Precautionary Pricing and Markup Cyclicality*

Thomas Hasenzagl — FRB of Richmond
Luis Pérez — SMU

PRELIMINARY AND INCOMPLETE

[Click here for the latest version]

This version: April 28, 2026.
First version: January 30, 2026.

Abstract

Why are markups cyclical, and why does their cyclicality vary across firms, sectors, and business cycle episodes? This paper develops a theory of pricing under uncertainty in which risk-averse firms set prices before costs and demand are realized. In this environment, pricing serves a precautionary role: optimal markups incorporate a risk premium that reflects the value of limiting exposure to uninsurable cost and demand risk. A central implication is that markup dynamics depend on the composition of risk firms face. Cost risk and demand risk affect pricing incentives in systematically different ways, so whether markups are procyclical or countercyclical depends on which source of risk dominates and how it co-moves with economic activity. We test these predictions using industry-level data for U.S. manufacturing over 1976–2018. Consistent with the theory, a one-standard-deviation increase in cost uncertainty raises prices by 5.8%, while a comparable increase in demand uncertainty lowers prices by 6.6%. The structural calibration of the dynamic GE model and the corresponding quantitative exercises—impulse responses to uncertainty shocks, welfare costs, the historical 1976–2018 path, and untargeted cross-industry validation—are work in progress and will appear in a subsequent draft.

Keywords: Markups, Pricing under Risk, Uncertainty, Business Cycles.
JEL codes: E32, E31, D82, L11.

*We thank Anmol Bhandari, V. V. Chari, Loukas Karabarbounis, and Chris Phelan for valuable comments, as well as seminar participants at SMU. All errors are our own. Emails: thomas.hasenzagl@gmail.com, luisperez@smu.edu.

1 Introduction

Why are markups cyclical, and why does their cyclicality differ across firms, sectors, and business-cycle episodes? Empirical findings are strikingly heterogeneous: in some contexts markups appear countercyclical, in others procyclical, or even close to acyclical. Despite broad agreement on the cyclical behavior of consumption, investment, and employment, there is no consensus on the cyclical behavior of markups.

This paper reconciles these conflicting findings. We argue that markup cyclicality reflects differences in the composition of risks firms face. Firms that set prices under uncertainty use pricing to manage their exposure to risk. When costs are uncertain, a higher price limits exposure to high-cost states; when demand is uncertain, a lower price sustains sales in weak-demand states. Cost risk therefore raises markups, and demand risk lowers them. Whether markups rise or fall over the business cycle depends on which type of risk dominates and how that risk covaries with economic activity. This paper develops and tests a theory of *markup cyclicality* built on this mechanism.

The dominant interpretation of markups in macroeconomics emphasizes market power. In models of imperfect competition, firms set prices above marginal cost because they face downward-sloping demand curves (Chamberlin, 1933; Dixit and Stiglitz, 1977). A large literature studies how such markups affect production, misallocation, inflation, and welfare, and how they respond to realized movements in demand and costs (e.g., see Blanchard and Kiyotaki, 1987; Edmond, Midrigan and Xu, 2023). In this view, markup fluctuations reflect realized shocks or changes in competitive conditions.[1]

A distinct explanation for markups emphasizes risk. Since Knight (1921), economists have argued that profits compensate firms for bearing uncertainty that cannot be fully insured or contracted away. Modern risk-based models formalize this insight by showing that uninsurable productivity risk distorts input choices and generates wedges between prices and marginal costs (Boar, Gorea and Midrigan, 2025; Di Tella, Malgieri and Tonetti, 2025). In these frameworks, markups emerge indirectly from distorted real allocations; pricing itself is not the object of analysis.

We instead analyze pricing as the locus of risk management. When firms set prices before uncertainty is resolved, risk aversion introduces a trade-off between expected profits and profit volatility directly in the pricing decision. Under this mechanism, markups do not solely reflect rents from market power. Instead, optimal pricing incorporates a risk premium that reflects the firm's exposure to profit risk. Market

[1] Throughout, we use "markups" in the broad macroeconomic sense of wedges between prices and marginal costs. Such wedges can reflect market power, but may also arise from a variety of other frictions—including nominal rigidities and price adjustment costs, search and matching frictions, customer markets and price discrimination, and informational frictions. See, for example, Calvo (1983); Golosov and Lucas (2007); Atkeson and Burstein (2008); Menzio (2024); Angeletos and La'O (2013).

power and other pricing frictions determine the risk-neutral component of the markup, while the risk premium incorporates the precautionary value of limiting exposure to uninsurable risk. Understanding markup cyclicality therefore requires analyzing optimal pricing under uncertainty.

We develop a general theory of pricing under uncertainty in which risk-averse firms set prices ex ante, before costs and demand are realized, and quantities adjust endogenously to meet demand. Optimal pricing decomposes into two components: a risk-neutral markup—reflecting market power and other familiar pricing forces—and a risk premium that captures the precautionary value of pricing under uncertainty. This decomposition isolates the effect of risk from other determinants of markups and delivers sharp sign restrictions across distinct sources of uncertainty. Cost risk increases risk-adjusted marginal cost and raises markups. Demand risk reduces risk-adjusted marginal revenue and lowers markups.

Main Contributions. Our first contribution is theoretical. Under weak assumptions on preferences and demand, we establish these sign restrictions without relying on particular functional forms. Our focus is not on explaining the level of markups, but on their variation over the business cycle. The theory shows that markup cyclicality is an equilibrium outcome determined by the composition of risk firms face. This framework offers a coherent interpretation of the mixed empirical evidence on markup cyclicality documented in the literature (see, for example, Bils, 1987; Hall, 1988; Basu and Fernald, 1997; Nekarda and Ramey, 2020; De Loecker, Eeckhout and Unger, 2020; Hasenzagl and Pérez, 2023), which standard market-power models and existing risk-based theories struggle to reconcile.

Our second contribution is empirical. We test the theory's predictions using industry-level data for U.S. manufacturing covering 335 6-digit NAICS industries (matched to 209 of 231 manufacturing detail input-output codes) over 1976–2018. We construct Bartik-style shifters from the BEA input-output tables that measure each industry's exposure to upstream cost conditions and downstream demand conditions through its position in the production network. Uncertainty is measured as the rolling volatility of innovations to these shifters. A key feature of our empirical design is the inclusion of Bartik level controls—lagged levels of the cost and demand shifters—which absorb mechanical cost pass-through and ensure that the estimated coefficients capture the pure uncertainty channel. Consistent with the theory, we find that cost uncertainty significantly raises prices while demand uncertainty lowers them: in our preferred 2SLS specification, a one-standard-deviation increase in cost uncertainty raises prices by approximately 5.8%, while a comparable increase in demand uncertainty lowers prices by approximately 6.6%.

The reduced-form slopes identify the precautionary pricing channel at the industry level: cost uncertainty raises prices and demand uncertainty lowers them, with both channels well-identified and quantitatively meaningful. Mapping these reduced-form slopes into a quantitative decomposition of aggregate markup cyclicality requires a defensible firm-level markup measure and a fully calibrated structural model, neither of which we attempt in this draft. The structural calibration of the dynamic GE model—and the associated quantitative exercises (impulse responses to uncertainty shocks, the welfare cost of uncertainty, the historical 1976–2018 path, and untargeted cross-industry validation)—are work in progress and will appear in a subsequent draft.

Related Literature. Our paper relates to two strands of the literature. First, as discussed above, recent work shows that uninsurable risk can generate wedges commonly interpreted as markups (Boar, Gorea and Midrigan, 2025; Di Tella, Malgieri and Tonetti, 2025). In these models, imperfectly diversified entrepreneurs choose inputs—such as labor or capital—before productivity is realized, and risk aversion distorts these choices away from the risk-neutral benchmark, adding a risk component to price–cost wedges.

We differ in two fundamental respects. First, we study risk management through pricing rather than input choice: firms in our model set prices ex ante and quantities adjust to meet demand, so markups arise directly from optimal pricing under uncertainty—consistent with evidence that firms predominantly adjust prices rather than quantities in response to shocks (Flynn, Nikolakoudis and Sastry, 2024). Second, we allow for multiple sources of risk with systematically different effects on pricing incentives. While existing work focuses on productivity or cost risk, which pushes pricing distortions in a single direction, our framework accommodates both cost and demand risk and shows how their interaction determines markup cyclicality. Our model generates a risk component of markups similar to that obtained in models where risk distorts input choices when demand risk is absent and cost uncertainty arises solely from productivity shocks.

Second, our analysis complements recent work by Burstein, Carvalho and Grassi (2025). They develop a granular oligopolistic model in which firm-level shocks generate sectoral and aggregate markup dynamics through endogenous changes in market shares and demand elasticities. In their framework, markups are functions of firm size: as firms expand or contract following shocks, they move along a markup–size schedule implied by variable demand curvature. Shocks that alter firms' relative positions within a sector therefore generate both within-firm markup adjustments—through changes in the elasticity each firm faces—and aggregate markup movements through reallocation across heterogeneous firms. Empirically, they decompose markup fluctuations into within-firm and reallocation components and find that within-firm markup changes account for roughly sixty percent of sectoral markup movements for the median sector in French data. We provide a complementary structural mechanism for within-firm adjustments. In our framework, markups change because firms use prices to manage risk exposure through precautionary pricing, so aggregate cyclicality reflects both precautionary pricing and the reallocation forces emphasized in their model.

The remainder of the paper is organized as follows. Section 2 presents a simple example illustrating the core mechanism. Section 3 develops the general theory of pricing under risk and markup cyclicality. Section 4 embeds the static theory in a dynamic general-equilibrium model. Section 5 tests the qualitative predictions of the theory using U.S. manufacturing data. Sections 6 and 7 outline the structural calibration of the dynamic GE model and the associated quantitative exercises, both of which are work in progress and will appear in a subsequent draft. Section 8 concludes.

2 Precautionary Pricing: An Example

This section illustrates the core mechanism behind *precautionary pricing under uncertainty* in a tractable environment that delivers closed-form solutions and clear intuition. We focus on a special case with CARA preferences, linear demand, constant marginal cost, and normally distributed shocks. These assumptions yield an exact certainty-equivalent representation of the firm's pricing problem and make transparent how different sources of uncertainty distort pricing relative to the risk-neutral benchmark.

This environment shuts down wealth effects and isolates the risk-management role of pricing. Prices are the firm's only instrument for managing profit risk, so deviations from the risk-neutral benchmark reflect incentives to stabilize profits across states of the world. This knife-edge case therefore provides a clean prototype for the general theory developed in the next section, where we relax functional-form assumptions and allow for endogenous wealth effects.

Environment. A monopolistic firm sets its price p before demand and cost shocks are realized.[2] Demand is linear and subject to level uncertainty,

$$y(p;\xi) = a + \xi - bp,$$

[2]Appendix A provides a directed-search microfoundation for price commitment in which demand uncertainty arises endogenously from buyer arrival risk. Cost uncertainty is standard and reflects risk in input prices or own productivity risk.

where $a, b > 0$, and ξ captures demand risk. Marginal cost is constant in output but stochastic,

$$\kappa(\eta) = \overline{\kappa} + \eta,$$

where $\overline{\kappa} > 0$ and η is a cost shock. The shocks $\boldsymbol{\theta} = (\eta, \xi)$ are mutually independent, mean zero, and normally distributed.

Profits are given by

$$\pi(p; \boldsymbol{\theta}) = (p - \overline{\kappa} - \eta)(a + \xi - bp),$$

which can be decomposed as

$$\pi(p; \boldsymbol{\theta}) = \underbrace{(p - \overline{\kappa})(a - bp)}_{\text{deterministic}} + \underbrace{(p - \overline{\kappa})\xi - (a - bp)\eta - \eta\xi}_{\text{stochastic}}.$$

The firm is owned by a risk-averse entrepreneur with no outside income. Preferences over profits are of the CARA form,

$$U(\pi) = -\exp(-\alpha\pi), \qquad \alpha \geq 0.$$

Pricing problem. The firm chooses price before shocks are realized to maximize expected utility,

$$\max_{p>0} \quad \mathbb{E}[-\exp(-\alpha\pi(p))].$$

We first characterize the risk-neutral benchmark ($\alpha = 0$), and then show how risk aversion distorts pricing away from this benchmark.

Risk-neutral benchmark. Under risk neutrality, the optimal price satisfies

$$\frac{\mathrm{d}}{\mathrm{d}p}\mathbb{E}[\pi(p)] = a - 2bp + b\overline{\kappa} = 0,$$

which yields

$$p_{\mathrm{RN}} = \frac{a + b\overline{\kappa}}{2b}.$$

At this price, expected marginal profits are zero. As a result, a marginal change in price has no first-order effect on mean profits. Any deviation from p_{RN} under risk aversion must therefore arise entirely from how price changes redistribute profits across states of the world, rather than from changes in expected profits.

