show that sectors with greater price rigidity exhibit larger entry premia and that tariff-induced cost shocks pass through to the prices of new products, but not incumbents—direct evidence that long-run pricing frictions shape the margins through which firms adjust.

We develop an endogenous growth model with trend inflation and infrequent price adjustment that rationalizes these facts. Firms anticipate that they will be unable to adjust prices over a product's life cycle and therefore set elevated introductory markups—price overshooting—to insure against future erosion of real prices. Pricing frictions create persistent dispersion in markups across otherwise identical products, generating misallocation and depressing innovation incentives. At the same time, the option to reset prices when launching new products partly offsets these losses by raising the value of innovating.

Quantitatively, pricing frictions have sizable effects: they reduce growth, variety, and welfare, even when firms optimally overshoot prices. Removing overshooting leads to markedly lower innovation, higher exit, and significantly larger welfare losses, underscoring its central role as a margin of adjustment. Optimal-policy analysis reveals that R&D subsidies to incumbents are particularly effective because they accelerate product churning, compress the markup distribution, and alleviate misallocation—especially in high-friction sectors.

Taken together, our findings call for a reorientation of how nominal rigidities enter models of endogenous growth. Long-run pricing frictions, not just short-run stickiness, shape firms' innovation incentives, the dynamics of product turnover, and the consequences of trade shocks and stabilization policy. Understanding the interaction between inflation, pricing frictions, and innovation is therefore essential for evaluating the long-run effects of monetary and regulatory interventions.

Bibliography

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr, "Innovation, Reallocation, and Growth," *American Economic Review*, November 2018, *108* (11), 3450–3491.

Adam, Klaus and Henning Weber, "Estimating the optimal inflation target from trends in relative prices," *American Economic Journal: Macroeconomics*, 2023, *15* (3), 1–42.

Aghion, Philippe and Peter Howitt, "A Model of Growth through Creative Destruction," *Econometrica*, 1992, *60* (2), 323–351.

_ , Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li, "A Theory of Falling Growth and Rising Rents," *The Review of Economic Studies*, 2023, *90* (6), 2675–2702. Published online February 20, 2023.

Akcigit, Ufuk and William R. Kerr, "Growth through Heterogeneous Innovations," *Journal of Political Economy*, 2018, *126* (4), 1374–1443.

Anderson, Eric, Nir Jaimovich, and Duncan Simester, "Price stickiness: Empirical evidence of the menu cost channel," *Review of Economics and Statistics*, 2015, *97* (4), 813–826.

Argente, David and Chen Yeh, "Product life cycle, learning, and nominal shocks," *The Review of Economic Studies*, 2022, *89* (6), 2992–3054.

_ and Munseob Lee, "Cost of Living Inequality During the Great Recession," *Journal of the European Economic Association*, 06 2020, *19* (2), 913–952.

_ , _ , and Sara Moreira, "The life cycle of products: Evidence and implications," *Journal of Political Economy*, 2024, *132* (2), 337–390.

Bai, Liang and Sebastian Stumpner, "Estimating US Consumer Gains from Chinese Imports," *American Economic Review: Insights*, 2019, *1* (2), 209–224.

Beraja, Martin, Erik Hurst, and Juan Ospina, "The Aggregate Implications of Regional Business Cycles," *Econometrica*, 2019, *87* (6), 1789–1833.

Bils, Mark and Peter J. Klenow, "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 2004, *112* (5), 947–985.

Blanco, Andrés, Corina Boar, Callum Jones, and Virgiliu Midrigan, "Nonlinear Inflation Dynamics in Menu Cost Economies," Working Paper 2024-10, Federal Reserve Bank of Atlanta September 2024. FRB Atlanta Working Paper No. 2024-10.

Cafarella, Michael, Gabriel Ehrlich, Tian Gao, John C Haltiwanger, Matthew D Shapiro, and Laura Zhao, "Using machine learning to construct hedonic price indices," Technical Report, National Bureau of Economic Research 2023.

Calvo, Guillermo A., "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 1983, *12* (3), 383–398.

Campbell, Jeffrey R and Benjamin Eden, "Rigid prices: Evidence from us scanner data," *International*

Economic Review, 2014, *55* (2), 423–442.

Cavallo, Alberto, Gita Gopinath, Brent Neiman, and Jenny Tang, "Tariffs Pass Through as a Domestic Tax: Evidence from U.S. Trade Policy," *American Economic Review*, 2021, *111* (2), 686–719.

Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong, "The cyclicality of sales, regular and effective prices: Business cycle and policy implications," *American Economic Review*, 2015, *105* (3), 993–1029.

Deaton, Angus, *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*, Baltimore: Johns Hopkins University Press for the World Bank, 1997.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu, "How Costly Are Markups?," Working Paper 30753, National Bureau of Economic Research 2022. NBER Working Paper No. 30753.

Eichenbaum, Martin, Nir Jaimovich, and Sergio Rebelo, "Reference prices, costs, and nominal rigidities," *American Economic Review*, 2011, *101* (1), 234–262.

Fajgelbaum, Pablo D, Pinelopi K Goldberg, Patrick J Kennedy, and Amit K Khandelwal, "The Return to Protectionism," *Quarterly Journal of Economics*, 2020, *135* (1), 1–55.

Friedman, Milton, "The Role of Monetary Policy," *The American Economic Review*, 1968, *58* (1), 1–17.

Galí, Jordi, *Monetary Policy, Inflation, and the Business Cycle*, Princeton University Press, 2015.

Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J. Klenow, "How Destructive is Innovation?," *Econometrica*, September 2019, *87* (5), 1507–1541.

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

Grossman, Gene M. and Elhanan Helpman, *Innovation and Growth in the Global Economy*, Cambridge, MA: MIT Press, 1991.

Jones, Charles I., "The AI Dilemma: Growth versus Existential Risk," *AER: Insights*, 2024, *6* (4), 575–590.

Kaplan, Greg and Guido Menzio, "The morphology of price dispersion," *International Economic Review*, 2015, *56* (4), 1165–1206.

Kehoe, Patrick and Virgiliu Midrigan, "Prices are sticky after all," *Journal of Monetary Economics*, 2015, *75*, 35–53.

Klenow, Peter J and Oleksiy Kryvtsov, "State-dependent or time-dependent pricing: Does it matter for recent US inflation?," *The Quarterly Journal of Economics*, 2008, *123* (3), 863–904.

Klette, Tor Jakob and Samuel Kortum, "Innovating firms and aggregate innovation," *Journal of Political Economy*, 2004, *112* (5), 986–1018.

Liu, Xiaojie, "Consumer Information Friction and Monetary Non-Neutrality: A Sequential Search Approach," 2025. Available at SSRN: https://ssrn.com/abstract=5271515 or

http://dx.doi.org/10.2139/ssrn.5271515.

Nakamura, Emi and Jón Steinsson, "Five Facts about Prices: A Reevaluation of Menu Cost Models," *The Quarterly Journal of Economics*, 2008, *123* (4), 1415–1464.

Peters, Michael, "Heterogeneous markups, growth, and endogenous misallocation," *Econometrica*, 2020, *88* (5), 2037–2073.

Rebelo, Sergio, Miguel Santana, and Pedro Teles, "Behavioral sticky prices," *Journal of Monetary Economics*, 2025, p. 103828.

Romer, Paul M., "Endogenous Technological Change," *Journal of Political Economy*, October 1990, *98* (5, Part 2), S71–S102.

Werning, Iván, "Expectations and the Rate of Inflation," Working Paper 30260, National Bureau of Economic Research 2022. NBER Working Paper No. 30260.

Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, 2003.

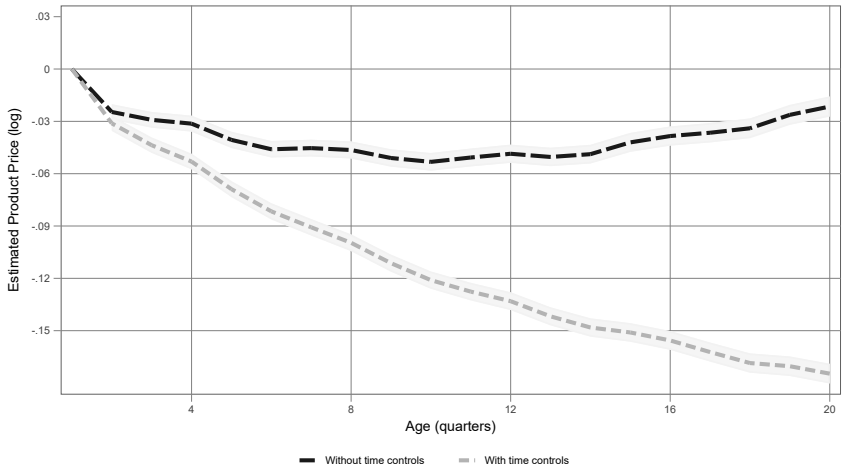
Innovation and Pricing Frictions

Appendix
For Online Publication

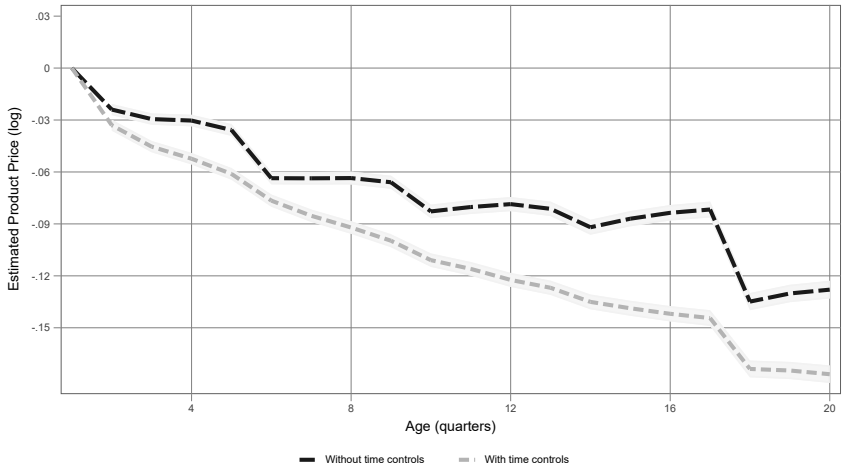
A Additional Results

Figure A1: Prices over the Product Life Cycle: Age effects for Semi-Balanced and Unbalanced Samples

(a) Semi-balanced Sample



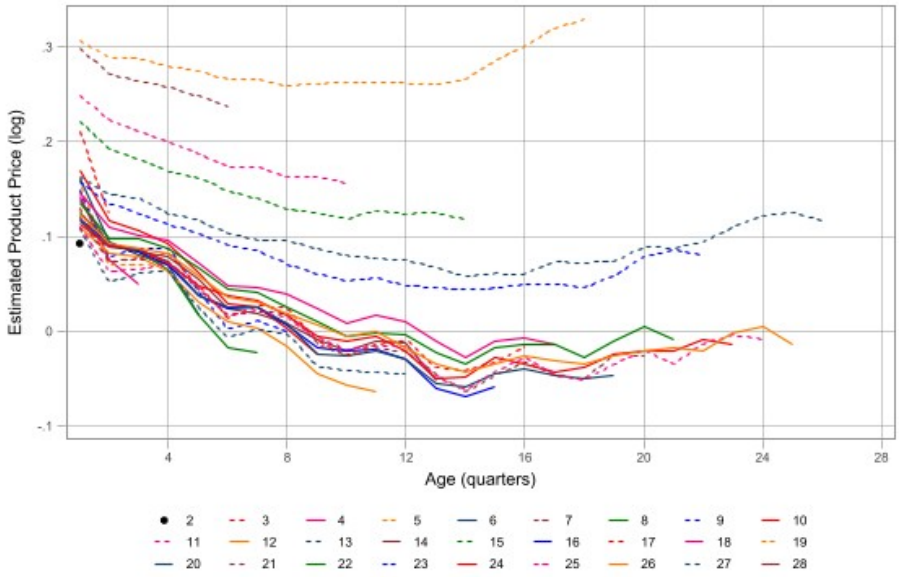
(b) Unbalanced Sample



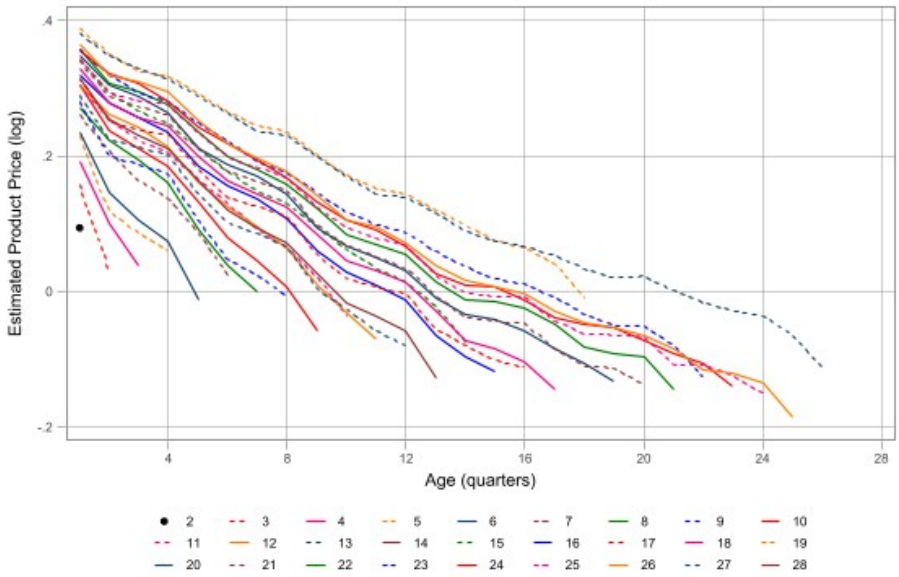
Notes: The figure shows the estimated age fixed effects of (log) prices of products computed using equations (3) and (4). The estimation uses all products. In panel (a) the age fixed effects are computed for cohorts 2006q3-2016q3, including products that exit before they complete 20 quarters of sales. In panel (b) the age fixed effects includes products from cohorts 2006q3 to 2021q4. The gray area indicates the 95% confidence interval.