Equivalently, the optimality condition can be written in output space as

$$\frac{a - 2y}{b} = \overline{\kappa},$$

which is the familiar condition "mean marginal revenue equals mean marginal cost" shown in Figure 1.

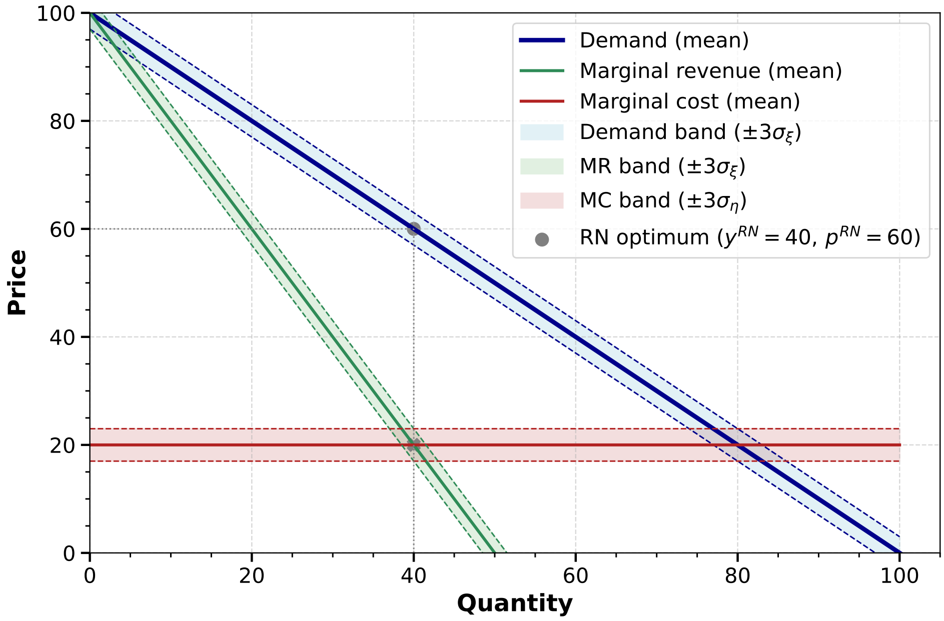

Figure 1: Optimal risk-neutral pricing.

Notes. Linear demand $y = a - bp$ with $a = 100$ and $b = 1$. Inverse demand is $p(y) = 100 - y$, marginal revenue is $\text{MR}(y) = 100 - 2y$, and marginal cost is $\overline{\kappa} = 20$. Both bands are plotted with $\pm 3\sigma_i$ from their mean value, with $\sigma_i = 1$ for $i \in \{\eta, \xi\}$.

Risk-adjusted pricing. With CARA utility and normally distributed profits, maximizing expected utility is equivalent to maximizing the certainty equivalent,

$$\max_{p>0} \quad \text{CE}(p) = \mathbb{E}[\pi(p)] - \frac{\alpha}{2}\text{Var}\big(\pi(p)\big).$$

This formulation makes the risk-management role of pricing explicit: prices are chosen not only to maximize expected profits but also to manage the dispersion of profits across states of the world. The optimal price satisfies

$$\frac{\mathrm{d}}{\mathrm{d}p}\mathbb{E}[\pi(p)] - \frac{\alpha}{2}\frac{\mathrm{d}}{\mathrm{d}p}\text{Var}(\pi(p)) = 0.$$

The second term captures the *risk adjustment*. Locally around the risk-neutral benchmark p_{RN}, pricing responds exclusively to how marginal price changes affect profit volatility. Whether risk raises or lowers prices depends entirely on the source of uncertainty.

2.1 Cost risk

Suppose only marginal cost is uncertain ($\xi = 0$). Profits reduce to

$$\pi(p;\eta) = (p - \overline{\kappa} - \eta)(a - bp),$$

where η is a mean zero shock with variance σ_η^2, and profit volatility depends on the scale of production,

$$\text{Var}(\pi(p)) = (a - bp)^2\sigma_\eta^2 = y^2\sigma_\eta^2,$$

so higher output amplifies exposure to cost risk.

The resulting pricing condition can be written as

$$(a - 2bp + b\overline{\kappa}) + \alpha b(a - bp)\sigma_\eta^2 = 0,$$

where the first term corresponds to the standard risk-neutral pricing condition and the second term captures the effect of cost risk under risk aversion. Relative to the risk-neutral benchmark, cost risk induces an upward distortion in price:

$$p_{\text{RA}} \approx p_{\text{RN}} + \frac{\alpha}{2} y_{\text{RN}}\sigma_\eta^2.$$

Cost risk raises price. Intuitively, when marginal costs are uncertain, profits are most volatile at high output levels, where adverse cost realizations are especially costly. By raising prices, the firm reduces quantity demanded and dampens its exposure to high-cost outcomes. Pricing therefore limits the firm's exposure to cost risk and generates higher markups.

The same intuition can be expressed in output space. Using inverse demand, the optimality condition can be rewritten as

$$\underbrace{\frac{a - 2y}{b}}_{\text{mean MR}(y)} = \underbrace{\overline{\kappa} + \alpha y\sigma_\eta^2}_{\text{risk-adjusted MC}(y)}.$$

Cost risk thus enters the firm's problem as an endogenous, output-dependent wedge: from the firm's perspective, pricing under uncertainty is equivalent to facing an increasing marginal cost schedule, even though technological marginal cost is constant. At the optimum, risk-adjusted marginal cost equals mean marginal revenue (see Figure 2).

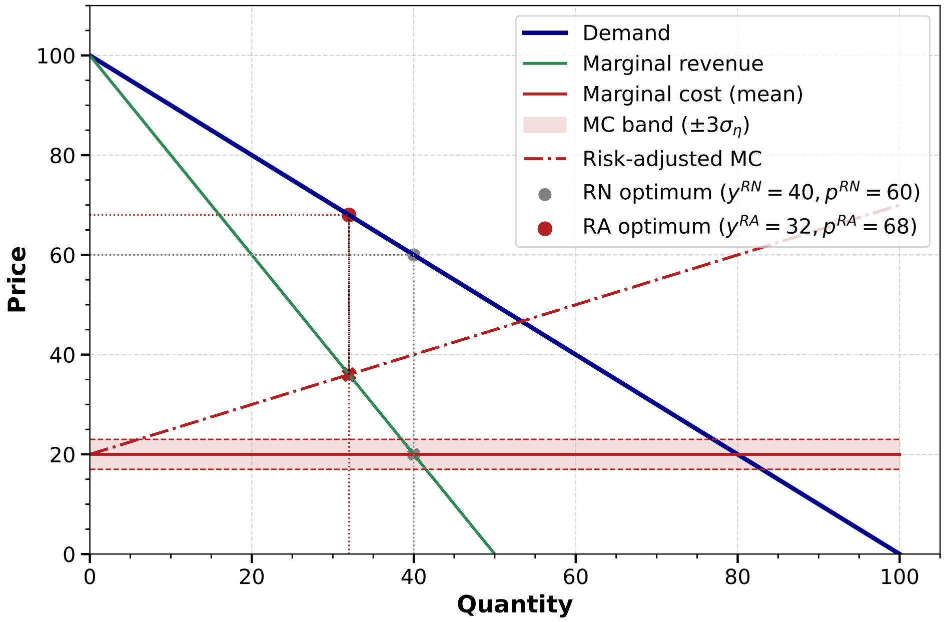

Figure 2: Optimal risk-averse pricing with cost risk.

Notes. This figure corresponds to linear demand $y = a - bp$, where $a = 100$ and $b = 1$. In this case, inverse demand is $p(y) = 100 - y$, marginal revenue is $MR(y) = 100 - 2y$, and risk-adjusted marginal cost is $\overline{\kappa} + \alpha y \sigma_\eta^2$, where $\overline{\kappa} = 20$, $\alpha = 1/2$ and $\sigma_\eta^2 = 1$. The marginal-cost band plotted in the figure is $\overline{\kappa} \pm 3\sigma_\eta$.

This characterization provides a formal foundation for Knight (1921)'s theory of profit. Knight argued that under perfect competition and no uncertainty, price equals marginal cost and profits are zero, and that positive profits arise only as compensation for bearing uninsurable risk. Our result makes this precise. In our setting, a risk-averse firm facing cost uncertainty sets price above expected marginal cost even in the absence of market power. In both cases—pricing under risk and pricing with market power—equilibrium output is lower than in the competitive, risk-neutral benchmark, but for fundamentally different reasons. Under market power, higher prices reflect rents from imperfect competition; under cost risk, higher prices reflect a risk premium that compensates for exposure to adverse cost realizations. The resulting profit is therefore not a rent from restricting supply, but a risk premium required to induce production under uncertainty.

2.2 Demand risk

Suppose instead that only demand is uncertain ($\eta = 0$). Firm profits are

$$\pi(p; \xi) = (p - \overline{\kappa})(a + \xi - bp),$$

where ξ is a mean-zero demand shock with variance σ_ξ^2. Profit volatility depends on the wedge between price and expected marginal cost,

$$\text{Var}(\pi(p)) = (p - \overline{\kappa})^2 \sigma_\xi^2,$$

so larger markup wedges $(p - \overline{\kappa})$ amplify exposure to demand risk.

The resulting pricing condition can be written as

$$(a - 2bp + b\overline{\kappa}) - \alpha(p - \overline{\kappa})\sigma_\xi^2 = 0.$$

Relative to the risk-neutral benchmark, demand risk induces a downward distortion in price. A first-order expansion around the risk-neutral price yields

$$p_{\text{RA}} \approx p_{\text{RN}} - \frac{\alpha}{2b}(p_{\text{RN}} - \overline{\kappa})\sigma_\xi^2.$$

Demand risk lowers prices. When demand is uncertain, a higher price—and hence a higher markup—amplifies profit volatility, since revenues then respond more strongly to demand realizations. A risk-averse firm therefore lowers its price ex ante, compressing the markup and reducing the sensitivity of profits to demand shocks. Pricing therefore limits the firm's exposure to demand risk, thereby stabilizing profits across demand realizations.

This mechanism is particularly transparent in output space. Using inverse demand $p(y) = (a - y)/b$, the optimality condition can be rewritten as

$$\underbrace{\frac{a - 2y}{b} + \frac{\alpha}{b}(p(y) - \overline{\kappa})\sigma_\xi^2}_{\text{risk-adjusted MR}(y)} = \underbrace{\overline{\kappa}.}_{\text{mean MC}}$$

Demand risk enters the firm's problem as an endogenous, output-dependent wedge on marginal revenue. From the firm's perspective, pricing under demand uncertainty is equivalent to facing a downward-shifted marginal revenue schedule, with the adjustment proportional to the markup wedge $p(y) - \overline{\kappa}$. Since this wedge is positive whenever markups are positive, the equality between risk-adjusted marginal revenue and marginal cost is satisfied at higher output and therefore lower price, as illustrated in Figure 3.

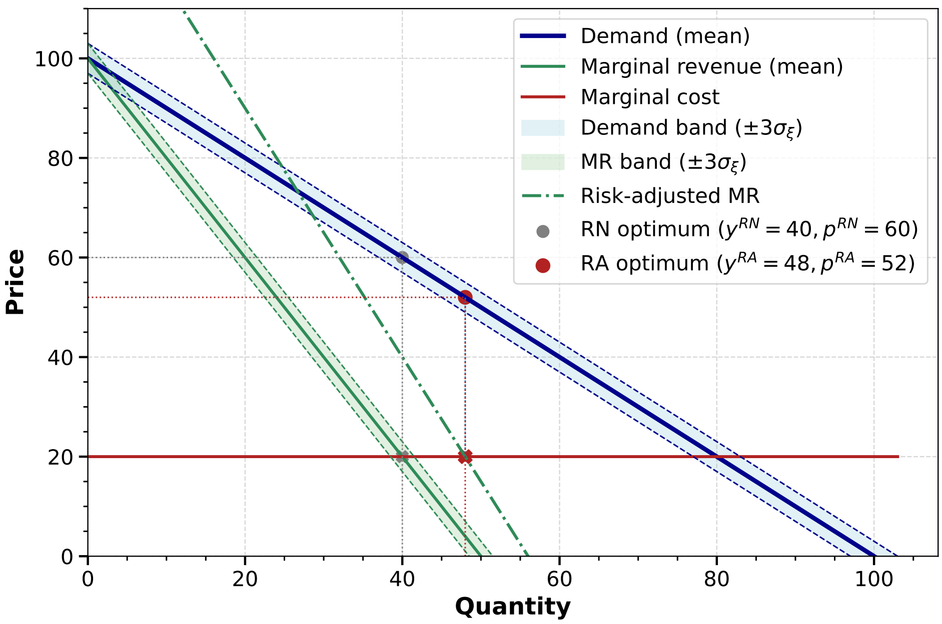

Figure 3: Optimal risk-averse pricing with demand-level risk.

Notes. This figure corresponds to linear demand $y = a - bp$, where $a = 100$ and $b = 1$. In this case, inverse demand is $p(y) = 100 - y$, marginal revenue is $MR(y) = 100 - 2y$, and risk-adjusted marginal revenue is $100 - 2y + \frac{\alpha}{b}(\mu - 1)\sigma_\xi^2$, where $\alpha = 1/2$ and $\sigma_\xi^2 = 1$.

2.3 Markup cyclicality and risk composition

When both sources of risk are active, the optimal price under linear demand has the following first-order approximation around the risk-neutral benchmark:

$$p_{\text{RA}} \approx p_{\text{RN}} + \alpha \mathcal{R}(\boldsymbol{\sigma}^2),$$

where the risk adjustment is given by

$$\mathcal{R}(\boldsymbol{\sigma}^2) = \frac{y_{\text{RN}}}{2}\sigma_\eta^2 - \frac{y_{\text{RN}}}{2b^2}\sigma_\xi^2.$$

The sign and magnitude of $\mathcal{R}(\boldsymbol{\sigma}^2)$ depend on the composition of risk and summarize the central mechanism of precautionary pricing. Different sources of uncertainty distort pricing in opposite directions because they shift profits toward different states of the world and therefore reshape the distribution of profit risk. Cost risk raises prices by increasing risk-adjusted marginal cost: high output magnifies exposure to unfavorable cost realizations, so raising prices reduces quantity demanded and limits exposure to high-cost states. In contrast, demand risk lowers prices by depressing risk-adjusted marginal revenue. Higher prices shift profits toward already favorable demand states, increasing volatility and inducing the firm to compress markups.

To connect these pricing effects to markups, it is useful to distinguish between ex-ante and ex-post notions of markups. The ex-ante markup, chosen by the firm at the time of setting price, is

$$\mu := \frac{p_{\text{RA}}}{\overline{\kappa}} \approx \mu_{\text{RN}} + \frac{\alpha}{\overline{\kappa}}\mathcal{R}(\boldsymbol{\sigma}^2),$$

and captures the firm's pricing decision under uncertainty. By contrast, the ex-post markup, realized after the cost shock is observed, is

$$\mu_{\text{ex-post}}(\eta) := \frac{p_{\text{RA}}}{\overline{\kappa} + \eta} \approx \frac{\overline{\kappa}}{\overline{\kappa} + \eta}\Big[\mu_{\text{RN}} + \frac{\alpha}{\overline{\kappa}}\mathcal{R}(\boldsymbol{\sigma}^2)\Big].$$

Ex-post markups fluctuate mechanically with realized cost shocks, but their expectation coincides with the ex-ante markup, $\mathbb{E}[\mu_{\text{ex-post}}] = \mu$. Our theory concerns the ex-ante markup, which reflects the firm's optimal pricing choice.[3]

Markup cyclicality arises because uncertainty itself varies over the business cycle. In the linear-demand environment, changes in uncertainty shift the firm's optimal price relative to marginal cost through the same risk-adjustment channel that governs precautionary pricing. If cost uncertainty rises in downturns, firms raise prices to reduce exposure to adverse cost realizations, generating countercyclical markups. If instead demand uncertainty rises in downturns, firms compress markups to stabilize profits, generating procyclical markups. When these forces offset each other, markups may appear acyclical even though firms actively adjust prices in response to risk.

A convenient summary of this logic is given by the decomposition

$$\text{Cov}(\mu_t, y_t) = \alpha \sum_{k\in\{\eta,\xi\}} \underbrace{\frac{\partial \mathcal{R}}{\partial \sigma_k^2}}_{\substack{\text{markup sensitivity}\\ \text{to risk } k}} \times \underbrace{\text{Cov}(\sigma_{kt}^2, y_t)}_{\text{cyclicality of risk } k},$$

which expresses markup cyclicality in terms of exposure—that is, how sensitive markups are to each type of uncertainty—and cyclicality—that is, how uncertainty covaries with economic activity. A central implication of the theory is thus that markup cyclicality reflects the composition of risk and the cyclical behavior of its components.

The linear-demand environment studied here deliberately abstracts from other mechanisms firms may use to manage risk. The entrepreneur has no income beyond firm profits, cannot trade financial assets, and has CARA preferences that eliminate wealth effects. Pricing is therefore the firm's only instrument for managing profit risk. In richer environments, firms may partially manage profit risk through financial markets

[3]Beyond the firm-level identity $\mathbb{E}[\mu_{\text{ex-post}}] = \mu$, the two coincide *exactly* at the cross-sectional aggregate under output-weighted harmonic sales-weighting—see Proposition 5 in Section 3.

or diversification of income sources, and preferences with non-constant absolute risk aversion introduce wealth effects that interact with pricing incentives. The next section shows that these considerations modify—but do not overturn—the central insight: markup cyclicality is determined by the composition of risk and by how each risk component covaries with the business cycle.

3 Precautionary Pricing: A Theory of Markup Cyclicality

Prices can be used to manage exposure to profit risk. By adjusting prices, firms owned by risk-averse entrepreneurs can partially manage profit fluctuations induced by demand and cost shocks. This precautionary pricing motive alters pricing incentives relative to the risk-neutral benchmark: cost uncertainty raises optimal prices, while demand uncertainty lowers them. The direction of markup cyclicality therefore depends on which type of uncertainty predominates over the business cycle.

Throughout, firms are owned by entrepreneurs who choose prices to maximize expected utility. Profits need not be the entrepreneur's only source of wealth, so the risk-management role of pricing depends on how price changes affect the distribution of wealth across states, not merely profits in isolation.

We proceed in three steps. First, we derive general pricing results without functional-form assumptions for preferences, demand, or technology. Second, we study pricing locally around the risk-neutral benchmark. Third, we specialize to canonical environments that link the general theory to the quantitative analysis.

3.1 General pricing under uncertainty

A firm chooses its price before the realization of demand and cost shocks. The firm is owned by an entrepreneur who sets price to maximize expected utility. We characterize optimal pricing in a fully general environment.

Let $\boldsymbol{\theta}$ denote a vector of demand and cost shocks. Demand for the firm's variety is $y(p;\boldsymbol{\theta}) = \mathcal{D}(p;\boldsymbol{\theta})$, total cost is $\mathcal{C}(y;\boldsymbol{\theta})$, and profits are

$$\pi(p;\boldsymbol{\theta}) = p\mathcal{D}(p;\boldsymbol{\theta}) - \mathcal{C}(\mathcal{D}(p;\boldsymbol{\theta});\boldsymbol{\theta}).$$

This profit specification encodes an assumption that we use throughout the paper:

- **A0. Price commitment / demand-determined quantity.** The firm's sole decision is its price p; once p is posted, the firm commits to supply the quantity demanded at that price, $y = \mathcal{D}(p;\boldsymbol{\theta})$. Equivalently, realized output is $\mathcal{F}^{\text{pre}}$-measurable up to the demand-shock realization in $\boldsymbol{\theta}$ and does not respond to cost-shock realizations separately from demand.

Assumption A0 is embedded in the profit function above through the substitution $y = \mathcal{D}(p;\boldsymbol{\theta})$ that appears in both the revenue and cost terms. It is the same price-commitment convention invoked in the dynamic environment of Section 4 and underlies the ex-ante pricing theory that follows. Let $\omega(\boldsymbol{\theta})$ denote the entrepreneur's non-profit resources, which may be stochastic and arbitrarily correlated with demand and cost shocks. Total resources are

$$\mathcal{W}(p;\boldsymbol{\theta}) = \omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta}).$$

In general equilibrium, the entrepreneur purchases consumption at price index $P(\boldsymbol{\theta})$, so utility can be written as $U(\mathcal{W}(p;\boldsymbol{\theta})/P(\boldsymbol{\theta}))$.[4] For notational simplicity, normalize $P(\boldsymbol{\theta}) \equiv 1$ and write the entrepreneur's problem as

$$\max_{p>0} \mathbb{E}_{\boldsymbol{\theta}}\left[U\left(\mathcal{W}(p;\boldsymbol{\theta})\right)\right].$$

Theorem 1 (Optimal pricing under cost and demand uncertainty). *Suppose that:*

A1. $U : \mathbb{R} \to \mathbb{R}$ *is twice continuously differentiable, strictly increasing, and strictly concave.*

A2. For almost all $\boldsymbol{\theta}$*,* $\mathcal{D}(\cdot;\boldsymbol{\theta})$ *is twice continuously differentiable with* $\mathcal{D}_p(p;\boldsymbol{\theta}) < 0$ *for all* $p > 0$*, and* $\mathcal{C}(\cdot;\boldsymbol{\theta})$ *is twice continuously differentiable with* $\mathcal{C}_y(y;\boldsymbol{\theta}) > 0$ *for all* $y > 0$.

A3. For almost all $\boldsymbol{\theta}$*, the profit function* $p \mapsto \pi(p;\boldsymbol{\theta})$ *is strictly concave on* $(0,\infty)$.

A4. There exists an open interval $I \subset (0,\infty)$ *and an integrable function* $M(\boldsymbol{\theta})$ *such that, for all* $p \in I$,

$$\left|U\Big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\Big)\right| \le M(\boldsymbol{\theta}) \qquad \text{and} \qquad \left|U'\Big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\Big)\,\pi_p(p;\boldsymbol{\theta})\right| \le M(\boldsymbol{\theta}).$$

The first inequality guarantees that the objective $F(p) = \mathbb{E}_{\boldsymbol{\theta}}[U(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta}))]$ *is finite on* I*; the second permits differentiation under the expectation via dominated convergence.*

A5. There exists $0 < \underline{p} < \overline{p} < \infty$ *with* $[\underline{p},\overline{p}] \subset I$ *such that*

$$\mathbb{E}_{\boldsymbol{\theta}}\Big[U'\Big(\omega(\boldsymbol{\theta}) + \pi(\underline{p};\boldsymbol{\theta})\Big)\,\pi_p(\underline{p};\boldsymbol{\theta})\Big] > 0, \qquad \mathbb{E}_{\boldsymbol{\theta}}\Big[U'\Big(\omega(\boldsymbol{\theta}) + \pi(\overline{p};\boldsymbol{\theta})\Big)\,\pi_p(\overline{p};\boldsymbol{\theta})\Big] < 0.$$

[4]Here $P(\boldsymbol{\theta})$ is any scalar summary of the composite-good price faced by the entrepreneur; it need not be an ideal (homothetic, cost-of-living) price index. For the quasilinear specification used later, P can be taken as the unit price of the numeraire and the normalization $P(\boldsymbol{\theta}) \equiv 1$ is without loss.

Then there exists a unique optimal price $p^* \in (\underline{p}, \overline{p})$ *characterized by*

$$\mathbb{E}_{\boldsymbol{\theta}}\Big[U'\Big(\omega(\boldsymbol{\theta}) + \pi(p^*;\boldsymbol{\theta})\Big)\,\pi_p(p^*;\boldsymbol{\theta})\Big] = 0,$$

where

$$\pi_p(p;\boldsymbol{\theta}) = y(p;\boldsymbol{\theta}) + \Big[p - \mathcal{MC}(p;\boldsymbol{\theta})\Big] y_p(p;\boldsymbol{\theta}), \qquad \mathcal{MC}(p;\boldsymbol{\theta}) := \mathcal{C}_y\Big(\mathcal{D}(p;\boldsymbol{\theta});\boldsymbol{\theta}\Big).$$

Proof. See Appendix B.1. □

Remarks. Theorem 1 provides a fully general characterization of optimal pricing under uncertainty with minimal assumptions on preferences, demand, and technology. Assumptions A1–A3 impose standard economic structure: strictly concave preferences and single-peaked profits ensure existence and uniqueness of the optimal price. Assumptions A4–A5 are technical regularity conditions that guarantee differentiability of expected utility and rule out corner solutions. Together, these assumptions encompass standard monopolistic-pricing environments under uncertainty.

A3 does not follow from A2 alone; it is a standing assumption on the profit function. Standard sufficient primitive conditions under which it holds include (a) log-concave inverse demand paired with a weakly convex cost function, (b) linear or exponential demand with a convex cost function, or (c) any demand function for which the inverse elasticity $\mathcal{D}(p;\boldsymbol{\theta})/[p\mathcal{D}_p(p;\boldsymbol{\theta})]$ is weakly decreasing in p, together with weakly convex marginal cost. The CARA-with-linear-demand specialisation used in Sections 3.4 and 4 satisfies A3 directly.

Assumptions A1 and A2 require C^2 regularity because the local pricing expansion in Theorem 2 invokes second derivatives of profit. Theorem 1 itself uses only the first-order conditions and so could be stated under C^1 regularity; we maintain C^2 throughout for uniformity across the main results.