Figure A2: Prices over the Product Life Cycle: by Duration

Panel (a) - Without time controls

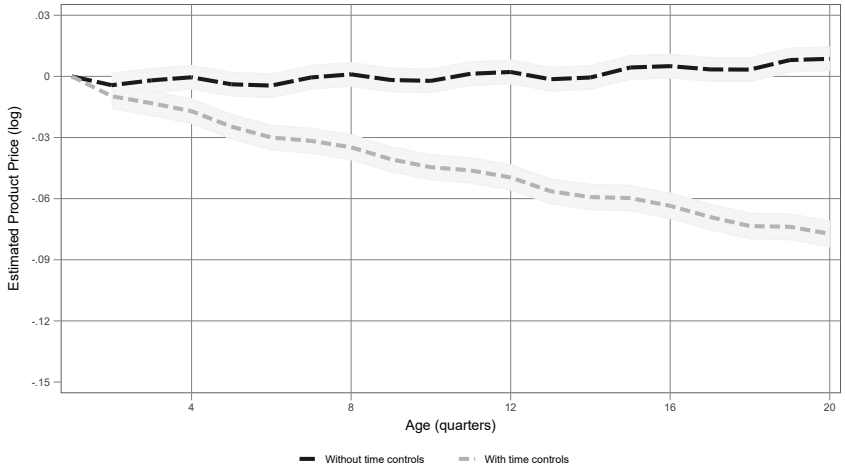


Panel (b) - With time controls



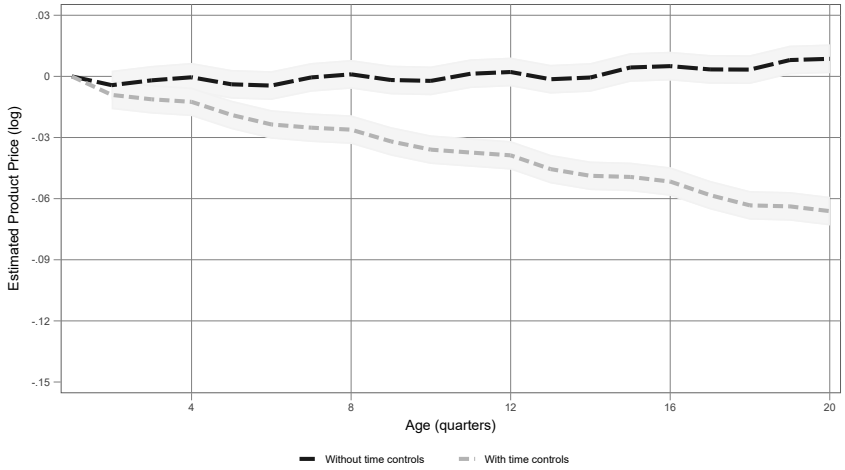
Notes: The figure shows the estimated age fixed effects of (log) prices of products computed using equations (3) and (4) where age effects are interacted with duration, for products that lasted between 2 and 28 quarters in the market.

Figure A3: Prices over the Product Life Cycle: Alternative Sample



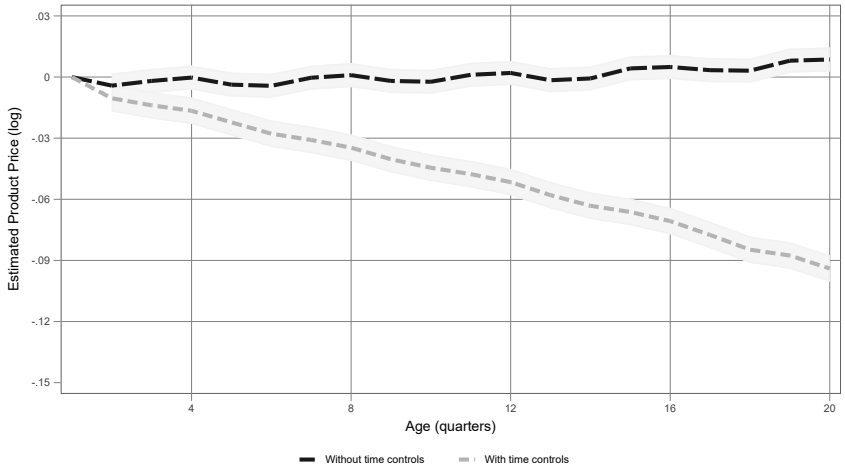
Notes: The figure shows the estimated age fixed effects of (log) prices of products computed using equations (3) and (4). The estimation uses all products and the age fixed effects are computed for products duration of more than 20 quarters. The gray area indicates the 95% confidence interval.

Figure A4: Prices over the Product Life Cycle: Alternative Cohort Effects



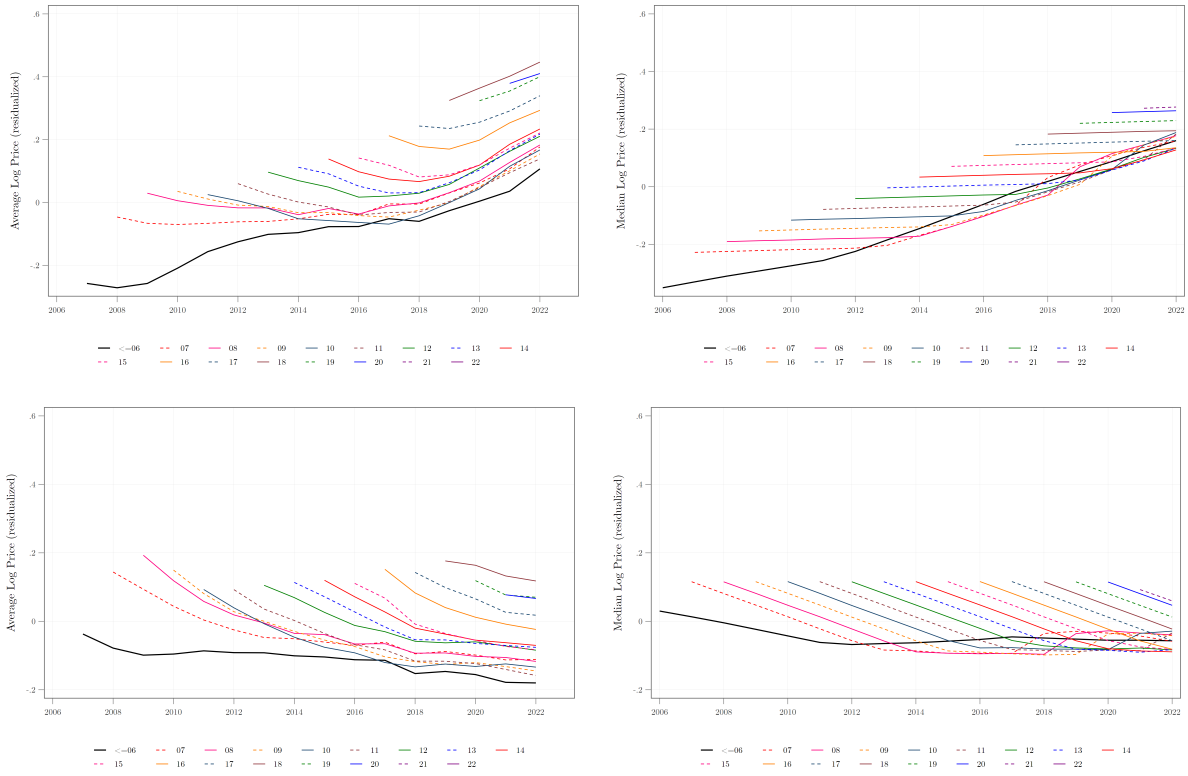
Notes: The figure shows the estimated age fixed effects of (log) prices of products computed using equations (3) and (4). The estimation uses all products and the age fixed effects are computed for products duration of more than 20 quarters. The gray area indicates the 95% confidence interval.

Figure A5: Prices over the Product Life Cycle: Firm Controls



Notes: The figure shows the estimated age fixed effects of (log) prices of products computed using equations (3) and (4). The estimation uses all products and the age fixed effects are computed for products duration of more than 20 quarters. The gray area indicates the 95% confidence interval.

Figure A6: Price Dynamics Fit between simulated data and real data



Notes: **The left panel plots the price dynamics that is estimated in the real data. The right panel plots the price dynamics that is estimated in the simulated data.**

Table A1: Inflation and Price Overshooting of New Products

	(1)	(2)	(3)	(4)	(5)	(6)
	Relative Price New Products					
Inflation	1.462***	1.472***	1.314***	1.325***	1.275***	1.180***
	(0.022)	(0.023)	(0.023)	(0.023)	(0.021)	(0.077)
Observations	58,201	58,201	58,201	58,201	58,201	7,052
R-squared	0.068	0.073	0.182	0.184	0.303	0.506
Time FE	N	Y	Y	Y	Y	Y
Sector FE	N	N	Y	Y	Y	Y
Controls	N	N	N	Y	Y	Y
Weights	N	N	N	N	Y	Y
Period	07q1-22q4	07q1-22q4	07q1-22q4	07q1-22q4	07q1-22q4	21q1-22q4

Notes: The table presents the association between inflation and price overshooting of new products using sector (module) x time variation. Inflation and price overshooting are defined as $\Delta_{m,t} = \frac{P_{m,t} - P_{m,t-4}}{P_{m,t-4}}$ and $PO_{m,t} = \frac{P^{new}_{m,t} - p^{incumbents}_{m,t}}{p^{incumbents}_{m,t}}$, respectively. Controls includes product introduction rate of incumbent firms. Weights are revenue weights.

Table A2: Relative Price New Products and Tariff Exposure

	Relative Price New Products					
	(1)	(2)	(3)	(4)	(5)	(6)
Tariff Exposure	1.3644**	1.3603**	1.5102**			
	(XXX)	(XXX)	(XXX)			
Change Tariff Exposure				1.1827*	1.1721*	1.7296**
				(XXX)	(XXX)	(XXX)
time × category	yes	yes	yes	yes	yes	yes
firm × category	yes	yes	yes	yes	yes	yes
controls	no	yes	yes	no	yes	yes

Notes: The table reports the estimated coefficients of tariff exposure from equation 6 with $k = 2$, using firm–module–year variation for the period 2010–2022. Price overshooting is defined as $\tilde{P}_{i,j,t+k} = \left(P^{\text{new}}_{i,j,t+k} / P^{\text{inc}}_{i,j,t+k}\right) - 1$. Categories correspond to NielsenIQ product modules. Columns (1) and (4) include time–category and firm–category fixed effects. Columns (2) and (5) additionally include controls for the quantity coverage of the tariff exposure measure. Columns (3) and (6) further include the interaction of coverage and tariff variables; the reported coefficients correspond to the estimated effect of tariff exposure evaluated at the average coverage level.

Table A3: Tariff Exposure and Relative Prices of New Products

	k						
	-2	-1	0	1	2	3	4
Tariff Exposure	0.4823	-0.4651	0.2609	1.5726*	2.7002***	3.4068***	-1.1363
	(1.05)	(0.88)	(0.79)	(0.92)	(1.02)	(1.17)	(1.28)
Tariff Exposure ×	-0.7557	0.7307	-0.1452	-1.5020**	-1.7996**	-2.2246**	1.0332
Frequency Price Changes	(0.867)	(0.662)	(0.581)	(0.727)	(0.767)	(0.887)	(1.001)
Observations	51463	60210	66820	59182	49463	42222	35503

Notes: The table shows the estimated coefficients of equation $y_{i,j,t+k} = \beta T_{i,j,t} + \gamma T_{i,j,t} \times Z_j + \omega_{j,t+k} + \phi_{i,j} + \text{controls}_{i,j,t+k} + \varepsilon_{i,j,t+k}$, $k = -2, ..., 0, ..., 4$. The variable $T_{i,j,t}$ is the level of tariff exposure. The outcome is the relative price of new products $\tilde{P}_{i,j,t+k} = \left(P^{\text{new}}_{i,j,t+k} / P^{\text{inc}}_{i,j,t+k}\right) - 1$. The variable Z_j is the standardized frequency of price changes at the product category level. The vertical bands represent +/- 1.96 × st. error of each point estimate.

B Robustness for Alternative Samples

B1. Alternative Samples

We evaluate four sets of measurement problems: (1) entry and exit of stores; (2) private label products; (3) incomplete spells; (4) coverage of some product categories.

(1) Entry and Exit of Stores

The baseline sample includes products sold in all stores. The number of store available in the dataset every year varies, and this raises two measurement issues. First, estimates of products' entries and exits might be affected by the entries and exits of stores in the sample. Second, estimates of prices over the life cycle can also vary with the entry and exit of stores with varying price levels (Kaplan and Menzio, 2015). Therefore, for robustness of the main results, we consider a balanced sample of stores during our sample period 2006-2022. Table B3 presents the descriptive statistics of this alternative sample. Figures B1 and B2 present the main results.

Table B1: Summary Statistics: Robustness Entry and Exit Stores – E

	All	By Censoring Type			
		Not Censored	Right	Left	Right & Left
Total # of products	675,455	308,383	203,961	116,559	46,552
Duration (quarters)					
average	18	9.7	20	18	68
less than 4	31	44	21	28	0
above 20	27	11	32	31	100
above 28	20	5.4	23	21	100
Revenue (quarterly, $1,000)					
mean	94	30	158	27	294
25th percentile	0.6	0.3	1.6	0.2	5.6
median	4.9	2.6	13	1.4	32
75th percentile	36	16	93	10	183
90th percentile	177	66	365	51	694
95th percentile	404	136	719	119	1,292
Price (quarterly)					
mean	0.047	0.039	0.12	-0.061	0.026
25th percentile	-0.35	-0.39	-0.26	-0.44	-0.32
median	0.079	0.088	0.14	-0.0022	0.0081
75th percentile	0.50	0.53	0.54	0.37	0.37
90th percentile	0.93	0.98	0.96	0.79	0.79
95th percentile	1.20	1.30	1.20	1.10	1.10

Notes: The table presents summary statistics for the products included in the baseline sample for 2006q1–2022q4. Products already active in 2006q1–2006q2 are left-censored, and products with sales in 2022q3–2022q4 are right-censored. "Not Censored" includes products observed from entry to exit. "Right Censored" indicates products still sold at the end of the sample. "Left Censored" refers to products already on the market at the start of the sample. "Right & Left" includes products with both entry and exit censored. Duration is measured in quarters. Revenue and price are calculated per quarter. Revenue is in thousands of dollars. Prices are differenced from the sector×quarter median price.