Importantly, no restriction is imposed on the joint distribution of $(\omega(\boldsymbol{\theta}), \mathcal{D}(\cdot;\boldsymbol{\theta}), \mathcal{C}(\cdot;\boldsymbol{\theta}))$. In particular, non-profit income $\omega(\boldsymbol{\theta})$ may be stochastic and arbitrarily correlated with demand and cost shocks. One may therefore interpret $\omega(\boldsymbol{\theta})$ as capturing returns on pre-existing assets or other income sources that provide partial or full insurance against these shocks. Pricing decisions then respond to the residual risk borne by the firm.

The first-order condition admits a natural economic interpretation in terms of risk-adjusted marginal revenue and marginal cost.

Corollary 1 (Risk-adjusted pricing formula)**.** *Define the risk-adjusted expectation at p by*

$$\mathbb{E}_p^{RA}[X] := \frac{\mathbb{E}\Big[U'\Big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\Big)X\Big]}{\mathbb{E}\Big[U'\Big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\Big)\Big]}.$$

Under the assumptions of Theorem 1, the optimal price satisfies

$$\mathcal{MR}_{RA}(p^*) = \mathcal{MC}_{RA}(p^*),$$

where risk-adjusted marginal revenue and marginal cost are given by

$$\mathcal{MR}_{RA}(p) := p + \frac{\mathbb{E}_p^{RA}[y(p;\boldsymbol{\theta})]}{\mathbb{E}_p^{RA}[y_p(p;\boldsymbol{\theta})]}, \qquad \mathcal{MC}_{RA}(p) := \frac{\mathbb{E}_p^{RA}[\mathcal{MC}(p;\boldsymbol{\theta})y_p(p;\boldsymbol{\theta})]}{\mathbb{E}_p^{RA}[y_p(p;\boldsymbol{\theta})]}.$$

Equivalently, these objects admit the decomposition

$$\mathcal{MR}_{RA}(p) = \underbrace{p + \frac{\mathbb{E}[y(p;\boldsymbol{\theta})]}{\mathbb{E}[y_p(p;\boldsymbol{\theta})]}}_{\mathcal{MR}_{mean}(p)} + \underbrace{\left(\frac{\mathbb{E}_p^{RA}[y]}{\mathbb{E}_p^{RA}[y_p]} - \frac{\mathbb{E}[y]}{\mathbb{E}[y_p]}\right)}_{\mathcal{RA}_{MR}(p)}$$

$$\mathcal{MC}_{RA}(p) = \underbrace{\frac{\mathbb{E}\Big[\mathcal{MC}(p;\boldsymbol{\theta})\, y_p(p;\boldsymbol{\theta})\Big]}{\mathbb{E}[y_p(p;\boldsymbol{\theta})]}}_{\mathcal{MC}_{mean}(p)} + \underbrace{\left(\frac{\mathbb{E}_p^{RA}[\mathcal{MC}\, y_p]}{\mathbb{E}_p^{RA}[y_p]} - \frac{\mathbb{E}[\mathcal{MC}\, y_p]}{\mathbb{E}[y_p]}\right)}_{\mathcal{RA}_{MC}(p)}.$$

Proof. See Appendix B.8. □

Interpretation. Optimal pricing equates risk-adjusted marginal revenue and risk-adjusted marginal cost. Relative to the risk-neutral benchmark, risk aversion tilts the effective probability distribution toward states in which marginal utility is high—that is, states in which total resources are low. This reweighting modifies both marginal revenue and marginal cost by assigning greater weight to adverse realizations of demand and cost shocks.

Prices therefore serve as an instrument for managing exposure to risk. By adjusting prices, firms trade off expected profits against lower exposure to states in which marginal utility is high. When utility is linear, the risk-adjusted expectations collapse to physical expectations and the standard risk-neutral pricing rule is recovered.

3.2 Pass-through and local pricing

The exact characterizations of Theorem 1 and Corollary 1 are fully general but do not reveal how specific sources of uncertainty affect prices and markups, nor whether risk aversion leads firms to raise or lower markups. To obtain sharper economic insight, we now study pricing locally around the risk-neutral benchmark. When risk aversion is small, the optimal price can be expressed as a perturbation of the risk-neutral optimum. This local approximation yields a transparent expression for the risk adjustment in prices and allows us to isolate the channels through which uncertainty shapes markup cyclicality. Before providing this approximation, we introduce terminology that connects the curvature of the risk-neutral problem to classic cost/tax pass-through.

Definition 1 (Cost pass-through)**.** *Let $\overline{\kappa}$ be a deterministic cost shifter entering the profit function $\pi(p;\boldsymbol{\theta},\overline{\kappa})$. Let the risk-neutral price $p_{RN}(\overline{\kappa})$ solve*

$$\mathbb{E}_{\boldsymbol{\theta}}\Big[\pi_p(p_{RN}(\overline{\kappa});\boldsymbol{\theta},\overline{\kappa})\Big] = 0.$$

The (risk-neutral) cost pass-through is defined as

$$\varphi := \frac{\partial p_{RN}(\overline{\kappa})}{\partial \overline{\kappa}}.$$

Remark. This definition is fully general and does not impose restrictions on preferences, demand, on technology. Differentiating the risk-neutral first-order condition and applying the implicit function theorem yields

$$\varphi = -\frac{\mathbb{E}_{\boldsymbol{\theta}}[\pi_{p\overline{\kappa},\mathrm{RN}}]}{\mathbb{E}_{\boldsymbol{\theta}}[\pi_{pp,\mathrm{RN}}]}.$$

If the cost shifter enters profits through marginal cost one-for-one (as with a unit tax or an additive input-price shock), then $\pi_{p\overline{\kappa},\mathrm{RN}} = -y_{p,\mathrm{RN}}$, and

$$\varphi = \frac{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,\mathrm{RN}}]}{\mathbb{E}_{\boldsymbol{\theta}}[\pi_{pp,\mathrm{RN}}]}.$$

Since $y_{p,\mathrm{RN}} < 0$ and $\pi_{pp,\mathrm{RN}} < 0$, we have $\varphi > 0$. Thus, pass-through is a sufficient statistic for the local curvature cost of changing prices in the risk-neutral problem: larger φ means the risk-neutral objective is locally flatter, allowing the firm to move prices more (at lower expected-profit cost) to manage profit risk through pricing.

Theorem 2 (Local pricing formula around the risk-neutral price)**.** *Suppose the assumptions of Theorem 1 hold. Let the risk-neutral price p_{RN} be the unique interior maximizer of expected* *profits, characterized by*

$$\mathbb{E}_{\boldsymbol{\theta}}[\pi_p(p_{RN};\boldsymbol{\theta})] = 0, \qquad \mathbb{E}_{\theta}[\pi_{pp}(p_{RN};\boldsymbol{\theta})] < 0.$$

In addition, assume:

A6. Profits have finite fourth moments uniformly in a neighborhood of p_{RN}:

$$\sup_{p\in\mathcal{N}(p_{RN})} \mathbb{E}_{\boldsymbol{\theta}}\Big[|\pi(p;\boldsymbol{\theta})|^4\Big] < \infty.$$

A7. The mapping $p \mapsto \mathbb{E}_{\boldsymbol{\theta}}[\pi(p;\boldsymbol{\theta})]$ is twice continuously differentiable on $\mathcal{N}(p_{RN})$.

A8. Preferences belong to a one-parameter family $\{U_\alpha\}_{\alpha\geq 0}$ with $U_0(\mathcal{W}) = \mathcal{W}$ and absolute risk aversion

$$A_\alpha(\mathcal{W}) := -\frac{U''(\mathcal{W})}{U'(\mathcal{W})} = \alpha A(\mathcal{W}),$$

where $A : \mathbb{R} \to \mathbb{R}_+$ is continuously differentiable.

Define benchmark objects evaluated at p_{RN} as

$$\begin{aligned} \mathcal{W}_{RN}(\boldsymbol{\theta}) &:= \omega(\boldsymbol{\theta}) + \pi(p_{RN};\boldsymbol{\theta}), \qquad & \pi_{RN}(\boldsymbol{\theta}) &:= \pi(p_{RN};\boldsymbol{\theta}), \\ \pi_{p,RN}(\boldsymbol{\theta}) &:= \pi_p(p_{RN};\boldsymbol{\theta}), & y_{p,RN}(\boldsymbol{\theta}) &:= y_p(p_{RN};\boldsymbol{\theta}). \end{aligned}$$

Then, for sufficiently small α, the optimal price $p(\alpha)$ under preferences U_α satisfies

$$p(\alpha) = p_{RN} + \alpha\mathcal{R} + \mathcal{O}(\alpha^2) + \alpha \cdot o_{\|\Sigma\|\to 0}(1),$$

where the last term is absent under CARA preferences (in which case A is constant and the mean-value step in the proof is exact); under variable-A preferences (e.g., CRRA) it represents the higher-order residual of the integral-form expansion of U'_α and vanishes under joint small-(α, Σ) asymptotics. where the first-order risk adjustment admits the covariance representation

$$\mathcal{R} = \varphi\Bigg[\underbrace{\frac{\overline{A}\cdot \mathrm{Cov}_{\boldsymbol{\theta}}\big(\pi_{RN}, \pi_{p,RN}\big)}{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,RN}]}}_{\text{pure risk term}} + \underbrace{\frac{\mathrm{Cov}_{\boldsymbol{\theta}}\big(A(\mathcal{W}_{RN}), \pi_{RN}\pi_{p,RN}\big)}{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,RN}]}}_{\text{wealth-effects term}}\Bigg],$$

with $\overline{A} := \mathbb{E}_{\theta}[A(\mathcal{W}_{RN})]$.

Proof. See Appendix B.2. □

Remarks. Although the entrepreneur's objective is defined over total resources, price affects utility only through profits. Accordingly, the local expansion and curvature conditions are naturally stated in terms of expected profits, and at $\alpha = 0$ the expected-utility problem collapses to profit maximization. Also, while the result is stated as a local approximation, it is exact in leading cases of interest—such as CARA preferences with normally distributed shocks, or more generally whenever profits are affine in the underlying shocks—so that higher-order risk corrections vanish.

The risk adjustment $\mathcal{R}$ admits a transparent decomposition. First, it makes explicit that the magnitude of precautionary pricing is disciplined by (i) cost-pass through φ and (ii) local demand slope $\mathbb{E}[y_{p,\mathrm{RN}}]$. All else equal, stronger pass-through amplifies the effect of risk on prices, while steeper demand (more price-sensitive consumers) dampens it. The first term within brackets captures pure risk exposure: pricing responds to the covariance between profits and marginal profits, scaled by average absolute risk aversion. The second term captures wealth effects: when absolute risk aversion varies with total resources, optimal pricing additionally reflects how states with high marginal profits co-vary with wealth. Because $\mathcal{W}_{\mathrm{RN}} = \omega + \pi_{\mathrm{RN}}$, non-profit income $\omega(\boldsymbol{\theta})$ acts as an insurance channel shaping residual risk, and no restriction is imposed on its joint distribution with demand and cost shocks.

Intuition. At the risk-neutral price p_{RN}, expected marginal profits are zero, so a marginal change in price has no first-order effect on mean profits. A risk-averse firm, however, also cares about how a price change redistributes profits across states. Locally around p_{RN}, the direction of the optimal price adjustment is therefore governed entirely by how marginal price changes affect the riskiness of profits.

The covariance term $\mathrm{Cov}_{\boldsymbol{\theta}}(\pi_{\mathrm{RN}}, \pi_{p,\mathrm{RN}})$ captures this effect. It measures whether marginal price increases change profits disproportionately in already good states or already bad states. If this covariance is negative, raising price increases profits primarily in low-profit states and reduces volatility; price increases therefore shift profits toward states in which profits would otherwise be low, and the firm optimally raises price relative to the risk-neutral benchmark. If the covariance is positive, raising price amplifies risk by shifting income toward already favorable states, and the firm optimally lowers price to smooth profits across states.

When absolute risk aversion varies with total resources, an additional wealth-effects term appears. This term captures how the risk-management value of pricing depends on the correlation between non-profit income and the shocks affecting profits. Under CARA preferences, absolute risk aversion is constant and this channel is absent; under CRRA preferences, it generally survives and can either amplify or attenuate the pure risk effect, depending on how marginal profits co-move with wealth. The next corollary formalizes these observations.

Corollary 2 (CARA and CRRA risk corrections)**.** *Under the assumptions of Theorem 2:*

(i) ***CARA****. If $U_\alpha(\mathcal{W}) = -\exp(-\alpha\mathcal{W})$, then absolute risk aversion is constant, so $A(\mathcal{W}) \equiv 1$ and the wealth-effects term vanishes, and the risk adjustment simplifies to*

$$\mathcal{R} = \varphi \cdot \frac{Cov_{\boldsymbol{\theta}}(\pi_{RN}, \pi_{p,RN})}{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,RN}]}.$$

(ii) ***CRRA****. If $U(\mathcal{W}) = \mathcal{W}^{1-\sigma}/(1-\sigma)$, with $U(\mathcal{W}) = \log \mathcal{W}$ for $\sigma = 1$, then $A_\sigma(\mathcal{W}) = \sigma/\mathcal{W} = \sigma \cdot A(\mathcal{W})$ with $A(\mathcal{W}) = 1/\mathcal{W}$, and the general decomposition applies with a non-vanishing wealth-effects term.*

Remark. Although CRRA preferences are not globally nested in the family $\{U_\alpha\}_{\alpha \geq 0}$, they satisfy the local scaling property of absolute risk aversion required for Theorem 2 and are therefore covered by the local approximation. Under CARA preferences, pricing responds solely to pure risk exposure. Under CRRA preferences, wealth effects interact with the precautionary pricing motive through non-profit income.

To isolate the core precautionary pricing mechanism and its implications for markup cyclicality, we therefore focus first on the pure-risk component of the risk adjustment. The local pricing formula shows that risk affects prices through the covariance between profits and marginal profits. Decomposing this covariance allows us to identify how distinct sources of uncertainty—cost risk, demand-level risk, and demand-elasticity risk—shape pricing incentives, both individually and through their interactions. Focusing on the pure-risk term characterizes the precautionary role of pricing independently of wealth effects, which depend on the joint distribution of profits and non-profit income and are discussed separately.

Proposition 1 (Decomposition of the pure risk adjustment term by source)**.** *Let $\boldsymbol{\theta} = (\eta, \xi, \psi)$ collect a marginal-cost shifter η, a demand-level shifter ξ, and a demand-elasticity shifter ψ. Suppose the assumptions of Theorem 2 hold. In addition, assume:*

A9. Shocks have mean zero, are mutually independent, and have finite fourth moments.

A10. For each p in a neighborhood of p_{RN}, the mapping $\boldsymbol{\theta} \mapsto \pi(p; \boldsymbol{\theta})$ is three times continuously differentiable in a neighborhood of $\boldsymbol{\theta} = \mathbf{0}$.

Let $\sigma_k^2 = Var(\theta_k)$ and $\Sigma = Var(\boldsymbol{\theta})$. Let $\mathcal{R}$ denote the risk adjustment from Theorem 2, and write

$$\mathcal{R} = \mathcal{R}_{risk} + \mathcal{R}_{wealth},$$

where $\mathcal{R}_{risk}$ is the pure-risk component.

For $k, l \in \{\eta, \xi, \psi\}$, define the following exposures, evaluated at $(p, \boldsymbol{\theta}) = (p_{RN}, \mathbf{0})$:

$$\pi_k := \partial_k \pi, \qquad \pi_{pk} := \partial_k \pi_p, \qquad \pi_{kl} := \partial_{kl} \pi, \qquad \pi_{p,kl} := \partial_{kl} \pi_p.$$

Then the pure-risk component admits the first-order expansion

$$\mathcal{R}_{risk} = \underbrace{\mathcal{R}^{(\eta)}(\sigma_\eta^2)}_{\text{cost risk}} + \underbrace{\mathcal{R}^{(\xi)}(\sigma_\xi^2)}_{\text{demand-level risk}} + \underbrace{\mathcal{R}^{(\psi)}(\sigma_\psi^2)}_{\text{demand-elasticity risk}} + o\left(\|\Sigma\|\right),$$

where

$$\mathcal{R}^{(k)}(\sigma_k^2) := \varphi \overline{A} \cdot \frac{\pi_k \pi_{pk}}{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,RN}]} \sigma_k^2,$$

$\overline{A} := \mathbb{E}_{\boldsymbol{\theta}}[A(\mathcal{W}_{RN})]$, and $\|\cdot\|$ denotes any matrix norm. At the next order in Σ — assuming additionally that π is four-times continuously differentiable at $\boldsymbol{\theta} = \mathbf{0}$ — the covariance picks up cross-risk interaction contributions

$$\mathcal{R}^{(kl)}(\sigma_k^2 \sigma_l^2) := \varphi \, \frac{\overline{A}}{2} \, \frac{\pi_{kl} \, \pi_{p,kl}}{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,RN}]} \, \sigma_k^2 \sigma_l^2, \qquad k \neq l,$$

which are $O(\|\Sigma\|^2)$ and therefore suppressed under the leading-order small-noise scaling.

Under CARA preferences, $A(\cdot)$ is constant and $\mathcal{R}_{wealth} = 0$, so the above decomposition characterizes the risk adjustment to leading order in Σ.

Proof. See Appendix B.3. □

Scope of the second-order correction. A fully rigorous $O(\|\Sigma\|^2)$ decomposition would augment the cross-risk interactions $\mathcal{R}^{(kl)}$ with (a) diagonal centered-quartic contributions of the form $\frac{1}{4} \pi_{kk} \, \pi_{p,kk} \, \mathrm{Var}_{\boldsymbol{\theta}}(\theta_k^2)$ and (b) third-derivative contributions involving $\pi_{kkl} \, \pi_{pl} \, \mathbb{E}_{\boldsymbol{\theta}}[\theta_k^2 \theta_l]$, both of which are of the same order as $\mathcal{R}^{(kl)}$. We do not pursue this expansion because the empirical and calibration analysis in Sections 5–6 relies only on the first-order sourcewise decomposition; under any small-noise scaling $\Sigma = \varepsilon \, \Sigma_0$ with $\varepsilon \to 0$, all second-order terms (interactions, centered-quartic, and third-derivative) are $O(\varepsilon^4)$ and asymptotically negligible relative to the $O(\varepsilon^2)$ sourcewise channels.

Interpretation. Proposition 1 decomposes the pure-risk component into contributions associated with individual sources of uncertainty and their interactions. The terms $\mathcal{R}^{(k)}$ capture the effect of introducing uncertainty in shock k in isolation, holding all other sources of risk fixed. The interaction terms $\mathcal{R}^{(kl)}$ capture how the presence of one source of risk modifies the pricing response to another. The pure-risk component depends only on the joint behavior of profits and marginal profits; wealth effects depend instead on the joint distribution of profits and non-profit income and are analyzed separately.

All pure-risk effects operate through the covariance between profits and marginal profits. Because the denominator of the risk adjustment reflects only the curvature of the risk-neutral problem, the sign and magnitude of each component are governed by how a marginal price change redistributes profits across states of the world. In this sense, pricing under risk is analogous to a portfolio-choice problem: prices are adjusted to shift profits toward states in which they are relatively low.

When multiple sources of uncertainty are present, the total effect of any one risk combines its pure effect with interaction terms. Interactions arise because a price change may stabilize profits only for particular combinations of shocks, so the risk-management properties of pricing depend on the joint distribution of uncertainties rather than on any single variance in isolation. As a result, interaction terms can attenuate or amplify pure risk effects, and in some cases overturn their sign. Overall, risk-averse pricing reflects a *composition of precautionary motives* across cost risk, demand-level risk, and demand-elasticity risk.

The next proposition provides sufficient—and economically interpretable—conditions under which each source of uncertainty raises or lowers prices via the pure-risk term.

Proposition 2 (Sufficient conditions to sign each risk source in pure-risk component)**.** *Let $\boldsymbol{\theta} = (\eta, \xi, \psi)$, where η is a marginal cost shifter, ξ is a demand-level shifter, and ψ is a demand-elasticity shifter. Suppose the assumptions of Proposition 1 hold, and that*

$$\mathbb{E}_{\boldsymbol{\theta}}\Big[\pi_{pp}(p_{RN}; \boldsymbol{\theta})\Big] < 0.$$

Define the second-order pure-risk contribution of source $k \in \{\eta, \xi, \psi\}$ to the price adjustment by

$$\mathcal{R}^{(k)}_{risk} := \mathcal{R}^{(k)}(\sigma_k^2) + \frac{1}{2}\sum_{l \neq k} \mathcal{R}^{(kl)}(\sigma_k^2 \sigma_l^2),$$

where $\mathcal{R}^{(k)}$ and $\mathcal{R}^{(kl)}$ are defined in Proposition 1.

Assume that, locally at $(p, \boldsymbol{\theta}) = (p_{RN}, \mathbf{0})$, the following exposure conditions hold:

$$\pi_\eta < 0, \quad \pi_{p\eta} > 0; \qquad \pi_\xi > 0, \quad \pi_{p\xi} \geq 0; \qquad \pi_\psi < 0, \quad \pi_{p\psi} < 0.$$

Assume further that interactions involving demand shifters do not overturn these signs, in the sense that

$$\pi_{\xi l}\pi_{p,\xi l} \geq 0 \;\; \textit{for all} \;\; l \neq \xi, \qquad \pi_{\psi l}\pi_{p,\psi l} \geq 0 \;\; \textit{for all} \;\; l \neq \psi.$$

Then, to second order in uncertainty:

1. *Cost-uncertainty raises prices provided the interaction contributions from demand shocks are not too large: there exist thresholds* $\overline{\sigma}^2_\xi(\sigma^2_\psi), \overline{\sigma}^2_\psi(\sigma^2_\xi) > 0$ *such that* $\mathcal{R}^{(\eta)}_{risk} > 0$ *whenever* $\sigma^2_\xi \leq \overline{\sigma}^2_\xi(\sigma^2_\psi)$ *and* $\sigma^2_\psi \leq \overline{\sigma}^2_\psi(\sigma^2_\xi)$*; equivalently, in the small-noise scaling* $\Sigma = \varepsilon\Sigma_0$ *the conclusion holds for all sufficiently small* ε*.*
2. *Demand-level uncertainty lowers prices:* $\mathcal{R}^{(\xi)}_{risk} \leq 0$*, with strict inequality if* $\pi_{p\xi} > 0$*.*
3. *Demand-elasticity uncertainty lowers prices:* $\mathcal{R}^{(\psi)}_{risk} < 0$*.*

Proof. See Appendix B.4 □

Remarks. Proposition 2 relies on four assumptions. First, the existence of an interior risk-neutral optimum ensures that pricing distortions are locally governed by deviations from a well-defined benchmark. Second, the sign restrictions on profit and marginal-profit exposures formalize economically transparent mechanisms: higher marginal costs reduce profits but raise the marginal value of increasing prices; stronger demand raises both profits and marginal profits; and higher demand elasticity depresses both. These conditions are satisfied in standard monopolistic-pricing environments and simply encode how prices redistribute profits across states of the world.

Third, the restriction on the interaction terms rules out cases in which joint demand shocks overturn these basic risk-management properties of pricing. This condition holds in widely used demand systems—including linear, isoelastic, and separable specifications—where interactions preserve the sign of marginal risk-adjustment incentives. Fourth, the elasticity-risk threshold reflects a local dominance condition: cost risk leads to raising prices unless fluctuations in demand elasticity are sufficiently strong to offset the risk-reducing effect of price increases.

Importantly, the proposition only characterizes the pure-risk component of price adjustment. Wealth effects, which arise when absolute risk aversion varies with total resources, are analyzed separately. Focusing on the pure-risk channel isolates the core precautionary pricing mechanism embedded in price setting and clarifies how distinct sources of uncertainty shape markup cyclicality.

3.3 Markup cyclicality

The local pricing results have direct implications for markup cyclicality. Because firms set prices *ex ante*, before the realization of shocks, the relevant object for understanding markup dynamics is the *ex-ante* markup implied by the firm's pricing decision, rather than the realized markup that fluctuates mechanically with ex-post cost realizations.

Formally, let $\overline{\kappa}_i > 0$ denote a firm-specific, time-invariant deterministic scale for marginal cost (a baseline technology parameter), and define the ex-ante markup as

$$\mu_{it}(\alpha) := \frac{p_{it}(\alpha)}{\overline{\kappa}_i},$$

where $p_{it}(\alpha)$ is the price chosen by firm i in period t under risk aversion α. Time-invariance of $\overline{\kappa}_i$ is what all specialisations in the canonical cases below require: Corollaries 3–6 all use the constant firm-specific scale $\overline{\kappa}_i$ around which realized marginal cost fluctuates, $\kappa_{it} = \overline{\kappa}_i + \eta_{it}$. Time-variation in the realized marginal cost is absorbed into the additive shock η_{it}; all time-series variation in μ_{it} under the theorem works through the ex-ante optimal price $p_{it}(\alpha)$, not through a time-varying denominator.

By Theorem 2, the optimal ex-ante markup admits the first-order representation

$$\mu_{it}(\alpha) = \mu_{it,\text{RN}} + \frac{\alpha}{\overline{\kappa}_i}\mathcal{R}_{it}, \qquad \mu_{it,\text{RN}} := \frac{p_{it,\text{RN}}}{\overline{\kappa}_i},$$

where $\mathcal{R}_{it}$ denotes the risk-adjustment implied by the distribution of shocks faced by firm i in period t.