Table B2: Summary Statistics: Robustness Entry and Exit Stores – Full

	All	By Censoring Type			
		Not Censored	Right	Left	Right & Left
Total # of products	1,789,877	937,703	372,209	395,221	84,743
Duration (quarters)					
average	21	15	22	25	68
less than 4	22	28	18	14	0
above 20	37	26	39	48	100
above 28	26	15	30	35	100
Revenue (quarterly, $1,000)					
mean	51	23	105	23	207
25th percentile	0.2	0.1	0.7	0	2.1
median	1.6	1.2	5.6	0.5	14
75th percentile	14	9.4	45	5	95
90th percentile	79	44	217	35	444
95th percentile	199	98	478	92	930
Price (quarterly)					
mean	0.097	0.14	0.17	-0.059	0.022
25th percentile	-0.36	-0.34	-0.25	-0.51	-0.36
median	0.13	0.18	0.18	-0.0014	0.015
75th percentile	0.62	0.69	0.64	0.43	0.42
90th percentile	1.10	1.20	1.10	0.90	0.85
95th percentile	1.40	1.50	1.40	1.20	1.10

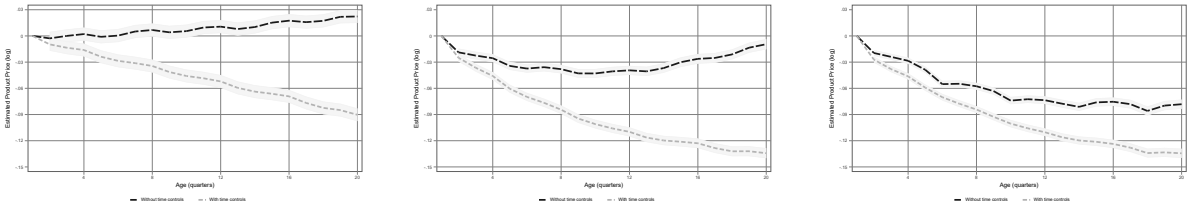
Notes: Same definitions as in Table **??**. Duration is in quarters. Revenue and price are quarterly; revenue is in thousands of dollars; prices are differenced from the sector×quarter median price.

Table B3: Summary Statistics: Robustness Entry and Exit Stores – Full

	All	By Censoring Type			
		Not Censored	Right	Left	Right & Left
Total # of products	1,789,877	937,703	372,209	395,221	84,743
Duration (quarters)					
average	21	15	22	25	68
less than 4	22	28	18	14	0
above 20	37	26	39	48	100
above 28	26	15	30	35	100
Revenue (quarterly, $1,000)					
mean	51	23	105	23	207
25th percentile	0.2	0.1	0.7	0.0	2.1
median	1.6	1.2	5.6	0.5	14
75th percentile	14	9.4	45	5	95
90th percentile	79	44	217	35	444

Figure B1: Price Life Cycle: Robustness Entry and Exit of Stores

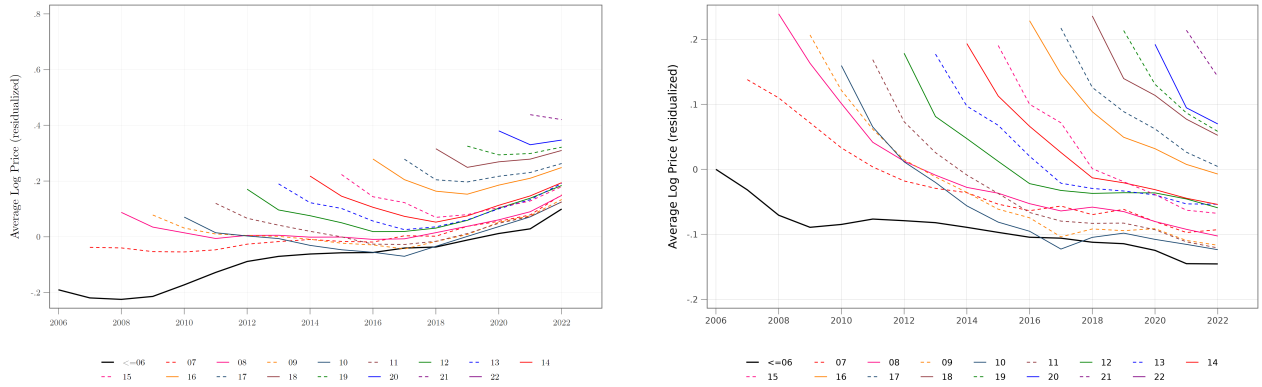
(a) Balanced cohorts w/o exit (b) Balanced cohorts w/ exit (c) Unbalanced cohorts w/exit



Notes: The figure shows the estimated age fixed effects of (log) prices of products computed using equations (3) and (4). We use the sector and sector x time effects, as well as iteration of sector and sector x quarter fixed effects with firm-level fixed effects. The estimation uses all branded products sold in the balanced set of stores active over the entire period 2006-2022. In (a) the age fixed effects are computed for balanced cohort and products duration of more than 20 quarters. In (b) the age fixed effects are computed for balanced cohorts but individual products can exit before 20 quarters of activity. In (c) we use an unbalanced sample of cohorts with exit. Main text provides details on the cohorts used.

Figure B2: Prices by Cohort of Products: Robustness Entry and Exit of Stores

(a) Without Time Controls (b) With Time Controls



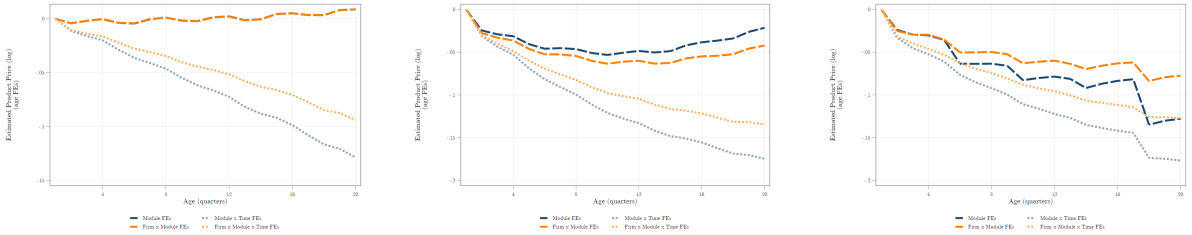
Notes: The figures display average log prices by cohort over time using the baseline annual-level sample. In panel (a), log prices are residualized using sector fixed effects. In panel (b), log prices are residualized using sector-by-year fixed effects. The estimation uses all branded products sold in the balanced set of stores active over the entire period 2006-2022.

(2) Private Label Products

The baseline sample excludes private label products. In order to protect the identity of the retailer, Nielsen alters the UPCs associated with private label goods. As a result, multiple private label items are mapped to a single UPC that makes it difficult to interpret the entry and exit patterns of these items since it is not possible to determine the producer of these goods. We evaluated the impact of this exclusion by having including these products in robustness samples. Figures B3 and B4 present the main results.

Figure B3: Price Life Cycle: Robustness Private Label

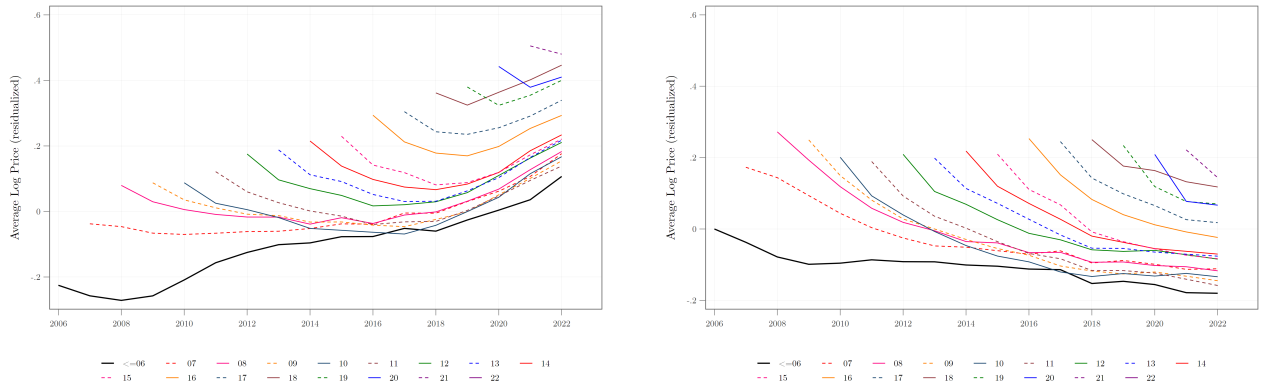
(a) Balanced cohorts w/o exit (b) Balanced cohorts w/ exit (c) Unbalanced cohorts w/exit



Notes: The figure shows the estimated age fixed effects of (log) prices of products computed using equations (3) and (4). We use the sector and sector x time effects, as well as iteration of sector and sector x quarter fixed effects with firm-level fixed effects. The estimation uses all branded products and private label products. In (a) the age fixed effects are computed for balanced cohort and products duration of more than 20 quarters. In (b) the age fixed effects are computed for balanced cohorts but individual products can exit before 20 quarters of activity. In (c) we use an unbalanced sample of cohorts with exit. Main text provides details on the cohorts used.

Figure B4: Prices by Cohort of Products: Robustness Private Label

(a) Without Time Controls (b) With Time Controls



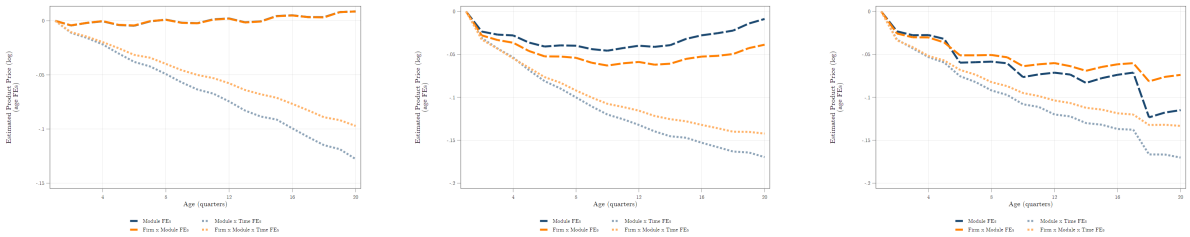
Notes: The figures display average log prices by cohort over time using the baseline annual-level sample. In panel (a), log prices are residualized using sector fixed effects. In panel (b), log prices are residualized using sector-by-year fixed effects. The estimation uses all branded products sold in the balanced set of stores active over the entire period 2006-2022.

(3) Incomplete Spells

In a robustness sample, we exclude products with missing quarters to rule out the possibility that our results are driven by seasonal products, promotional items, or products with very little sales. Figures B5 and B6 present the main results.

Figure B5: Price Life Cycle: Robustness Incomplete spells

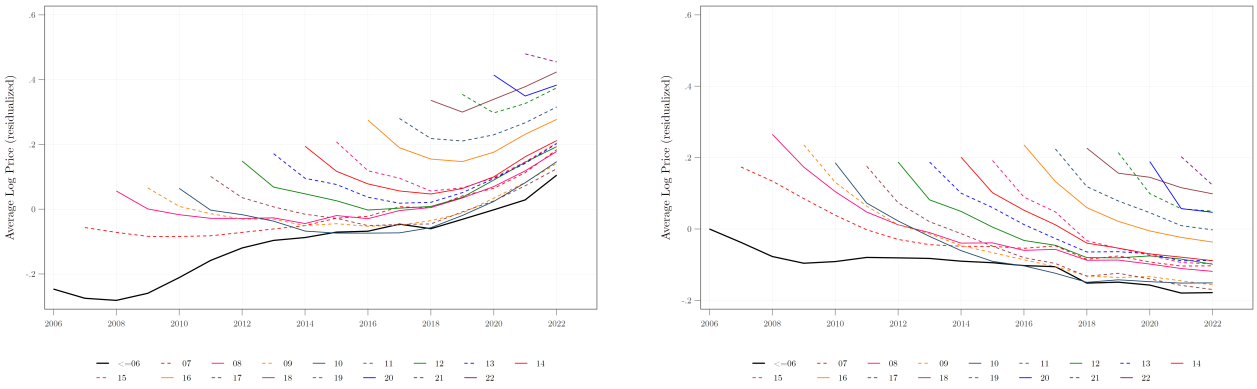
(a) Balanced cohorts w/o exit (b) Balanced cohorts w/ exit (c) Unbalanced cohorts w/exit



Notes: The figure shows the estimated age fixed effects of (log) prices of products computed using equations (3) and (4). We use the sector and sector x time effects, as well as iteration of sector and sector x quarter fixed effects with firm-level fixed effects. The estimation included all branded products, excluding products with incomplete observations. In (a) the age fixed effects are computed for balanced cohort and products duration of more than 20 quarters. In (b) the age fixed effects are computed for balanced cohorts but individual products can exit before 20 quarters of activity. In (c) we use an unbalanced sample of cohorts with exit. Main text provides details on the cohorts used.