Markup cyclicality may therefore arise from time variation in uncertainty. As the distribution of shocks evolves over the business cycle, the risk adjustment $\mathcal{R}_{it}$ varies over time, inducing endogenous movements in ex-ante markups—even holding expected costs and demand fixed.

Markup cyclicality can be defined at different levels of aggregation, depending on the economic question of interest. In what follows, we consider three measures of markup cyclicality: the covariance between firm-level ex-ante markups and firm-level output, the covariance between the aggregate sectoral markup and sectoral output, and the covariance between the aggregate (economy-wide) markup and aggregate output. Firm-level cyclicality isolates pricing under uncertainty, sectoral cyclicality combines pricing behavior with within-sector reallocation, and aggregate cyclicality further reflects between-sector shifts in economic activity.

Firm-level markup cyclicality. At the firm-level, markup cyclicality is captured by $\text{Cov}(\mu_{it}, y_{it})$, which measures how a firm's ex-ante markup co-moves with its output over the business cycle. This object isolates the within-firm component of markup dynamics and provides the most direct link between the pricing mechanism and firm-level activity.

Proposition 3 (Firm-level markup cyclicality)**.** *Suppose the assumptions of Theorem 2 hold period by period for each firm i, and let the distribution of shocks vary over time through uncertainty parameters $\boldsymbol{\sigma}_{it}^2 = (\sigma_{\eta,it}^2, \sigma_{\xi,it}^2, \sigma_{\psi,it}^2)$. Then, for sufficiently small α,*

$$Cov(\mu_{it}, y_{it}) = Cov(\mu_{it,RN}, y_{it}) + \frac{\alpha}{\overline{\kappa}_i} \sum_{k \in \{\eta,\xi,\psi\}} \underbrace{\frac{\partial \mathcal{R}_{it}}{\partial \sigma_{k,it}^2}}_{\substack{\text{markup sensitivity} \\ \text{to risk } k}} \times \underbrace{Cov(\sigma_{k,it}^2, y_{it})}_{\text{cyclicality of risk } k} + \mathcal{O}(\alpha^2) + \mathcal{O}(\alpha \, \|\Delta \boldsymbol{\sigma}_{it}^2\|^2).$$

Proof. See Appendix B.5. □

Interpretation. Markup cyclicality reflects the interaction between pricing under risk and the cyclical behavior of uncertainty. Each source of risk affects markup dynamics through two distinct components. The *exposure channel*, $\partial \mathcal{R}_{it}/\partial \sigma_{k,it}^2$, measures how sensitive optimal pricing is to uncertainty of type k, and the *cyclical channel*, $\mathrm{Cov}(\sigma_{k,it}^2, y_{it})$, captures how that uncertainty co-moves with firm-level output over the business cycle.

As a result, different risks need not have symmetric effects on markup dynamics. Even if multiple risks are equally volatile on average, their impact on markups depends on exposure and cyclical comovement. Markups may therefore be procyclical, countercyclical, or acyclical depending on the composition and cyclicality of risk.

Sectoral markup cyclicality. We construct aggregate markups using the harmonic sales-weighted markup. For any sector s and period t, let $\mathcal{I}_{st}$ denote the set of active firms. The sectoral markup is

$$\mu_{st}^{\text{hsw}} = \left(\sum_{i \in \mathcal{I}_{st}} \omega_{it} \mu_{it}^{-1} \right)^{-1}, \qquad \omega_{it} = \frac{p_{it} y_{it}}{\sum_{j \in \mathcal{I}_{st}} p_{jt} y_{jt}}, \quad \sum_{i \in \mathcal{I}_{st}} \omega_{it} = 1.$$

Proposition 4 (Sectoral markups and sectoral markup cyclicality)**.** *Suppose the assumptions of Theorem 2 hold period by period for each firm i. For sufficiently small α,*

$$\mu_{st}^{hsw} = \mu_{st,RN}^{hsw} + \alpha \left(\mu_{st,RN}^{hsw} \right)^2 \sum_{i \in \mathcal{I}_{st}} \omega_{it} \frac{1}{\overline{\kappa}_{it}} \frac{\mathcal{R}_{it}(\boldsymbol{\sigma}_{it}^2)}{\mu_{it,RN}^2} + \mathcal{O}(\alpha^2).$$

Moreover, the cyclicality of the sectoral markup with sectoral output y_{st} satisfies

$$Cov(\mu_{st}^{hsw}, y_{st}) = Cov(\mu_{st,RN}^{hsw}, y_{st}) + \alpha\left(\mu_{st,RN}^{hsw}\right)^2 \sum_{i \in \mathcal{I}_{st}} \frac{\omega_{it}}{\overline{\kappa}_{it}\mu_{it,RN}^2}\left[\sum_{k \in \{\eta,\xi,\psi\}} \frac{\partial \mathcal{R}_{it}}{\partial \sigma_{k,it}^2} \times Cov(\sigma_{k,it}^2, y_{st})\right].$$

Proof. See Appendix B.6. □

Interpretation. To first order, sectoral markups equal the risk-neutral aggregate markup plus a weighted average of firm-level risk adjustments. Consequently, time variation in uncertainty generates aggregate markup fluctuations even if the cross-sectional distribution of revenue shares remains unchanged.

Our analysis adopts the same harmonic sales-weighted markup used by Burstein, Carvalho and Grassi (2025). The economic mechanisms, however, differ. Burstein, Carvalho and Grassi emphasize reallocation: aggregate shocks shift market shares across firms with heterogeneous markups, and the aggregate markup changes mechanically as a result. In contrast, our framework explains why firm-level markups move in the first place. Because firms set prices ex ante under uncertainty, time-varying risk induces systematic movements in firm-level markups even holding technology, preferences, and the distribution of market power fixed. In this sense, reallocation governs how markups aggregate, while uncertainty how markups adjust over the cycle.

The complementarity between the two mechanisms is most transparent in statistical decompositions of the (inverse) aggregate markup that separate within- from between-firm components. A standard decomposition used in empirical work is

$$\Delta\left(\mu_{st}^{\text{hsw}}\right)^{-1} = \underbrace{\sum_{i \in \mathcal{I}_s} \overline{\omega}_i \Delta\mu_{it}^{-1}}_{\substack{\text{within} \\ \text{(pricing)}}} + \underbrace{\sum_{i \in \mathcal{I}_s} \Delta\omega_{it}\overline{\mu}_i^{-1}}_{\substack{\text{between} \\ \text{(reallocation)}}},$$

where bars on top of variables denote midpoint (Tornqvist) averages across periods. This identity follows by applying the elementary Tornqvist formula $\Delta(ab) = \overline{a}\,\Delta b + \overline{b}\,\Delta a$ to each summand $\omega_{it}\mu_{it}^{-1}$ and collecting terms (see e.g. Melitz and Polanec, 2015 for a textbook treatment in the reallocation-decomposition literature).

The first term captures movements in firm-level markups holding revenue shares fixed. The second term captures shifts in revenue shares across firms with heterogeneous markups. The theory of Burstein, Carvalho and Grassi highlights the latter channel. Our theory provides a structural foundation for the former by linking changes in firm-level markups to time-varying uncertainty.

Ex-ante versus ex-post aggregation. Proposition 4 aggregates the *ex-ante* firm-level markups $\mu_{it} = p_{it}/\overline{\kappa}_{it}$. Standard empirical markup measurement (e.g., De Loecker, Eeckhout and Unger, 2020), however, recovers the *ex-post* realized markup $\mu_{it}^{\text{ep}} = p_{it}/\kappa_{it}$. A natural concern is whether aggregating ex-post markups delivers the same aggregate cyclicality predictions as aggregating the theoretical ex-ante objects. The following proposition shows that, under the price-commitment timing of Theorem 1, the two harmonic sales-weighted aggregates coincide exactly. The argument relies on a law-of-large-numbers averaging of mean-zero cost shocks, so we pass to the continuum-of-firms limit $|\mathcal{I}_{st}| \to \infty$ and replace the sums of Proposition 4 with integrals over the unit-mass of active firms; the finite-N counterpart holds approximately, with a residual of order $1/\sqrt{N_{st}}$.

Proposition 5 (Ex-ante / ex-post aggregation equivalence)**.** *Maintain the assumptions of Theorem 1 (including the price-commitment setup, Assumption A0), and let $\overline{\kappa}_{it} = \mathbb{E}[\kappa_{it} \mid \mathcal{F}_{it}^{pre}]$ denote the conditional expectation of marginal cost at the time of pricing. Then, in the continuum-of-firms limit, the output-weighted harmonic sales-weighted markup computed with ex-post realized costs equals the corresponding aggregate computed with ex-ante expected costs:*

$$\mu_{st}^{hsw,\,ep} := \frac{\int y_{it}\,di}{\int y_{it}\,\kappa_{it}/p_{it}\,di} = \frac{\int y_{it}\,di}{\int y_{it}\,\overline{\kappa}_{it}/p_{it}\,di} =: \mu_{st}^{hsw,\,ea}.$$

In particular, the sectoral cyclicality decomposition in Proposition 4 applies verbatim to the ex-post aggregate. In finite samples of size N_{st}, the two coincide up to a residual of order $1/\sqrt{N_{st}}$ by the law of large numbers.

Proof. See Appendix B.7. □

Implication. The aggregate markup and its cyclicality are invariant to whether firm-level markups are measured ex ante or ex post. This reconciles the theoretical ex-ante focus of our pricing model with the ex-post aggregation conventions of the empirical markup literature: the DLW/DLE aggregate markup is already the correct measurement target for our theoretical predictions. The ex-ante/ex-post distinction matters at the *firm level* — where realized and expected markups differ by the cost-shock realization η_{it} — but disappears in the aggregate under harmonic sales-weighting. This is a consequence of the cross-sectional averaging of mean-zero cost shocks, combined with the price-commitment timing that makes y_{it} and p_{it} orthogonal to η_{it} in the cross-section.

Aggregate markup cyclicality. The same logic applies at the aggregate level. The economy-wide harmonic sales-weighted markup aggregates firm-level markups using firm sales shares in total economy-wide revenue. Equivalently, it can be written as the harmonic aggregation of sectoral markups using sectoral revenue shares: partitioning firms into sectors $s = 1, \ldots, S$ and denoting sectoral sales by $Y_{st} = \sum_{i \in s} y_{it} p_{it}$,

$$\mu_t^{\text{hsw}} = \frac{\sum_i y_{it} p_{it}}{\sum_i y_{it} \kappa_{it}} = \frac{\sum_s Y_{st}}{\sum_s Y_{st}/\mu_{st}^{\text{hsw}}},$$

where $\mu_{st}^{\text{hsw}} = Y_{st}/\sum_{i \in s} y_{it}\kappa_{it}$ is the sectoral harmonic sales-weighted markup. The identity follows directly from the definition of μ_{st}^{hsw} via $Y_{st}/\mu_{st}^{\text{hsw}} = \sum_{i \in s} y_{it}\kappa_{it}$ and summing over sectors. It makes clear that aggregate markup fluctuations can be understood as combining movements within sectors and reallocation across sectors.

Aggregate markup dynamics therefore reflect three components: (i) firm-level pricing, including responses to uncertainty, (ii) reallocation across firms within sectors, and (iii) reallocation across sectors.

3.4 Canonical cases

We now illustrate the general markup-cyclicality mechanism in canonical monopolistic-pricing environments. These cases deliver closed-form expressions for prices, markups, and their cyclicality, and make transparent how different sources of uncertainty map into procyclical or countercyclical markups. Corollary 3 nests the linear-demand example of Section 2 within the general theory and provides an explicit mapping from each component of uncertainty to the risk adjustment $\mathcal{R}_{it}$. We then extend the analysis to isoelastic demand and to CRRA preferences, where wealth effects interact with the precautionary pricing motive.