Figure B6: Prices by Cohort of Products: Robustness Private Label

(a) Without Time Controls (b) With Time Controls



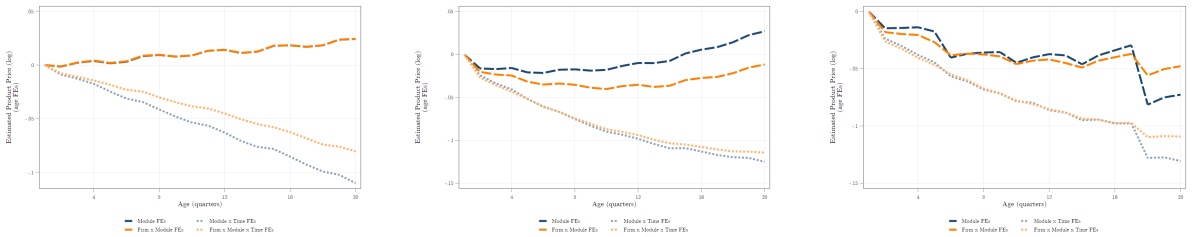
Notes: The figures display average log prices by cohort over time using the baseline annual-level sample. In panel (a), log prices are residualized using sector fixed effects. In panel (b), log prices are residualized using sector-by-year fixed effects. The estimation included all branded products, excluding products with incomplete observations.

(4) Coverage of Some Product Categories

We include robustness where we exclude Alcohol and General Merchandise because these are the departments for which the coverage in Nielsen RMS is smaller and less likely to be representative. Figures B7 and C1 present the main results.

Figure B7: Price Life Cycle: Robustness Coverage of Some Product Categories

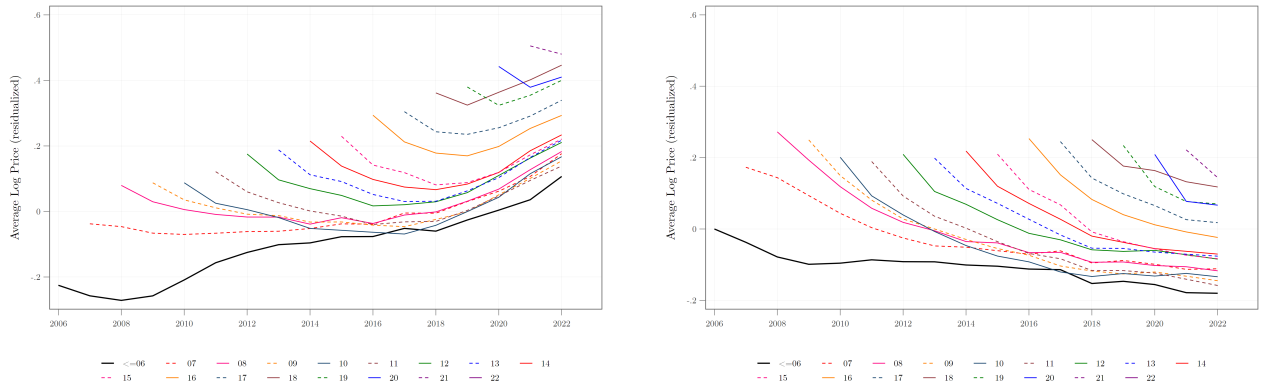
(a) Balanced cohorts w/o exit (b) Balanced cohorts w/ exit (c) Unbalanced cohorts w/exit



Notes: The figure shows the estimated age fixed effects of (log) prices of products computed using equations (3) and (4). We use the sector and sector x time effects, as well as iteration of sector and sector x quarter fixed effects with firm-level fixed effects. The estimation uses excludes all branded products except incomplete observations, and products from Alcohol and General Merchandise Sectors. In (a) the age fixed effects are computed for balanced cohort and products duration of more than 20 quarters. In (b) the age fixed effects are computed for balanced cohorts but individual products can exit before 20 quarters of activity. In (c) we use an unbalanced sample of cohorts with exit. Main text provides details on the cohorts used.

Figure B8: Prices by Cohort of Products: Robustness Coverage of Categories

(a) Without Time Controls (b) With Time Controls



Notes: The figures display average log prices by cohort over time using the baseline annual-level sample. In panel (a), log prices are residualized using sector fixed effects. In panel (b), log prices are residualized using sector-by-year fixed effects. The estimation uses excludes all branded products except incomplete observations, and products from Alcohol and General Merchandise Sectors.

C Frequency and Magnitude of Price Changes

C.1 Working Datasets

The raw Nielsen-RMS data has average weekly prices and value at product (upc $\times$ upc-version) $\times$ store $\times$ week. We start by scaling prices (and quantities) such that we use the same unit measure within each product category. We group products of the Nielsen RMS data by assigning an ID to a product (upc $\times$ upc-version) at a given store (store_code_uc). To measure price changes, we compute first differences, and detect changes, their magnitude and nature (positive/negative, small/large, and whether it reverts to prior level). Our baseline working dataset aggregates information to product-store-year level.

There are two important measurement issues. First, we have to handle missings in price information. Nielsen RMS does not measure prices when there is no quantity sold in a particular week $l = 1, ..., 52$. Our baseline measure of frequency of price changes only uses consecutive weeks with price information. We build two other measures for robustness, where we make different assumptions for the missing observations. Suppose we observe prices for unit i (product $\times$ store) at time t and $t-k$, p_{it-k} and p_{it}, we make the following assumptions: (i) assume $p_{i,t-k+1} = p_{i,t}$, price in $t-k+1$ changes to $p_{i,t}$ and remains the same until t; (ii) assume $p_{i,t-k} = p_{i,t-1}$, all prices between $p_{i,t-k}$ and $p_{i,t}$ do not change and stay at $p_{i,t-k}$.

The second measurement issue is related to the nature of price changes. The literature identified three broad classes of price changes: regular price changes for identical items, temporary sales, and price changes due to product substitution (Nakamura and Steinsson, 2008). Instead of trying to identify regular price changes, our approach consists in estimating all price changes and price changes likely associated with temporary price discounts within product $\times$ store, and directly measure product substitution. Nielsen RMS does not directly measure if a product $\times$ store is on temporary sale. As in Coibion et al. (2015), we develop distinct measures of temporary price reductions using information on the amount and duration of a price reduction.

C.1.1 Key Variables

Frequency and size of price changes – The measure of frequency of price changes is computed as:

$$\Delta_{it} = \frac{\sum_{l}^{\Omega_{it}} 1\{p_{it,l} - p_{it,l-1} \neq 0\}}{\sum_{l}^{\Omega_{it}} 1_{t,l}} \tag{61}$$

where for the baseline Ω_{it} is the set of consecutive weeks with price information, while for the robustness Ω_{it} includes all weeks. The baseline measures the share of price changes out of all price first differences that can be computed using Nielsen RMS, without imputations. The underlying

assumption of the baseline is that the missings in first differences are random. The underlying assumption of the robustness measure is that there were no price changes when no price data is observed.

For both baseline and robustness, we build measures that condition on whether the change is negative or positive, and if the differences are larger than 5% or 10%. We also build measures that capture decline in prices that are likely temporary.

We compute measures of the amount of price changes during the time t. We consider

$$M_{it}^{(1)} = \sum_{l}^{\Omega_{it}} |\ln p_{it,l} - \ln p_{it,l-1}| \tag{62}$$

$$M_{it}^{(2)} = \frac{\sum_{l}^{\Omega_{it}} |p_{it,l} - p_{it,l-1}|}{\bar{p}_{it}} \tag{63}$$

where $\bar{p}_{it}$ is the average price of product. We build measures that only account differences are larger than 5% or 10%. We also build measures that capture decline in prices that are likely temporary.

Temporary price discounts – We apply a sales filter similar to that in Nakamura and Steinsson (2008) and Coibion et al. (2015). Specifically, we consider a product on sale if a price reduction is followed by a price increase of similar magnitude within four weeks. We consider a version where the price increase should be of exact amount such that prices return to previous level, and another version where the price increase to within a level very close to the initial price.

C.1.2 Implementation

We build working datasets for each sector. These sector working datasets are used on the statistics and life cycle exercises.

C.2 Descriptive Statistics

In this section, we provide summary statistics on the frequencies of product price change in the United States from 2006 to 2020. The summary statistics are computed from a panel data set defined by upc $\times$ store_code_uc $\times$ year $\times$ module consisting of 328 modules in the Nielsen scanner data. We include one baseline frequency of price changes, two robust checks for frequencies of price changes, the frequency of positive price changes, the frequency of negative price changes, and three measures for the frequencies of sales in our analysis.

In Table 1, we calculate the median and average frequencies and durations of the price changes in the panel weighted by the revenue of products in each module. We follow Nakamura and Steinsson (2008) and compute the duration of price changes as $\frac{-1}{ln(1-f)}$. The statistics suggest that distributions for the frequencies and durations of price changes are right-skewed.

	All Price Changes					Temporary Price Changes		
	Baseline	Opt 1	Opt 2	Opt +	Opt -	Baseline	Opt 1	Opt 2
Median								
Frequency	0.212	0.078	0.078	0.083	0.095	0	0	0
Duration	2.611	5.150	5.150	5.150	5.039	7.739	2.804	2.466
Average								
Frequency	0.291	0.163	0.162	0.139	0.152	0.048	0.143	0.174
Duration	4.962	11.040	11.040	8.186	7.982	9.536	3.557	3.309

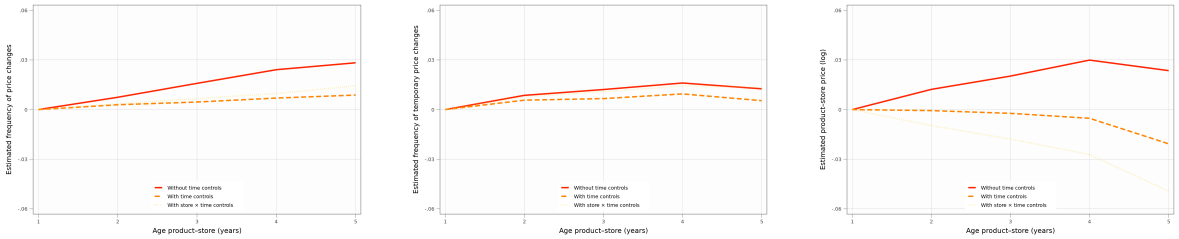
Table C1: The Frequency and Duration of Price Changes from 2006-2020

C.3 Life Cycle

In this section, we provide life cycle estimates on the frequencies of product price change in the United States from 2006 to 2020. The summary statistics are computed from a panel data set defined by upc × store_code_uc × year × module consisting of 328 modules in the Nielsen scanner data.

Figure C1: Prices by Cohort of Products: Robustness Coverage of Categories

(a) All Price Changes (b) Temporary Price Changes (c) Level Prices (log)



Notes: The figures displays the life cycle of upc x store for the balanced sample in the first 5 years of activity.

D Additional Model

D.1 Omitted Proof in Prices and Profits

The quality-adjusted price index is given by,

$$\hat{P}_t = \frac{W_t}{A_t} m_t N_t^{\frac{1}{1-\sigma}} Q_t^{-1} \tag{64}$$

The profit as share of nominal GDP is given by,

$$\frac{\pi_{it}}{\hat{P}_t Y_t} = \frac{1}{\hat{P}_t}(p_{it} - W_t) q_{it}^{\sigma-1} (\frac{p_{it}}{\hat{P}_t})^{-\sigma}$$

Plug in the price index, we get the profit. Next, we derive the aggregate output. The labor demand of firm i is given by,

$$A_t L_{it} = q_{it}^{\sigma-1} p_{it}^{-\sigma} P_t^{\sigma} Y_t \tag{65}$$

$$= \hat{q}_{it}^{\sigma-1} \mu_{it}^{-\sigma} {m_t}^{\sigma} N_t^{\frac{\sigma}{1-\sigma}} \cdot \frac{Y_t}{Q_t} \tag{66}$$

The aggregate labor demand is,

$$A_t L = \int \frac{1}{N} \hat{q}_{it}^{\sigma-1} \mu_{it}^{-\sigma} \, di \cdot {m_t}^{\sigma} N^{\frac{1}{1-\sigma}} \frac{Y_t}{Q_t} \tag{67}$$

Rearrange,

$$A_t \frac{{m_t}^{-\sigma}}{\mathbb{E}_q[\mu_{it}^{-\sigma}]} N_t^{\frac{1}{\sigma-1}} Q_t L = Y_t \tag{68}$$

D.2 Proof for Lemma 2

Proof. Fix a product line with state $(\hat{q}, \mu)$. Under Assumption **??**, flow profits and all R&D costs are proportional to $\hat{q}^{\sigma-1}$, and successful innovation scales quality multiplicatively while allowing a price reset to μ_0. Consider the line-level HJB (in quality-adjusted units). Guess $\Gamma(\hat{q}, \mu) = \hat{q}^{\sigma-1} A(\mu)$. Then for any innovation type $k \in \{I, E, N\}$,

$$\Gamma(\lambda^k \hat{q}, \mu_0) - \Gamma(\hat{q}, \mu) = \hat{q}^{\sigma-1} \Big((\lambda^k)^{\sigma-1} A(\mu_0) - A(\mu) \Big),$$

and the corresponding innovation objective

$$\max_{x^k \geq 0} \Big\{ x^k \big[\Gamma(\lambda^k \hat{q}, \mu_0) - \Gamma(\hat{q}, \mu) \big] - \xi^k (x^k)^{\frac{1}{1-a}} \hat{q}^{\sigma-1} \Big\}$$

is $\hat{q}^{\sigma-1}$ times a function of μ alone. The markup drift term depends only on μ (via $\dot{\mu}/\mu = -\pi$), and the hat-normalization contributes only through the aggregate growth rate, which is common across

lines on the BGP. Hence every term in the line HJB is homogeneous of degree $\sigma - 1$ in $\hat{q}$, so dividing the HJB by $\hat{q}^{\sigma-1}$ yields a closed one-dimensional equation for $A(\mu)$, verifying (122).

Because the reduced problem depends on the line state only through μ, the associated optimal policies are also functions of μ alone. In particular, the optimal reset markup μ_0 and the exit cutoff μ_b are identical for all $\hat{q}$, so (μ_0, μ_b) are common across product lines. Moreover, by assumption internal, external, and entrant innovations are feasible on both active and inactive lines and their success hazards and step sizes do not depend on μ or the line's active status. Therefore the stochastic process for $\hat{q}$ is governed solely by the innovation block and is independent of the markup process and exit decisions, implying independence of the stationary distributions.