Corollary 3 (Markup cyclicality with CARA utility, linear demand, and constant marginal cost)**.** *Let preferences be* $U_\alpha(\mathcal{W}_{it}) = -\exp(-\alpha \mathcal{W}_{it})$. *Marginal cost is* $\kappa_{it} = \overline{\kappa}_i + \eta_{it}$ *with* $\overline{\kappa}_i > 0$, *and demand is linear*

$$\mathcal{D}_{it}(p_{it}; \boldsymbol{\theta}_{it}) = a + \xi_{it} - (b + \psi_{it}) p_{it},$$

where $a, b > 0$ *and* $a > b\overline{\kappa}_i$. *Assume the shocks* $\eta_{it}, \xi_{it}, \psi_{it}$ *are mutually independent, mean-zero, with time-varying variances* $\boldsymbol{\sigma}_{it}^2 = (\sigma_{\eta,it}^2, \sigma_{\xi,it}^2, \sigma_{\psi,it}^2)$ *and finite fourth moments.*

The risk-neutral price and markup are

$$p_{i,RN} = \frac{a + b\overline{\kappa}_i}{2b}, \qquad \mu_{i,RN} = \frac{p_{i,RN}}{\overline{\kappa}_i}.$$

Moreover, for sufficiently small α, *the optimal ex-ante markup satisfies*

$$\mu_{it}(\alpha) = \mu_{i,RN} + \frac{\alpha}{\overline{\kappa}_i} \mathcal{R}_{it}(\boldsymbol{\sigma}_{it}^2) + \mathcal{O}(\alpha^2),$$

where the (first-order) risk adjustment is

$$\mathcal{R}_{it}(\boldsymbol{\sigma}_{it}^2) = \underbrace{\frac{a - b\overline{\kappa}_i}{4}\sigma_{\eta,it}^2}_{\text{pure cost risk}} - \underbrace{\frac{a - b\overline{\kappa}_i}{4b^2}\sigma_{\xi,it}^2}_{\text{pure demand-level risk}} - \underbrace{\frac{a(a + b\overline{\kappa}_i)(a - b\overline{\kappa}_i)}{8b^4}\sigma_{\psi,it}^2}_{\text{pure elasticity risk}} - \underbrace{\frac{a + b\overline{\kappa}_i}{4b^2}\sigma_{\eta,it}^2\sigma_{\psi,it}^2}_{\text{cost} \times \text{elasticity risk}}.$$

Consequently, firm-level markup cyclicality is

$$\begin{aligned} Cov(\mu_{it}, y_{it}) = \frac{\alpha}{\overline{\kappa}_i}\Bigg[&\frac{a - b\overline{\kappa}_i}{4}Cov(\sigma_{\eta,it}^2, y_{it}) - \frac{a - b\overline{\kappa}_i}{4b^2}Cov(\sigma_{\xi,it}^2, y_{it}) \\ &- \frac{a(a + b\overline{\kappa}_i)(a - b\overline{\kappa}_i)}{8b^4}Cov(\sigma_{\psi,it}^2, y_{it}) - \frac{a + b\overline{\kappa}_i}{4b^2}Cov(\sigma_{\eta,it}^2\sigma_{\psi,it}^2, y_{it})\Bigg] + \mathcal{O}(\alpha^2). \end{aligned}$$

Proof. See Appendix B.9. □

CARA utility and linear demand. Corollary 3 provides a transparent benchmark. Cost risk raises markups because higher prices reduce exposure to low-profit, high-cost states. By contrast, both demand-level and demand-elasticity risk lower markups because price increases shift profits toward already favorable demand realizations and thereby amplify risk. The pass-through in this case is $\varphi = 1/2$. Hence the coefficients above can be read as "exposure terms" scaled by φ. Since the risk-neutral markup is constant over time, markup cyclicality is entirely driven by the interaction between risk exposure and the cyclical behavior of uncertainty. Countercyclical cost uncertainty makes markups countercyclical, whereas countercyclical demand uncertainty makes markups procyclical. When multiple risks co-move with the cycle, observed markup dynamics reflect their relative strength and interaction.

Corollary 4 (Markup cyclicality with CARA utility, isoelastic demand, and constant marginal cost)**.** *Let preferences be $U_\alpha(\mathcal{W}_{it}) = -\exp(-\alpha\mathcal{W}_{it})$. Marginal cost is $\kappa_{it} = \overline{\kappa}_i + \eta_{it}$ with $\overline{\kappa}_i > 0$, and demand is isoelastic*

$$\mathcal{D}_{it}(p_{it}; \boldsymbol{\theta}_{it}) = (1 + \xi_{it})p^{-(\varepsilon + \psi_{it})},$$

where $\varepsilon > 1$ and $\boldsymbol{\theta}_{it} = (\eta_{it}, \xi_{it}, \psi_{it})$. Assume the shocks $\eta_{it}, \xi_{it}, \psi_{it}$ are mutually independent, mean-zero, with time-varying variances $\boldsymbol{\sigma}_{it}^2 = (\sigma_{\eta,it}^2, \sigma_{\xi,it}^2, \sigma_{\psi,it}^2)$ and finite fourth moments.

The risk-neutral price (for small elasticity uncertainty) is

$$p_{it,RN} = \overline{p}_i + \frac{\overline{\kappa}_i L_i}{(\varepsilon - 1)^2}\sigma_{\psi,it}^2 + o(\sigma_{\psi,it}^2), \qquad \overline{p}_i := \frac{\varepsilon}{\varepsilon - 1}\overline{\kappa}_i, \qquad L_i := \ln \overline{p}_i,$$

where $\overline{p}_i$ is the deterministic monopoly price, and the corresponding markup is $\mu_{it,RN} := p_{it,RN}/\overline{\kappa}_i$.

Moreover, for sufficiently small α, the optimal ex-ante markup satisfies

$$\mu_{it}(\alpha) = \mu_{it,RN} + \frac{\alpha}{\overline{\kappa}_i}\mathcal{R}_{it}(\boldsymbol{\sigma}^2_{it}) + \mathcal{O}(\alpha^2),$$

where the first-order risk adjustment decomposes as

$$\mathcal{R}_{it}(\boldsymbol{\sigma}^2_{it}) = \underbrace{C_{i\eta}\,\sigma^2_{\eta,it}}_{\text{pure cost risk}} - \underbrace{\Theta_{i\psi}\,\sigma^2_{\psi,it}}_{\text{pure elasticity risk}} + \underbrace{C_{i\eta}\sigma^2_{\eta,it}\sigma^2_{\xi,it}}_{\text{cost}\times\text{demand-level risk}} \underbrace{-\Theta_{i\eta\psi}\,\sigma^2_{\eta,it}\sigma^2_{\psi,it}}_{\text{cost}\times\text{elasticity risk}}.$$

The coefficients (evaluated at $\overline{p}_i$) are

$$C_{i\eta} := \frac{\left(\frac{\varepsilon-1}{\varepsilon}\right)^{\varepsilon-1}}{\overline{\kappa}_i^{\varepsilon}}, \qquad \Theta_{i\psi} := \frac{(\overline{p}_i-\overline{\kappa}_i)^2}{\varepsilon-1}\,\overline{p}_i^{-\varepsilon}\,L_i, \qquad \Theta_{i\eta\psi} := \frac{\overline{p}_i^{-\varepsilon}}{\varepsilon-1}\left(2L_i - 2\varepsilon L_i^2\right).$$

Consequently, firm-level markup cyclicality is

$$\begin{aligned}Cov(\mu_{it}, y_{it}) = \frac{L_i}{(\varepsilon-1)^2}Cov(\sigma^2_{\psi,it}, y_{it}) + \frac{\alpha}{\overline{\kappa}_i}\Big[&C_{i\eta}Cov(\sigma^2_{\eta,it}, y_{it})\\ &- \Theta_{i\psi}Cov(\sigma^2_{\psi,it}, y_{it}) + C_{i\eta}Cov(\sigma^2_{\eta,it}\sigma^2_{\xi,it}, y_{it}) - \Theta_{i\eta\psi}Cov(\sigma^2_{\eta,it}\sigma^2_{\psi,it}, y_{it})\Big] + \mathcal{O}(\alpha^2).\end{aligned}$$

Proof. See Appendix B.10. □

CARA utility and isoelastic demand. Corollary 4 shows that the risk-management role of pricing depends on demand curvature. As in the linear-demand case, cost risk raises markups. Elasticity risk lowers markups: when demand becomes more price-sensitive, high prices increase profit volatility, strengthening the incentive to compress markups. A key difference is that pure demand-level risk does not affect prices at first order under isoelastic demand, since level shocks scale revenues proportionally across prices; demand-level uncertainty matters only through interaction with cost risk. The pass-through in this case is time-varying due to the presence of elasticity risk. Holding the uncertainty environment fixed,

$$\varphi_{it} = \frac{\varepsilon}{\varepsilon-1} + \frac{1}{(\varepsilon-1)^2}(L_i+1)\sigma^2_{\psi,it},$$

and it reduces to $\varphi = \varepsilon/(\varepsilon-1)$ in the absence of elasticity risk. Markup cyclicality therefore depends primarily on the cyclical behavior of cost and elasticity risk, with elasticity risk also affecting the risk-neutral benchmark price.

Corollary 5 (Markup cyclicality with CRRA utility, linear demand, and constant marginal cost)**.** *Consider the same environment as in Corollary 3, except that preferences are* *of the CRRA form*

$$U(\mathcal{W}_{it}) = \frac{\mathcal{W}_{it}^{1-\sigma}}{1-\sigma}, \qquad \sigma > 0,$$

so that absolute risk aversion satisfies $A_\sigma(\mathcal{W}_{it}) = \sigma/\mathcal{W}_{it} = \sigma A(\mathcal{W}_{it})$ *with* $A(\mathcal{W}_{it}) = 1/\mathcal{W}_{it}$*, where* $\mathcal{W}_{it}(\boldsymbol{\theta}_{it}) = \omega_{it}(\boldsymbol{\theta}_{it}) + \pi_{it}(p_{it};\boldsymbol{\theta}_{it})$ *is total wealth. In this case, the ex-ante markup satisfies*

$$\mu_{it}(\sigma) = \mu_{i,RN} + \frac{\sigma}{\kappa_i}\mathcal{R}_{it} + \mathcal{O}(\sigma^2),$$

where $\mathcal{R}_{it} = \mathcal{R}_{risk}(\boldsymbol{\sigma}_{it}^2) + \mathcal{R}_{wealth}(\mathcal{W}_{it,RN})$*, and*

$$\mathcal{R}_{risk}(\boldsymbol{\sigma}_{it}^2) = \overline{A}_{it}\mathcal{R}_{it}^{CARA}(\boldsymbol{\sigma}_{it}^2), \qquad \overline{A}_{it} := \mathbb{E}_{\boldsymbol{\theta}_{it}}\left[\frac{1}{\mathcal{W}_{it,RN}}\right],$$

$$\mathcal{R}_{wealth}(\mathcal{W}_{it,RN}) = \frac{1}{2}\cdot\frac{Cov_{\boldsymbol{\theta}_{it}}\left(\frac{1}{\mathcal{W}_{it,RN}}, \pi_{it,RN}\pi_{p,it,RN}\right)}{\mathbb{E}_{\boldsymbol{\theta}_{it}}[y_{p,it,RN}]}.$$

Here, $\mathcal{R}_{it}^{CARA}(\boldsymbol{\sigma}_{it}^2)$ *denotes the risk adjustment from Corollary 3.*

Consequently, firm-level markup cyclicality is

$$Cov(\mu_{it}, y_{it}) = \frac{\sigma}{\kappa_i}\left[Cov\left(\overline{A}_{it}\mathcal{R}_{it}^{CARA}(\boldsymbol{\sigma}_{it}^2), y_{it}\right) + Cov\left(\mathcal{R}_{wealth}(\mathcal{W}_{it,RN}), y_{it}\right)\right] + \mathcal{O}(\sigma^2).$$

Proof. See Appendix B.11. □

CRRA preferences: pure risk-scaling and wealth effects. Under CRRA preferences, precautionary pricing inherits the CARA structure through a pure-risk component and adds a wealth-effects component. The pure-risk channel preserves the same "risk kernel" as in the CARA benchmark; CRRA preferences scale this kernel by average absolute risk aversion, defined as the expected inverse of risk-neutral wealth. Hence, the *composition* of the pricing response to risk is unchanged, while its *strength* varies endogenously with wealth. In addition, CRRA preferences introduce a wealth-effects term that depends on how state-contingent risk aversion co-moves with profit exposure, summarized by the covariance between inverse wealth and the product of profits and marginal profits. Markup cyclicality therefore reflects three forces: time variation in risk components, time variation in effective risk aversion, and the covariance between wealth and profit exposure that governs wealth effects.

Additional risk-sharing instruments and the role of pricing. The CRRA case also clarifies when pricing is likely to play a quantitatively important role in managing risk. The wealth-effects channel depends on the joint distribution of total wealth and profit risk, and in particular on how non-profit income co-moves with profits and marginal profits. If entrepreneurs have access to additional risk-sharing instruments—such as diversified equity holdings, commodity and input hedging contracts, revenue insurance, or long-term contracts with cost pass-through clauses—then fluctuations in total wealth are partially offset. In this case, variation in effective risk aversion is dampened and the wealth-effects channel is attenuated.

Importantly, stabilizing wealth does not eliminate profit risk itself. As long as profits remain stochastic and cannot be fully hedged, pricing continues to respond to risk through the pure-risk channel. Pricing ceases to respond to risk only under complete markets for firm-level shocks, in which case both the pure-risk and wealth-effects channels vanish.