Finally, the per-line external innovation choice reduces to

$$\max_{x^E \geq 0} \left\{ x^E \, A(\mu_0)(\lambda^E)^{\sigma-1} - \xi^E (x^E)^{\frac{1}{1-\alpha}} \right\},$$

whose first-order condition gives

$$x^{E*} = \left(\frac{1-\alpha}{\xi^E} A(\mu_0)(\lambda^E)^{\sigma-1} \right)^{\frac{1-\alpha}{\alpha}}.$$

Substituting back yields the maximized value

$$\Lambda^E = \xi^E \frac{\alpha}{1-\alpha} \left(x^{E*} \right)^{\frac{1}{1-\alpha}},$$

establishing (127). ■

D.3 Proof for Proposition 1:

Proof. In the inactive region $\mu \leq \mu_b$, the HJB implies $\varphi A(\mu) = \Lambda^I(\mu_0)$, hence

$$A(\mu) = \frac{\Lambda^I(\mu_0)}{\varphi} \quad \text{for all } \mu \leq \mu_b,$$

so $A'(\mu_b^-) = 0$. By smooth pasting at the free boundary, $A'(\mu_b^+) = A'(\mu_b^-) = 0$, hence $A'(\mu_b) = 0$.

Evaluate the active-region HJB at $\mu = \mu_b$ and use $A(\mu_b) = \Lambda^I(\mu_0)/\varphi$ and $A'(\mu_b) = 0$:

$$\varphi \cdot \frac{\Lambda^I(\mu_0)}{\varphi} = -\pi \mu_b A'(\mu_b) + m^{\sigma-1} N^{-1} (\mu_b - 1) \mu_b^{-\sigma} + \Lambda^I(\mu_0) + \gamma A(\mu_0).$$

Cancelling $\Lambda^I(\mu_0)$ on both sides and using $A'(\mu_b) = 0$ yields

$$0 = m^{\sigma-1} N^{-1} (\mu_b - 1) \mu_b^{-\sigma} + \gamma A(\mu_0),$$

which is exactly (42). ■

D.4 Proof for Proposition 2:

Proof. Let $C \equiv m^{\sigma-1}N^{-1}$, $a \equiv \varphi/\pi$, $L \equiv \Lambda(\mu_0)$ and $A_0 \equiv A(\mu_0)$.

Step 1 (Solve the ODE on $\{\mu > \mu_b\}$). For $\mu > \mu_b$ the HJB is

$$\varphi A(\mu) = -\pi\mu A'(\mu) + C(\mu-1)\mu^{-\sigma} + L + \gamma A_0. \tag{69}$$

Rewrite as $A'(\mu) + \frac{a}{\mu}A(\mu) = \frac{C}{\pi}\big(\mu^{-\sigma} - \mu^{-\sigma-1}\big) + \frac{L+\gamma A_0}{\pi}\mu^{-1}$. Using integrating factor μ^a and assuming $a \neq \sigma$ and $a \neq \sigma - 1$,

$$A(\mu) = \frac{C}{\pi}\Big[\frac{\mu^{1-\sigma}}{a-\sigma+1} - \frac{\mu^{-\sigma}}{a-\sigma}\Big] + \frac{L+\gamma A_0}{\varphi} + K\mu^{-a}, \tag{70}$$

for a constant K. Define

$$\xi_1 = \frac{C}{\pi}(a-\sigma+1)^{-1}, \qquad \xi_2 = \frac{C}{\pi}(a-\sigma)^{-1}.$$

Step 2 (Value matching at μ_b). For $\mu \leq \mu_b$, the HJB gives $\varphi A(\mu) = L$, hence $A(\mu) = L/\varphi$ and value matching implies $A(\mu_b) = L/\varphi$. Plugging $\mu = \mu_b$ into the general solution yields

$$K\mu_b^{-a} = -\xi_1\mu_b^{1-\sigma} + \xi_2\mu_b^{-\sigma} - \frac{\gamma A_0}{\varphi}.$$

Substituting K back and using $\mu_b^{1-\sigma}(\mu_b/\mu)^a = \mu^{1-\sigma}(\mu_b/\mu)^{a+1-\sigma}$ and $\mu_b^{-\sigma}(\mu_b/\mu)^a = \mu^{-\sigma}(\mu_b/\mu)^{a-\sigma}$, we obtain for $\mu > \mu_b$,

$$A(\mu) = \frac{L}{\varphi} + \xi_1\mu^{1-\sigma}\Big(1 - \big(\tfrac{\mu_b}{\mu}\big)^{a+1-\sigma}\Big) - \xi_2\mu^{-\sigma}\Big(1 - \big(\tfrac{\mu_b}{\mu}\big)^{a-\sigma}\Big) + \frac{\gamma A_0}{\varphi}\Big(1 - \big(\tfrac{\mu_b}{\mu}\big)^{a}\Big). \tag{71}$$

Thus the claimed form holds with $\xi_3 = \gamma A_0/\varphi$.

Step 3 (Smooth pasting pins down μ_b). Since $A(\mu) \equiv L/\varphi$ is constant on $\{\mu \leq \mu_b\}$, $A'(\mu_b^-) = 0$ and smooth pasting implies $A'(\mu_b^+) = 0$. Evaluate the continuation HJB at μ_b using $A(\mu_b) = L/\varphi$ and $A'(\mu_b) = 0$:

$$L = C(\mu_b - 1)\mu_b^{-\sigma} + L + \gamma A_0 \quad \Rightarrow \quad C(\mu_b - 1)\mu_b^{-\sigma} + \gamma A_0 = 0.$$

Step 4 (Value maximization pins down μ_0). If $\mu_0 > \mu_b$ is an interior maximizer, then $A'(\mu_0) = 0$. Evaluating the continuation HJB at μ_0 and using $A(\mu_0) = A_0$ gives

$$\varphi A_0 = C(\mu_0 - 1)\mu_0^{-\sigma} + L + \gamma A_0 \quad \Rightarrow \quad (\varphi - \gamma)A_0 = C(\mu_0 - 1)\mu_0^{-\sigma} + \Lambda(\mu_0).$$

■

D.5 Proof for Proposition 9

Proof. Consider the stationary HJB in the inaction region:

$$0 = -(\phi + \gamma - g^Q)A(\mu) - \pi\mu A'(\mu) + \tilde{\pi}(\mu) + \tilde{\Lambda}^I(\mu) + \gamma A(\mu^u). \tag{72}$$

Evaluating (72) at $\mu = \mu^u$ and using $A(\mu^u) = A(\mu^u)$ yields

$$0 = -(\phi + \gamma - g^Q)A(\mu^u) + \tilde{\pi}(\mu^u) + \tilde{\Lambda}^I(\mu^u) + \gamma A(\mu^u),$$

which simplifies to

$$(\phi - g^Q)A(\mu^u) = \tilde{\pi}(\mu^u) + \tilde{\Lambda}^I(\mu^u).$$

This proves (137).

Next, because μ^u is an interior maximizer of A, it satisfies

$$A'(\mu^u) = 0, \qquad A''(\mu^u) \leq 0.$$

Differentiating (72) with respect to μ gives

$$0 = -(\phi + \gamma - g^Q)A'(\mu) - \pi\big(A'(\mu) + \mu A''(\mu)\big) + \tilde{\pi}'(\mu) + \tilde{\Lambda}^{I\prime}(\mu).$$

Evaluating at $\mu = \mu^u$ and using $A'(\mu^u) = 0$ yields

$$\tilde{\pi}'(\mu^u) + \tilde{\Lambda}^{I\prime}(\mu^u) = \pi\mu^u A''(\mu^u) \leq 0. \tag{73}$$

It is easy to show $\tilde{\Lambda}^{I\prime}(\mu^u) = 0$, this implies

$$\tilde{\pi}'(\mu^u) \leq 0.$$

Now

$$\tilde{\pi}(\mu) = m^{\sigma-1}(\mu - 1)\mu^{-\sigma}$$

is strictly concave and uniquely maximized at

$$\mu^* = \frac{\sigma}{\sigma - 1}.$$

Hence

$$\tilde{\pi}'(\mu) > 0 \quad \text{for } \mu < \mu^*, \qquad \tilde{\pi}'(\mu^*) = 0, \qquad \tilde{\pi}'(\mu) < 0 \quad \text{for } \mu > \mu^*.$$

Therefore, $\tilde{\pi}'(\mu^u) \leq 0$ implies $\mu^u \geq \mu^*$. If $A''(\mu^u) < 0$ and $\pi > 0$, then the right-hand side of (73) is strictly negative, so $\tilde{\pi}'(\mu^u) < 0$, which implies

$$\mu^u > \mu^*.$$

This proves price overshooting.

Finally, when pricing frictions vanish, the benefit of choosing a reset markup above the static optimum disappears. If $\pi \to 0$, markup drift vanishes after adjustment; if $\gamma \to \infty$, free adjustment opportunities arrive arbitrarily frequently. In either case, the dynamic problem converges to the static pricing problem, so the reset markup converges to the unique maximizer of current profit:

$$\mu^u \to \mu^*.$$

■

D.6 Proof for Proposition 4

Proof. **Markup drift and hazards.** Markups drift deterministically due to inflation. Over a small interval Δt,

$$\mu_{t+\Delta t} = \mu_t(1 - \pi\Delta t) + o(\Delta t), \qquad \Longleftrightarrow \qquad \mu_t \geq \mu \;\Leftrightarrow\; \mu_{t+\Delta t} \geq \mu(1 + \pi\Delta t) \text{ absent events.}$$

Let

$$\kappa \equiv x^I + \tau + \psi, \qquad \lambda \equiv x^I + \tau + \gamma + \psi, \qquad \eta \equiv \frac{\lambda}{\pi}.$$

Here λ is the total *refresh* hazard for the markup while a line is active: internal innovation, takeovers (incumbent external and entrant), Calvo resets, and exogenous obsolescence all restart the pricing cycle at the reset markup μ_u. The parameter κ is the total *reactivation* hazard from the inactive pool into active status: internal innovation, takeovers, and obsolescence-driven replacement all (re)introduce a line at μ_u.

Unconditional markup distribution. Let $\bar{H}(\mu) \equiv \Pr(\text{markup} \geq \mu)$ denote the stationary survival function over the full population of product lines (active and inactive), with support $\mu \in [0, \mu_u]$. For any $\mu < \mu_u$, stationarity implies

$$\bar{H}(\mu) = (1 - \lambda\Delta t)\,\bar{H}\big(\mu(1 + \pi\Delta t)\big) + o(\Delta t), \tag{74}$$

because with probability $\lambda\Delta t$ an event occurs over Δt that refreshes the markup to μ_u (internal/external innovation, a Calvo reset, or obsolescence and replacement), and otherwise the markup drifts deterministically. Using the expansion

$$\bar{H}(\mu(1 + \pi\Delta t)) = \bar{H}(\mu) + \pi\mu\bar{H}'(\mu)\Delta t + o(\Delta t)$$

in (74) and letting $\Delta t \to 0$ yields the ODE

$$\pi\mu\bar{H}'(\mu) = \lambda\,\bar{H}(\mu). \tag{75}$$

Imposing the boundary condition $\bar{H}(\mu_u) = 1$ gives

$$\bar{H}(\mu) = \left(\frac{\mu}{\mu_u}\right)^{\eta}, \qquad \mu \in [0, \mu_u], \qquad \eta = \frac{x^I + \tau + \gamma + \psi}{\pi}. \tag{76}$$

Active-line distribution on $[\mu_b, \mu_u]$. Let $H(\mu)$ be the *unnormalized* stationary CDF over active products,

$$H(\mu) \equiv \text{mass of active products with markup} \leq \mu, \qquad \mu \in [\mu_b, \mu_u],$$

so that $H(\mu_b) = 0$ and $H(\mu_u) = N$, where N is the mass of active products. For $\mu \in [\mu_b, \mu_u)$, stationarity implies

$$H(\mu) = (1 - \lambda\Delta t)\, H\big(\mu(1 + \pi\Delta t)\big) - (1 - \lambda\Delta t)\, H\big(\mu_b(1 + \pi\Delta t)\big) + o(\Delta t), \tag{77}$$

because (i) drift maps the threshold μ to $\mu(1 + \pi\Delta t)$ absent events, (ii) events with total hazard λ remove mass from the interior by resetting it to μ_u, and (iii) the lower boundary μ_b truncates the active support. Expanding and letting $\Delta t \to 0$ yields

$$\pi\mu H'(\mu) = \lambda\big(H(\mu) - H(\mu_b)\big) = \lambda H(\mu), \qquad \mu \in (\mu_b, \mu_u), \tag{78}$$

since $H(\mu_b) = 0$. Thus the active density takes the power form

$$h(\mu) \equiv H'(\mu) = K\,\mu^{\eta - 1}, \qquad \eta = \frac{\lambda}{\pi}. \tag{79}$$

Normalizing by $H(\mu_u) = N$ gives

$$H(\mu) = N\,\frac{\left(\frac{\mu}{\mu_b}\right)^{\eta} - 1}{\left(\frac{\mu_u}{\mu_b}\right)^{\eta} - 1}, \qquad \mu \in [\mu_b, \mu_u], \qquad \eta = \frac{x^I + \tau + \gamma + \psi}{\pi}. \tag{80}$$

Mass of active products. The outflow from active to inactive equals the probability flux hitting the lower boundary μ_b:

$$\text{outflow} = \pi\,\mu_b\,h(\mu_b) = \pi\,\mu_b\,K\,\mu_b^{\eta - 1} = \lambda N \cdot \frac{1}{\left(\frac{\mu_u}{\mu_b}\right)^{\eta} - 1},$$

where the last equality uses the normalization implied by (80). The inflow from inactive to active is the reactivation rate κ times the inactive mass $(1 - N)$:

$$\text{inflow} = \kappa(1 - N), \qquad \kappa = x^I + \tau + \psi,$$

because internal innovation, takeovers, and obsolescence-driven replacement all (re)introduce lines at μ_u. Stationarity of the active mass, inflow = outflow, therefore implies

$$N = \frac{\kappa}{\kappa + \dfrac{\lambda}{\left(\frac{\mu_u}{\mu_b}\right)^{\eta} - 1}}, \qquad \kappa = x^I + \tau + \psi, \quad \lambda = x^I + \tau + \gamma + \psi, \quad \eta = \frac{\lambda}{\pi}. \tag{81}$$

Average markup and aggregate efficiency (active lines). Let the stationary conditional density of markups among active lines be

$$f_A(\mu) = \frac{\eta\, \mu^{\eta-1}}{\mu_u^{\eta} - \mu_b^{\eta}}, \qquad \mu \in [\mu_b, \mu_u], \qquad \eta \equiv \frac{x^I + \tau + \gamma + \psi}{\pi}. \tag{82}$$

(Here η uses the total refresh hazard for active markups: internal innovation x^I, takeovers τ, Calvo resets γ, and obsolescence replacement ψ.)

Aggregate markup. Define the average markup over active lines by

$$m_t \equiv \left(\frac{1}{N_t} \int_{i \in \mathcal{N}_t} \mu_{it}^{1-\sigma} \hat{q}_{it}^{\sigma-1}\, di \right)^{\frac{1}{1-\sigma}} = \left(\mathbb{E}_A \left[\mu^{1-\sigma} \right] \right)^{\frac{1}{1-\sigma}}, \tag{83}$$

where the second equality uses independence between μ and $\hat{q}$ across active lines (so $\mathbb{E}_A[\hat{q}^{\sigma-1}] = 1$ under the relative-quality normalization). Using (82),

$$\begin{aligned} \mathbb{E}_A \left[\mu^{1-\sigma} \right] &= \int_{\mu_b}^{\mu_u} \mu^{1-\sigma} f_A(\mu)\, d\mu = \frac{\eta}{\mu_u^{\eta} - \mu_b^{\eta}} \int_{\mu_b}^{\mu_u} \mu^{\eta-\sigma}\, d\mu \\ &= \frac{\eta}{\eta - (\sigma-1)} \frac{\mu_u^{\eta-(\sigma-1)} - \mu_b^{\eta-(\sigma-1)}}{\mu_u^{\eta} - \mu_b^{\eta}}, \qquad \eta \neq \sigma - 1. \end{aligned} \tag{84}$$

Therefore,

$$m_t = \left[\frac{\eta}{\eta - (\sigma-1)} \frac{\mu_u^{\eta-(\sigma-1)} - \mu_b^{\eta-(\sigma-1)}}{\mu_u^{\eta} - \mu_b^{\eta}} \right]^{\frac{1}{1-\sigma}}. \tag{85}$$

Letting $\rho \equiv \mu_b / \mu_u \in (0, 1)$, this can be written as

$$\frac{m_t}{\mu_u} = \left[\frac{\eta}{\eta - (\sigma-1)} \frac{1 - \rho^{\eta-(\sigma-1)}}{1 - \rho^{\eta}} \right]^{-\frac{1}{\sigma-1}}. \tag{86}$$

Aggregate efficiency. Define aggregate efficiency over active lines by

$$v_t \equiv \frac{m_t^{-\sigma}}{\mathbb{E}_A[\mu^{-\sigma}]}. \tag{87}$$

Again using (82),

$$\begin{aligned} \mathbb{E}_A \left[\mu^{-\sigma} \right] &= \int_{\mu_b}^{\mu_u} \mu^{-\sigma} f_A(\mu)\, d\mu = \frac{\eta}{\mu_u^{\eta} - \mu_b^{\eta}} \int_{\mu_b}^{\mu_u} \mu^{\eta-\sigma-1}\, d\mu \\ &= \frac{\eta}{\eta - \sigma} \frac{\mu_u^{\eta-\sigma} - \mu_b^{\eta-\sigma}}{\mu_u^{\eta} - \mu_b^{\eta}}, \qquad \eta \neq \sigma. \end{aligned} \tag{88}$$

Combining (85) and (88) gives

$$v_t = \frac{\left[\frac{\eta}{\eta-(\sigma-1)} \frac{\mu_u^{\eta-(\sigma-1)} - \mu_b^{\eta-(\sigma-1)}}{\mu_u^{\eta} - \mu_b^{\eta}} \right]^{-\frac{\sigma}{1-\sigma}}}{\frac{\eta}{\eta-\sigma} \frac{\mu_u^{\eta-\sigma} - \mu_b^{\eta-\sigma}}{\mu_u^{\eta} - \mu_b^{\eta}}}. \tag{89}$$

Equivalently, in terms of $\rho = \mu_b/\mu_u$,

$$v_t = \frac{\eta-\sigma}{\eta}\,\frac{1-\rho^{\eta}}{1-\rho^{\eta-\sigma}}\left[\frac{\eta}{\eta-(\sigma-1)}\,\frac{1-\rho^{\eta-(\sigma-1)}}{1-\rho^{\eta}}\right]^{\frac{\sigma}{\sigma-1}}. \tag{90}$$

■

D.7 Proof for Proposition 7:

The growth rate involves random destruction of existing products and introduction of new products of higher quality. Since these processes are random, the average quality is just Q_t and Q_t times a function of innovation step.

$$Q_t^{\sigma-1} = \int_{\mathcal{N}} q_{it}^{\sigma-1}\,di$$

$$\begin{aligned}(\sigma-1)\ln Q_{t+\Delta t} &= \ln\left[(1-\tau\Delta t - x^I\Delta t)\int_{\mathcal{N}} q_{it}^{\sigma-1}\,di + \tau\Delta t\int_{\mathcal{N}}(1+\lambda^E)^{\sigma-1}q_{it}^{\sigma-1}\,di + x^I\Delta t\int_{\mathcal{N}}(1+\lambda^I)^{\sigma-1}q_{it}^{\sigma-1}\,di\right]\\ &= \ln\left(\int_{\mathcal{N}} q_{it}^{\sigma-1}\,di + \tau\Delta t\left[(1+\lambda^E)^{\sigma-1}-1\right]\int_{\mathcal{N}} q_{it}^{\sigma-1}\,di + x^I\Delta t\left[(1+\lambda^I)^{\sigma-1}-1\right]\int_{\mathcal{N}} q_{it}^{\sigma-1}\,di\right)\\ &= (\sigma-1)\ln Q_t + \tau\left[(1+\lambda^E)^{\sigma-1}-1\right]\Delta t + x^I\left[(1+\lambda^I)^{\sigma-1}-1\right]\Delta t\end{aligned}$$

The consumption share is straightforward. Now we turn to expected product lifetime

Consider a product with an initial markup μ_0 that decays deterministically at rate π toward an exit threshold μ_b. The product also faces an exogenous shock rate $\psi + x^I + \tau$ and a reset rate γ that restores the markup to μ_0.

The time T required for the markup to reach the floor μ_b in the absence of shocks or resets is:

$$T = \frac{\ln(\mu_0) - \ln(\mu_b)}{\pi} \tag{91}$$

The expected lifetime $E[L]$ is the sum of the expected time spent in the current cycle and the probability-weighted future lifetime if a reset occurs:

$$E[L] = \int_0^T e^{-(\psi+x^I+\tau+\gamma)t}\,dt + P(\text{Reset before } T)E[L] \tag{92}$$

The expected survival time within a single cycle (before a shock, reset, or reaching T) is:

$$\int_0^T e^{-(\psi+x^I+\tau+\gamma)t}\,dt = \frac{1-e^{-(\psi+x^I+\tau+\gamma)T}}{\psi+x^I+\tau+\gamma} \tag{93}$$

The probability that a reset occurs before an exogenous shock and before time T is:

$$P(\text{Reset before } T) = \frac{\gamma}{\psi+x^I+\tau+\gamma}\left(1-e^{-(\psi+x^I+\tau+\gamma)T}\right) \tag{94}$$

Combining the above and solving for $E[L]$ yields the final expression:

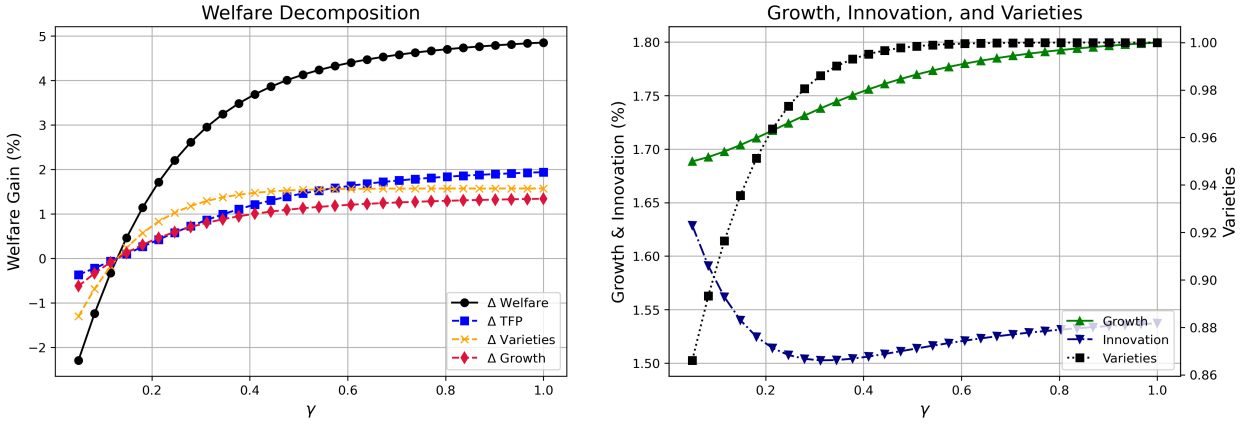
$$E[L] = \frac{1 - e^{-(\psi + x^I + \tau + \gamma)T}}{\psi + x^I + \tau + \gamma e^{-(\psi + x^I + \tau + \gamma)T}} \tag{95}$$

where $T = \frac{\ln(\mu_0/\mu_b)}{\pi}$.

D.8 Welfare and Growth Decomposition with $\mu_0 = \mu^*$

Below we show the welfare and growth when the reset markup is constrained to be μ^*.

Figure D1: Welfare and Growth vs. Pricing Friction



Notes: **The left panel plots the welfare change compared to the baseline welfare in the range of $\gamma = 0.1$ to $\gamma = 1.5$, and decomposition into misallocation, varieties and growth. The right panel plots the growth rate and its decomposition into innovation rate and number of varieties.**

E Model Extension

This section develops a model of creative destruction with endogenous innovation, embedded in a monetary general equilibrium environment with trend money growth and nominal pricing frictions. The economy is in continuous time and admits a balanced growth path (BGP) on which real allocations are stationary after appropriate normalization. The innovation block follows the multi-product framework in Klette and Kortum (2004) and the quality-ladder structure in Aghion and Howitt (1992); the monetary side adopts a money-in-the-utility formulation and menu-cost combined with free adjustments (Golosov and Lucas, 2007; Nakamura and Steinsson, 2008).

E.1 Monetary Economy

We adopt preferences at the intersection of Golosov and Lucas (2007), which facilitates the analysis of dynamics, and balanced-growth preferences, which imply constant labor supply and labor share along the BGP. Time is continuous. A representative household consumes the final good C_t, holds nominal money balances M_t, and owns a portfolio of nominal assets A_t (claims to firms and government bonds) that yields the nominal return R_t. Also, the real interest rate is denoted by $r_t = R_t - \pi_t$. Household also supplies labor L_t and get nominal wage W_t. The preference is:

$$\max_{\{C_t, L_t, M_t, B_t\}_{t=0}^{\infty}} \mathbb{E}_0 \int_0^{\infty} e^{-\rho t} \left[\ln C_t - L_t + \chi \ln \left(\frac{M_t}{\hat{P}_t} \right) \right] \tag{96}$$

subject to

$$\dot{A}_t + \dot{M}_t = W_t L_t + R_t A_t - \hat{P}_t C_t,$$

where $\rho > 0$ is the discount rate, and $\hat{P}_t$ is the quality-adjusted price index defined below. The first-order conditions imply,

$$\hat{P}_t C_t = W_t, \quad \hat{P}_t C_t = \chi^{-1} R_t M_t, \quad \frac{\dot{C}_t}{C_t} = R_t - \frac{\dot{\hat{P}}_t}{\hat{P}_t} - \rho. \tag{97}$$

Suppose the money supply increases at a constant rate $\bar{p}i$, our assumptions on household preference lead to the following results:

Proposition 6. *Equilibrium aggregates satisfy*

$$R_t = \rho + \bar{\pi}, \quad W_t = \chi^{-1} R M_t, \quad \hat{P}_t C_t = \chi^{-1} R M_t. \tag{98}$$

where $\bar{\pi}$ is a constant money growth rate.

E.2 Production, Product Lines, and Demand

There is a unit measure of product lines. The final good is a CES aggregate of quality-adjusted intermediate inputs:

$$Y_t = \left(\int (q_{it} y_{it})^{\frac{\sigma-1}{\sigma}} \, di \right)^{\frac{\sigma}{\sigma-1}}, \tag{99}$$

where $\sigma > 1$ is the elasticity of substitution, q_{it} denotes product quality, and y_{it} is the quantity of intermediate input i. Intermediate goods are produced with labor under a linear technology,

$$y_{it} = A_t \, \ell_{it}, \tag{100}$$

where ℓ_{it} is labor used in line i and A_t is an aggregate productivity shifter (normalized to one if desired). The demand for intermediate good i is:

$$y_{it} = q_{it}^{\sigma-1} \left(\frac{p_{it}}{\hat{P}_t} \right)^{-\sigma} Y_t, \tag{101}$$

where p_{it} is the nominal price of good i and the quality-adjusted price index is

$$\hat{P}_t = \left(\int \left(\frac{p_{it}}{q_{it}} \right)^{1-\sigma} di \right)^{\frac{1}{1-\sigma}}. \tag{102}$$

Define the quality index

$$Q_t \equiv \left(\int q_{it}^{\sigma-1} \, di \right)^{\frac{1}{\sigma-1}}, \tag{103}$$

and relative quality $\hat{q}_{it} \equiv q_{it}/Q_t$, which satisfies $\int \hat{q}_{it}^{\sigma-1} \, di = 1$.

E.3 Pricing, Markups, and Profits

The nominal marginal cost is same across firms, $\frac{W_t}{A_t}$. We define the product-level markup μ_{it} by

$$p_{it} = \mu_{it} \frac{W_t}{A_t}. \tag{104}$$

Given (101), markup dispersion affects aggregate allocations through the quality-adjusted price index. Define the quality-weighted aggregate markup

$$m_t \equiv \left(\int \mu_{it}^{1-\sigma} \hat{q}_{it}^{\sigma-1} \, di \right)^{\frac{1}{1-\sigma}} = \left(\mathbb{E}_{\hat{q}} \left[\mu_{it}^{1-\sigma} \right] \right)^{\frac{1}{1-\sigma}} = \frac{\hat{P}_t Q_t}{W_t / A_t}, \tag{105}$$

where $\mathbb{E}_{\hat{q}}[\cdot]$ denotes expectation under the quality-weighted cross-sectional distribution. On a BGP, both $\hat{P}_t Q_t$ and W_t grow at the inflation rate, implying that m_t is constant.

Combining (101) with labor-market clearing yields

$$Y_t = A_t \, v_t \, Q_t \, L_t, \qquad v_t \equiv \frac{m_t^{-\sigma}}{\mathbb{E}_{\hat{q}} \left[\mu_{it}^{-\sigma} \right]}. \tag{106}$$

where v_t measures the aggregate efficiency loss from the dispersion of markups across products. Compared to the efficient case where all firms charge the same markup, greater dispersion in markups implies that the household over-consumes low-markup goods and under-consumes high-markup goods, leading to misallocation and lower aggregate efficiency. The labor share in quality-adjusted terms satisfies

$$\frac{W_t L_t}{\hat{P}_t Y_t} = (v_t m_t)^{-1} = \frac{\mathbb{E}_{\hat{q}}\left[\mu_{it}^{1-\sigma}\right]}{\mathbb{E}_{\hat{q}}\left[\mu_{it}^{-\sigma}\right]}, \tag{107}$$

which is a sales-weighted harmonic mean of markups, consistent with the result in Edmond et al. (2022). Finally, the real profit flow of product i (in units of the quality-adjusted final good) is

$$\pi_{it} = m_t^{\sigma-1} Y_t \, \hat{q}_{it}^{\sigma-1} \left[(\mu_{it} - 1)\mu_{it}^{-\sigma}\right]. \tag{108}$$

Profits are proportional to market size Y_t and increasing in aggregate markups. Holding aggregate conditions fixed, higher relative quality raises profits, while profits are hump-shaped in the product's own markup with the static optimum at $\mu^* = \sigma/(\sigma - 1)$. With nominal pricing frictions (introduced below), firms face a dynamic trade-off between current and future profits.

E.4 Markup Dynamics and Price Adjustment

Between price adjustments, the nominal price is fixed. Therefore the markup evolves with nominal marginal cost. In general,

$$\frac{\dot{\mu}_{it}}{\mu_{it}} = -\frac{\dot{W}_t}{W_t} + \frac{\dot{A}_t}{A_t}. \tag{109}$$

Since the nominal wage grows at a constant rate $\bar{\pi}$, and we assume that the aggregate productivity does not grow, then this reduces to

$$\frac{\dot{\mu}_{it}}{\mu_{it}} = -\bar{\pi}. \tag{110}$$

A firm operating a product line with markup μ can adjust its price through three channels. First, with Poisson rate γ, it receives a free adjustment opportunity and resets its markup to the optimal level μ_{it}^u. Second, at any time, it can pay a menu cost $\kappa \hat{P}_t Y_t \hat{q}^{\sigma-1}$ to reset to the same level. As in standard menu-cost models, paying the menu cost becomes optimal when the markup drifts down to the lower adjustment threshold μ_{it}^b. We assume the menu cost is proportional to relative quality for tractability. Third, successful quality improvement replaces the existing product with a new vintage, at which point the firm resets the price of the new product.

E.5 Innovation

Innovation raises product quality and reallocates ownership of product lines through creative destruction. Incumbent firms engage in two types of R&D. *Internal innovation* upgrades product lines already owned by the firm. *External innovation* improves product lines not currently owned by the firm; upon success, the innovator acquires the upgraded line and displaces the previous owner. New entrants innovate in the same way and enter upon success. External and entrant innovation are *untargeted*: each innovation attempt draws a product line uniformly at random from the unit measure of lines, regardless of its quality. A firm's size is the measure of product lines it operates; when it loses a line, its portfolio shrinks, and it exits when it loses its last remaining line.

Successful innovation also interacts closely with pricing. Innovation provides a way to circumvent pricing frictions. Internal innovation allows the firm to relaunch an existing line as a new vintage, reset its price, and restore its markup. External innovation and entry play an analogous role at the portfolio level: successful innovators acquire newly upgraded lines and optimally price them upon introduction, thereby refreshing the firm's product mix with newly priced products. Thus, internal innovation relaxes pricing frictions at the line level, whereas external innovation does so at the portfolio level through creative destruction.

Following Klette and Kortum (2004), innovation is based on the firm's knowledge capital, proxied by the number of product lines in its portfolio. Formally, a firm operating n lines runs n research programs in parallel, so the arrival rate of innovation opportunities scales proportionally with n. Accordingly, we work with per-line innovation intensities and aggregate innovation costs by summing across the firm's portfolio.

For a product line of relative quality $\hat{q}$, an incumbent can invest in internal innovation at intensity $x^I > 0$, paying flow cost

$$c^I(x^I, \hat{q}) = \xi^I \left(x^I\right)^{\frac{1}{1-\alpha}} \hat{q}^{\sigma-1}, \tag{111}$$

where $\xi^I > 0$ is a cost shifter and $\alpha \in (0,1)$ governs the curvature of R&D costs. Upon success, quality improves multiplicatively,

$$q' = \lambda^I q, \qquad \lambda^I > 1, \tag{112}$$

and the firm resets its price when the upgraded product is introduced.

External innovation is untargeted. An incumbent chooses per-line intensity $x^E > 0$ and pays

$$c^E(x^E, \hat{q}) = \xi^E \left(x^E\right)^{\frac{1}{1-\alpha}} \hat{q}^{\sigma-1}, \tag{113}$$

if the quality of chosen line is $\hat{q}$. Upon success, it acquires the randomly drawn line, displacing the incumbent, upgrades quality to

$$q' = \lambda^E q, \qquad \lambda^E > 1, \tag{114}$$

and resets the price upon takeover.

Entrants behave analogously. An entrant chooses intensity $x^N > 0$ and pays

$$c^N(x^N, \hat{q}) = \xi^N \left(x^N\right)^{\frac{1}{1-\alpha}} \hat{q}^{\sigma-1}, \tag{115}$$

where $\xi^N > 0$. Upon success, it acquires a randomly drawn line, upgrades quality to $\lambda^N q$ with $\lambda^N > 1$, and sets the optimal reset price.

The economy-wide rate of creative destruction is

$$\tau \equiv x^E + x^N. \tag{116}$$

Since each product line undertakes untargeted external innovation at rate x^E and successful entrant innovations arrive at rate x^N, each successful external or entrant innovation transfers ownership of the affected line to the innovator. Hence τ summarizes the arrival rate at which product lines experience ownership turnover in the economy.

Because successive vintages within a line may differ in both quality and price, a standard CES demand system would in principle allow multiple vintages to coexist within a line. To preserve the canonical structure that at most one product is operated per line, we impose a simple selection device. Each product line has a single "slot" (shelf space) controlled by a retailer. The retailer receives an ε share of the operating profit generated by the stocked vintage and therefore chooses the vintage that maximizes the line's profit flow. We take $\varepsilon \to 0$, so the device affects selection but not equilibrium payoffs. Under this device, only one vintage is operated on each line in equilibrium, and when a new (improved) vintage arrives, the previously operated vintage is immediately deactivated.

E.6 Firm Problem

Let $\Omega_{it} = \{(\hat{q}_{jt}, \mu_{jt})\}_{j=1}^{M_{it}}$ denote the set of relative qualities and markups of the product lines operated by incumbent firm i. Because innovation costs are independent across lines, the firm's value is additive across products, as is standard in the literature (Klette and Kortum, 2004). Hence,

$$W_t(\Omega_{it}) = \sum_{j=1}^{M_{it}} V_t(\mu_{jt}, q_{jt}), \tag{117}$$

where $V_t(\mu)$ is the per-line value with markup μ. Each line is subject to exogenous firm exit at rate ψ and creative destruction at rate τ_t. In addition, the firm can undertake internal and external innovation. The (real) value of quality-normalized line then satisfies the time-dependent

quasi-variational inequality

$$0 = \max \Bigg\{ \partial_t V_t(\mu, \hat{q}) - \big(r_t + \tau_t + \psi + \gamma\big) V_t(\mu, \hat{q}) - \bar{\pi} \mu \frac{\partial V_t(\mu, \hat{q})}{\partial \mu} - g_t^Q \hat{q} \frac{\partial V_t(\mu, \hat{q})}{\partial \hat{q}} + \pi_t(\mu, \hat{q}) + \Lambda_t^I(\mu, \hat{q}) + \Lambda_t^E + \gamma V_t(\mu_t^u) - F_t,\ V_t(\mu_t^u, \hat{q}) - \kappa Y_t \hat{q}^{\sigma-1} - V_t(\mu, \hat{q}) \Bigg\}. \tag{118}$$

where g_t^Q is the growth rate of the quality index Q_t. The first term is the continuation value: markup drifts downward with inflation, the firm earns current flow profit $\pi_t(\mu, \hat{q})$, receives internal-innovation and external-innovation value $\Lambda_t^I(\mu, \hat{q})$, Λ_t^E, pays a fixed cost of operation F_t and obtains free price adjustment opportunities at Poisson rate γ. The second term is the value of paying the menu cost and resetting the markup.

The value of internal innovation $\Lambda^I(\mu, \hat{q})$ follows

$$\max_{x^I} \left\{ x^I \left[V_t(\lambda^I \hat{q}, \mu_t^u) - V_t(\hat{q}, \mu) - \xi^I (x^I)^{\frac{1}{1-\alpha}} \hat{q}^{\sigma-1} Y_t \right] \right\} \tag{119}$$

The value of external innovation $\Lambda^e(\mu, \hat{q})$ follows

$$\max_{x^E} \left\{ x^E \mathbb{E} \left[V_t(\lambda^E \hat{q}, \mu_t^u) - \xi^E (x^E)^{\frac{1}{1-\alpha}} \hat{q}^{\sigma-1} Y_t \right] \right\} \tag{120}$$

To maintain the tractability while preserving the key economic mechanism, we impose a disciplined restriction on the fixed cost of operation.

Assumption 2 (Fixed cost)**.** ***Fixed cost equals the external-innovation option value.*** *The per-period fixed operating cost equals the equilibrium per-line option value of external innovation:*

$$F_t = \Lambda_t^E. \tag{121}$$

The following lemma shows implications of this assumption.

Lemma 2 (Separability and common markup thresholds)**.** *Under Assumption 2, the following properties hold:*

- ***Separation.*** *The value of a product line factorizes as*

$$V_t(\hat{q}, \mu) = \hat{q}^{\sigma-1} A_t(\mu) Y_t \tag{122}$$

  where $A_t(\mu)$ depends only on the markup μ.

- ***Common reset and exit thresholds.*** *The optimal pricing policies are quality-invariant: there exist constants (μ^u, μ^b) such that the optimal reset markup and the exit cutoff are the same for all lines.*

- ***Independence of quality and markups.*** *The relative-quality process (and hence its stationary distribution) is independent of the markup distribution.*

Proof. See Appendix D. ■

Lemma 2 yields important tractability gains. Since operating profits, R&D costs and menu costs scale with $\hat{q}^{\sigma-1}$, the line value is homogeneous in relative quality. We normalize all the growing variables in the value function by Y_t to keep the equilibrium values constant on the BGP. Therefore, the value function factorizes as $V_t(\hat{q}, \mu) = \hat{q}^{\sigma-1} A_t(\mu)$. The dynamic pricing problem therefore reduces to the one-dimensional function $A(\mu)$, and the associated policies are *quality-invariant*: the optimal reset markup μ_t^u and adjustment markup μ_t^b are common across lines.

Define the effective discount rate $\phi_t \equiv r_t + \psi + \tau_t + (\sigma - 1) g_t^Q$. The QVI for $A_t(\mu)$ is

$$0 = \max \Bigg\{ -(\phi_t + \gamma - g_t^Y) A_t(\mu) + \partial_t A_t(\mu) - \bar{\pi} \mu \frac{\partial A_t(\mu)}{\partial \mu} + \tilde{\pi}_t(\mu) + \tilde{\Lambda}^I(\mu) + \gamma A_t(\mu_t^u),\ A_t(\mu_t^u) - \kappa - A_t(\mu) \Bigg\} \tag{123}$$

where g_t^Y is the growth rate of output, $\pi_t(\mu, \hat{q}) = \tilde{\pi}_t(\mu) \hat{q}^{\sigma-1} Y_t$ and $\Lambda^I(\mu, \hat{q}) = \tilde{\Lambda}^I(\mu) \hat{q}^{\sigma-1} Y_t$ The optimal reset markup solves

$$\mu_t^u \in \arg\max_{\mu'} A_t(\mu'). \tag{124}$$

Let μ_t^b denote the lower adjustment threshold, defined by

$$A_t(\mu_t^b) = A_t(\mu_t^u) - \kappa. \tag{125}$$

That is, once the markup falls to μ_t^b, it becomes optimal to pay the menu cost and reset to μ_t^u.

The per-line option value of internal innovation and the associated optimal innovation rate are

$$\tilde{\Lambda}_t^I(\mu) = \xi^I \frac{\alpha}{1-\alpha} \left(x_t^I(\mu) \right)^{\frac{1}{1-\alpha}} Y_t, \qquad x_t^I(\mu) = \left[\frac{1-\alpha}{\xi^I} \Big(A_t(\mu_t^u)(\lambda^I)^{\sigma-1} - A_t(\mu) \Big) \right]^{\frac{1-\alpha}{\alpha}}. \tag{126}$$

The optimal internal-innovation rate is increasing in the value gap between the current line and the reset state. When the current markup equals the reset markup, the innovation incentive is only driven by higher demand from the quality improvement.

The per-line option value of external innovation and optimal external-innovation rate equal

$$\tilde{\Lambda}_t^E = \xi^E \frac{\alpha}{1-\alpha} \left(x_t^E \right)^{\frac{1}{1-\alpha}}, \qquad x_t^E = \left(\frac{1-\alpha}{\xi^E} A(\mu_t^u)(\lambda^E)^{\sigma-1} \right)^{\frac{1-\alpha}{\alpha}}. \tag{127}$$

Since the external innovation is untargeted, it does not depend on current markup of this product.

E.7 Distribution Dynamics

Let $f_t(\mu)$ denote the cross-sectional distribution of markups on $[\mu^b, \mu_t^u]$. Between resets, inflation erodes markups at rate $\bar{\pi}\mu$, so the distribution drifts leftward. Firms also reset through free adjustment, entry, external innovation, and internal innovation, with interior reset hazard

$$\kappa_t(\mu) = \psi + \gamma + x_t^E + x_t^N + x_t^I(\mu).$$

Menu-cost adjustment works at the boundary μ_t^b. The Kolmogorov forward equation becomes

$$\begin{aligned}\partial_t f_t(\mu) = &-\partial_\mu\big(\bar{\pi}\mu\, f_t(\mu)\big) - \lambda_t^{\text{int}}(\mu) f_t(\mu) \\ &+ \delta(\mu - \mu_t^u)\left[\int_{\mu^b}^{\mu_t^u} \kappa_t(\tilde{\mu}) f_t(\tilde{\mu})\, d\tilde{\mu} + \bar{\pi}\mu_t^b f_t(\mu_t^b)\right], \qquad \mu \in [\mu_t^b, \mu_t^u]. \end{aligned} \tag{128}$$

The first term captures leftward drift due to inflation, the second term captures interior resets, and the last term collects all firms that either reset in the interior or hit the lower bound μ_t^b and therefore jump back to the reset markup μ_t^u. In particular, the term $\bar{\pi}\mu_t^b f_t(\mu_t^b)$ is the boundary flux of firms reaching the lower threshold and adjusting immediately.

Using the markup distribution and given the independence between quality and markup, the aggregate markup m_t is given by

$$m_t = \left(\int \mu^{1-\sigma} f_t(\mu) d\mu\right)^{\frac{1}{1-\sigma}}. \tag{129}$$

The aggregate efficiency v_t is

$$v_t \equiv \frac{m_t^{-\sigma}}{\int \mu_{it}^{-\sigma} f_t(\mu) d\mu}. \tag{130}$$

E.8 Entrants, Growth Rate and Consumption Share

There is a unit measure of potential entrants, which are brand-new firms in the data. Each entrant can innovate by introducing a new variety. They first randomly draw a quality from the quality distribution of existing products. Once they innovate successfully, they improve such quality by λ^N. The optimization problem for entrants is given by,

$$\max_{x^N} \left\{ x^N \mathbb{E}\left[V_t(\lambda^N \hat{q}, \mu^u) - \xi^N (x^N)^{\frac{1}{1-\alpha}} \hat{q}^{\sigma-1} Y_t \right] \right\} \tag{131}$$

where $E\Gamma$ is the expected value of entry. The optimal arrival rate of innovation by entrants is,

$$x^N = \left(\frac{1-\alpha}{\xi^N} A(\mu^u)(\lambda^N)^{\sigma-1}\right)^{\frac{1-\alpha}{\alpha}} \tag{132}$$

The goods market clearing implies

$$Y_t = C_t + D_t + F_t N_t + \kappa Y_t \pi_t \mu_t^b f_t(\mu_t^b), \tag{133}$$

where D_t is the aggregate R&D cost, $F_t N_t$ is the aggregate fixed cost, and $\pi_t \mu_t^b f_t(\mu_t^b)$ is the flow measure of product lines that hit the lower adjustment threshold and pay the menu cost.. The growth rate of quality index is driven by internal, external and entrant innovations.

Proposition 7. *The growth rate of the quality index is equal to*

$$g_t^Q = x_t^E \frac{(1+\lambda^E)^{\sigma-1}-1}{\sigma-1} + x_t^I \frac{(1+\lambda^I)^{\sigma-1}-1}{\sigma-1} + x_t^N \frac{(1+\lambda^N)^{\sigma-1}-1}{\sigma-1} \tag{134}$$

The share of consumption of total final goods,

$$\chi_c = 1 - \left(\xi^I \int \left(x_t^I(\mu)\right)^{\frac{1}{1-\alpha}} f_t(\mu)\, d\mu + \frac{1}{1-\alpha} \xi^E \left(x_t^E\right)^{\frac{1}{1-\alpha}} + \xi^N \left(x_t^N\right)^{\frac{1}{1-\alpha}} \right) - \kappa \int a_t(\mu) f_t(\mu)\, d\mu. \tag{135}$$

Proof. See Appendix D. ■

E.9 Equilibrium

The equilibrium of the above economy is as follows:

Definition 2 (Equilibrium)**.** *An equilibrium is a collection of paths*

$$\left\{ A_t(\mu), f_t(\mu), \mu_t^u, \mu_t^b, x_t^I(\mu), x_t^E, x_t^N, m_t, v_t, Y_t, C_t, \hat{P}_t, \pi_t, r_t \right\}_{t \geq 0}$$

such that (i) firms solve the pricing and innovation problem in (123)*; (ii) the markup distribution evolves according to* (**??**)*; (iii) aggregate growth satisfies the law of motion for Q_t in* (134)*; (iv) the resource constraint* (135) *holds; and (v) the nominal block satisfies* (97)*.*

Stationary equilibrium. A stationary equilibrium is a special case of (123) in which aggregate objects and the value function are time-invariant after scaling. Hence

$$\partial_t A_t(\mu) = 0, \qquad \partial_t f_t(\mu) = 0,$$

with constant inflation, constant innovation rates, constant growth rate, and a stationary markup distribution. We use notation without the time subscript t to denote stationary values.

E.10 Equilibrium Characterization

In this section, we first characterize the balance growth path. We characterize the equilibrium aggregates, the reset and lower adjustment markups and markup distribution. We also consider welfare comparison across long-run equilibria. Second, we characterize the equilibrium following unanticipated one-time permanent shocks in productivity and money supply.

E.11 Characterization of BGP

Equilibrium Aggregates. In the spirit of Golosov and Lucas (2007), the household preference and independence between markup and quality leads to the following proposition that characterizes equilibrium aggregates along the balanced growth path.

Proposition 8. *On the BGP, equilibrium aggregates satisfy*

$$R = \rho + \bar{\pi}, \quad W = \chi^{-1} RM, \quad L_t = L, \quad \pi = \bar{\pi}, \quad g^Y = g^Q = g, \quad \frac{\dot{\hat{P}}_t}{\hat{P}_t} = \pi - g, \tag{136}$$

where $\bar{\pi}$ denotes the money growth rate.

The proposition implies that balanced-growth output growth is entirely driven by growth in aggregate product quality. Inflation equals money growth, and nominal wages grow at the inflation rate. Hence the real wage $W/\hat{P}_t$ grows at rate g. Finally, under log utility, the income and substitution effects on labor supply exactly offset each other, implying a constant labor supply along the BGP.

Optimal Reset Markup and Markup Distribution. So far, we have characterized the value function taking the reset markup μ^u as given. The next proposition pins down the optimal reset markup and its main implications.

Proposition 9. *The optimal reset markup μ^u satisfies*

$$(\phi - g^Q) A(\mu^u) = m^{\sigma-1}(\mu^u - 1)(\mu^u)^{-\sigma} + \tilde{\Lambda}^I(\mu^u). \tag{137}$$

Moreover, the reset markup exhibits price overshooting*:*

$$\mu^u > \mu^* = \frac{\sigma}{\sigma - 1}.$$

As pricing frictions vanish, either because $\pi \to 0$ or because $\gamma \to \infty$, the reset markup converges to the frictionless benchmark:

$$\mu^u \to \mu^*.$$

Proof. See Appendix D. ■

Evaluating the HJB at $\mu = \mu^u$ given $A'(\mu^u) = 0$ yields (137). Although the option value of innovation depends on μ^u, it does not affect the marginal reset condition. Since μ^u is an interior maximizer of A, we have $A'(\mu^u) = 0$, which implies $\tilde{\Lambda}^{I\prime}(\mu^u) = 0$.

Differentiating the HJB with respect to μ and evaluating at $\mu = \mu^u$:

$$m^{\sigma-1}(\sigma - 1)\mu_u^{-\sigma-1}\left(\mu^u - \frac{\sigma}{\sigma - 1}\right) = -\pi \mu^u A''(\mu^u). \tag{138}$$

Because μ^u maximizes A, we have $A''(\mu^u) < 0$. Hence, whenever $\pi > 0$, the right-hand side of (138) is strictly positive, implying $\mu^u > \mu^*$. Intuitively, after a reset, inflation erodes the markup over time. Firms therefore choose a markup above the static optimum so that the expected markup over the duration of a price spell remains closer to the frictionless target.

In the limits $\pi \to 0$ or $\gamma \to \infty$, this dynamic consideration disappears. In the former case, the markup no longer drifts after a reset; in the latter, firms can reoptimize arbitrarily frequently. In either case, the reset problem collapses to the static pricing problem, and the optimal reset markup converges to μ^*.

The following proposition characterizes the markup distribution and its properties.

Proposition 10. *Let the stationary markup density on $[\mu^b, \mu^u]$ satisfy*

$$f(\mu) = \frac{1}{Z}\frac{1}{\mu}\exp\left(-\int_{\mu}^{\mu^u} \frac{\kappa(z)}{\bar{\pi} z}\, dz\right), \quad Z = \int_{\mu^b}^{\mu^u} \frac{1}{z}\exp\left(-\int_{z}^{\mu^u} \frac{\kappa(s)}{\bar{\pi} s}\, ds\right) dz. \tag{139}$$

where $\kappa(\mu) = \psi + \gamma + x^E + x^N + x^I(\mu)$. Then:

(i) if $\frac{\kappa_1(\mu)}{\bar{\pi}_1} \geq \frac{\kappa_0(\mu)}{\bar{\pi}_0}$ for all $\mu \in [\mu^b, \mu^u]$, the stationary distribution under κ_1 first-order stochastically dominates that under κ_0;

(ii) if $\inf_\mu \kappa(\mu)/\pi \to \infty$, then the stationary distribution converges weakly to a point mass at μ^*.

Intuitively, lower inflation implies that markups drift away from μ^u more slowly, while a higher reset hazard, through more frequent price adjustment or innovation, increases the share of products whose markups remain close to μ^u. In addition, Since $\kappa(\mu)$ is decreasing in μ, low-markup products are reset more aggressively than high-markup products, so the stationary distribution is tilted more toward μ^u.

More formally, a pointwise increase in $\kappa(\mu)$ raises $\kappa(\mu)$ for all $\mu \in (\mu^b, \mu^u)$, reduces the relative mass in low-markup states, and shifts the distribution toward μ^u in the sense of first-order stochastic dominance. Likewise, a lower inflation rate π has the same effect. As a result, markups become less dispersed, misallocation is mitigated, and aggregate efficiency v rises. In the limit, the distribution converges to a mass point at μ^* and the misallocation becomes zero and $v = 1$.

Long-Run Welfare. Following Acemoglu et al. (2018), we are interested in the welfare comparison between long-run equilibrium of different economies. Recall that consumption is a constant share of aggregate output, denoted by χ_c. To compare across economies, we set $Q_0 = 1$ and $A = 1$. The initial consumption is therefore $C_0 = \chi_c v$. We ignore the utility from money since it only serves as a device to introduce money demand. Then the welfare of the above economy can be obtained as,

$$U_0(v, g) = \int_0^\infty e^{-\rho t}\big[\log C_t - L\big]\, dt = \underbrace{\frac{1}{\rho}\log \chi_c + \frac{g}{\rho^2}}_{\text{growth-related}} \underbrace{+\frac{1}{\rho}\log v}_{\text{agg. efficiency}} \underbrace{-\frac{L}{\rho}}_{\text{labor disutility}}.$$

Welfare in this economy has three components: growth, misallocation, and labor disutility. Higher growth reflects faster innovation, but requires more R&D and thus lowers current consumption, introducing a tradeoff. Markup dispersion reduces welfare by distorting resource allocation across firms and lowering aggregate efficiency. Labor disutility also enters welfare, but labor supply is the same across the equilibria considered here.[14]

In comparing welfare across two economies, indexed by 1 and 2, with growth rates g_1 and g_2 and initial consumption levels C_0^1 and C_0^2, we measure the consumption-equivalent welfare change by the fraction κ of initial consumption in economy 2 that makes the representative household indifferent between the two economies. Formally, κ is defined by

$$U_0(\kappa C_0^2, g_2) = U_0(C_0^1, g_1).$$

Thus, we can decompose the log consumption-equivalent welfare change as

$$\log \kappa = \underbrace{\frac{g_1 - g_2}{\rho} + \log\left(\frac{\chi_{c1}}{\chi_{c2}}\right)}_{\text{growth-related}} + \underbrace{\log\left(\frac{v_1}{v_2}\right)}_{\text{misallocation}}. \tag{140}$$

[14] With linear labor disutility, labor supply is not uniquely pinned down. We select the equilibrium labor supply consistent with the limit as $\eta \to 0$.