We next show that these conclusions extend beyond linear demand by considering an isoelastic demand system.

Corollary 6 (Markup cyclicality with CRRA utility, isoelastic demand, and constant marginal cost)**.** *Consider the same environment as in Corollary 4, except that preferences are of the CRRA form*

$$U(\mathcal{W}_{it}) = \frac{\mathcal{W}_{it}^{1-\sigma}}{1-\sigma}, \qquad \sigma > 0,$$

so that absolute risk aversion satisfies $A_\sigma(\mathcal{W}_{it}) = \sigma/\mathcal{W}_{it} = \sigma A(\mathcal{W}_{it})$ *with* $A(\mathcal{W}_{it}) = 1/\mathcal{W}_{it}$*, where* $\mathcal{W}_{it}(\boldsymbol{\theta}_{it}) = \omega_{it}(\boldsymbol{\theta}_{it}) + \pi_{it}(p_{it};\boldsymbol{\theta}_{it})$ *is total wealth. In this case, the ex-ante markup satisfies*

$$\mu_{it}(\sigma) = \mu_{it,RN} + \frac{\sigma}{\overline{\kappa}_i}\mathcal{R}_{it} + \mathcal{O}(\sigma^2),$$

where $\mathcal{R}_{it} = \mathcal{R}_{risk}(\boldsymbol{\sigma}_{it}^2) + \mathcal{R}_{wealth}(\mathcal{W}_{it,RN})$*, and*

$$\mathcal{R}_{risk}(\boldsymbol{\sigma}_t^2) = \overline{A}_{it}\mathcal{R}_{it}^{CARA}(\boldsymbol{\sigma}_{it}^2), \qquad \overline{A}_{it} := \mathbb{E}_{\boldsymbol{\theta}}\left[\frac{1}{\mathcal{W}_{it,RN}}\right],$$

$$\mathcal{R}_{wealth}(\mathcal{W}_{it,RN}) = \varphi_{it} \cdot \frac{Cov_{\boldsymbol{\theta}_{it}}\left(\frac{1}{\mathcal{W}_{it,RN}}, \pi_{it,RN}\pi_{p,it,RN}\right)}{\mathbb{E}_{\boldsymbol{\theta}_{it}}[y_{p,it,RN}]}.$$

Here, $\mathcal{R}^{CARA}(\boldsymbol{\sigma}_t^2)$ *denotes the risk adjustment from Corollary 4.*

Consequently, firm-level markup cyclicality is

$$Cov(\mu_{it}, y_{it}) = \frac{L_i}{(\varepsilon-1)^2}Cov(\sigma_{\psi,it}^2, y_{it}) + \frac{\sigma}{\overline{\kappa}_i}\left[Cov\left(\overline{A}_{it}\mathcal{R}_{it}^{CARA}(\boldsymbol{\sigma}_{it}^2), y_{it}\right) + Cov\left(\mathcal{R}_{wealth}(\mathcal{W}_{it,RN}), y_{it}\right)\right] + \mathcal{O}(\sigma^2).$$

Proof. For the firm i at date t, fix the shock distribution — including the firm-specific elasticity variance $\sigma^2_{\psi,it}$ — and let $p_{\text{RN},it}(\sigma^2_{\psi,it})$ denote the risk-neutral price that solves the RN FOC $\mathbb{E}_{\boldsymbol{\theta}_{it}}[\pi_p(p_{\text{RN}};\boldsymbol{\theta}_{it})] = 0$ under the isoelastic-demand, constant-MC specification of Corollary 4. Unlike the linear-demand case, $p_{\text{RN},it}$ depends on $\sigma^2_{\psi,it}$ through the variance of the elasticity shifter; the expansion below is in the CRRA curvature σ and treats $p_{\text{RN},it}(\sigma^2_{\psi,it})$ as a fixed benchmark *for firm i at date t*. Cross-sectional and time-series variation in $\sigma^2_{\psi,it}$ then enters the cyclicality decomposition below through the covariance of $p_{\text{RN},it}(\sigma^2_{\psi,it})$ with y_{it}.

With $U(\mathcal{W}) = \mathcal{W}^{1-\sigma}/(1-\sigma)$, $A_\sigma(\mathcal{W}) = \sigma/\mathcal{W} = \sigma \cdot A(\mathcal{W})$, $A(\mathcal{W}) = 1/\mathcal{W}$, the preference family satisfies A8 of Theorem 2. For sufficiently small σ, the theorem delivers

$$p_{it}(\sigma) = p_{\text{RN},it} + \sigma\,\varphi_{it}\left[\frac{\overline{A}_{it}\,\text{Cov}_{\boldsymbol{\theta}_{it}}(\pi_{it,\text{RN}}, \pi_{p,it,\text{RN}})}{\mathbb{E}_{\boldsymbol{\theta}_{it}}[y_{p,it,\text{RN}}]} + \frac{\text{Cov}_{\boldsymbol{\theta}_{it}}\left(\frac{1}{\mathcal{W}_{it,\text{RN}}}, \pi_{it,\text{RN}}\pi_{p,it,\text{RN}}\right)}{\mathbb{E}_{\boldsymbol{\theta}_{it}}[y_{p,it,\text{RN}}]}\right] + \mathcal{O}(\sigma^2),$$

where all expressions in the bracket are evaluated at the firm's own $p_{\text{RN},it}(\sigma^2_{\psi,it})$, $\overline{A}_{it} = \mathbb{E}_{\boldsymbol{\theta}_{it}}[1/\mathcal{W}_{it,\text{RN}}]$, and $\varphi_{it} = \mathbb{E}_{\boldsymbol{\theta}_{it}}[y_{p,it,\text{RN}}]/\mathbb{E}_{\boldsymbol{\theta}_{it}}[\pi_{pp,it,\text{RN}}]$. Identifying the pure-risk and wealth-effects components from the bracket,

$$\mathcal{R}_{\text{risk}}(\boldsymbol{\sigma}^2_{it}) = \varphi_{it}\cdot\overline{A}_{it}\cdot\frac{\text{Cov}_{\boldsymbol{\theta}_{it}}(\pi_{it,\text{RN}}, \pi_{p,it,\text{RN}})}{\mathbb{E}_{\boldsymbol{\theta}_{it}}[y_{p,it,\text{RN}}]} = \overline{A}_{it}\mathcal{R}^{\text{CARA}}_{it}(\boldsymbol{\sigma}^2_{it}),$$
$$\mathcal{R}_{\text{wealth}}(\mathcal{W}_{it,\text{RN}}) = \varphi_{it}\cdot\frac{\text{Cov}_{\boldsymbol{\theta}_{it}}\left(\frac{1}{\mathcal{W}_{it,\text{RN}}}, \pi_{it,\text{RN}}\pi_{p,it,\text{RN}}\right)}{\mathbb{E}_{\boldsymbol{\theta}_{it}}[y_{p,it,\text{RN}}]},$$

using Corollary 4's identity $\varphi_{it}\cdot\text{Cov}_{\boldsymbol{\theta}_{it}}(\pi_{it,\text{RN}}, \pi_{p,it,\text{RN}})/\mathbb{E}_{\boldsymbol{\theta}_{it}}[y_{p,it,\text{RN}}] = \mathcal{R}^{\text{CARA}}_{it}(\boldsymbol{\sigma}^2_{it})$. The markup expansion $\mu_{it}(\sigma) = p_{it}(\sigma)/\overline{\kappa}_i$ follows from the deterministic $\overline{\kappa}_i > 0$, with $\mu_{it,\text{RN}} = p_{\text{RN},it}(\sigma^2_{\psi,it})/\overline{\kappa}_i$.

For the cyclicality expression, take covariances with y_{it} on both sides of the markup expansion:

$$\text{Cov}(\mu_{it}(\sigma), y_{it}) = \text{Cov}(\mu_{it,\text{RN}}, y_{it}) + \frac{\sigma}{\overline{\kappa}_i}\text{Cov}(\mathcal{R}_{it}, y_{it}) + \mathcal{O}(\sigma^2).$$

The first term captures the $\sigma^2_{\psi,it}$-dependence of the risk-neutral benchmark: under isoelastic demand with stochastic elasticity, Corollary 4 gives $\mu_{it,\text{RN}} = \mu_{it,\text{RN}}(\sigma^2_{\psi,it})$, so a first-order Taylor expansion around the mean $\bar{\sigma}^2_\psi$ yields $\text{Cov}(\mu_{it,\text{RN}}, y_{it}) = \frac{L_i}{(\varepsilon-1)^2}\text{Cov}(\sigma^2_{\psi,it}, y_{it}) + \mathcal{O}(\bar{\sigma}^4_\psi)$, where L_i absorbs the local derivative $\partial\mu_{\text{RN}}/\partial\sigma^2_\psi$ at the benchmark. The second term gives the remaining covariance by bilinearity of $\mathcal{R}_{it} = \overline{A}_{it}\mathcal{R}^{\text{CARA}}_{it} + \mathcal{R}_{\text{wealth}}$ in the two summand channels. Combining the two yields the stated cyclicality expression.

The first-order-in-σ expansion of $\partial_\sigma U'_\sigma$ used here is exact under CARA and a leading-order approximation under CRRA by the remark following Theorem 2; the higher-order corrections are absorbed into the $\mathcal{O}(\sigma^2)$ remainder. □

4 Precautionary Pricing in Dynamic General Equilibrium

The pricing theory in Section 3 characterizes how risk-averse firms adjust markups in response to uncertainty, but leaves open three quantitative questions. First, how much of the precautionary pricing motive survives when firms can self-insure through savings? Second, what are the aggregate consequences for output, consumption, and welfare of uncertainty-driven markup fluctuations? Third, how do markups respond dynamically to uncertainty shocks, and how persistent are these responses? To address these questions, we embed the precautionary pricing mechanism in a dynamic general equilibrium model with heterogeneous firms, incomplete markets, and an explicit savings technology.

4.1 Environment

Time is discrete, $t = 0, 1, 2, \ldots$ The economy consists of a unit mass of entrepreneur-workers, indexed by $i \in [0, 1]$, and two production sectors: a differentiated sector in which each entrepreneur owns one firm producing a distinct variety, and a competitive outside sector producing a homogeneous numeraire good. The idiosyncratic shock vector $\boldsymbol{\theta}_{it} = (z_{it}, \xi_{it})$ collects a productivity shock z_{it} entering variable cost through a decreasing-returns labor technology introduced below (which maps to the cost shifter η of Section 3 via a log-linearization around $\bar{z}$), and a demand-level shifter ξ_{it}. We abstract from demand-elasticity shocks and use the Melitz–Ottaviano demand system.

Outside sector. The outside sector is perfectly competitive and produces the numeraire good with constant returns to scale, $c_{0t} = A\,\ell_{0t}$. This pins the wage at $w_t = A$ and serves as the money metric in which all other prices are denominated.

Differentiated sector. Each entrepreneur i operates a firm producing a distinct variety with a decreasing-returns labor technology $y_{it} = z_{it}\,\ell_{it}^{1-\eta}$, where ℓ_{it} is labor hired at wage w_t, z_{it} is an idiosyncratic productivity shock, and $\eta \in [0, 1)$ is a span-of-control parameter in the sense of Lucas (1978). Inverting the technology, the labor requirement to serve quantity y is $\ell(y,z) = (y/z)^{1/(1-\eta)}$, giving the variable cost function

$$\mathcal{C}(y_{it}; z_{it}) = w_t \left(y_{it}/z_{it}\right)^{1/(1-\eta)},$$

which is strictly convex in y for $\eta > 0$ and reduces to $w_t y_{it}/z_{it}$ at $\eta = 0$. Firm i faces demand $y_{it} = \mathcal{D}(p_{it}; \xi_{it})$ and earns profits

$$\pi(p_{it}; \boldsymbol{\theta}_{it}) = p_{it} y_{it} - w_t \left(y_{it}/z_{it}\right)^{1/(1-\eta)}.$$

Under linear demand $y = a + \xi - bp$, strict concavity of π in p (Assumption A3 of Theorem 1) holds for all $\eta \in [0,1)$, so Assumptions A1–A5 of Theorem 1 hold at each $\boldsymbol{\theta}$ in the dynamic GE environment. Once shocks realize, firms commit to supply the quantity demanded at their posted price.

Preferences. Entrepreneurs have time-additive preferences with discount factor $\beta \in (0,1)$, CARA utility over the numeraire good, and quasilinear additive subutility over the differentiated bundle:

$$\mathcal{U}_{it} = \mathbb{E}_t \sum_{s=0}^{\infty} \tilde{\beta}^s \left[-\exp(-\alpha\, c_{0,it+s}) + \int_0^1 v(c_{ji,t+s}, \xi_{j,t+s})\, dj \right], \qquad v(y,\xi) = \frac{(a+\xi)y}{b} - \frac{y^2}{2b}, \tag{1}$$

where $\alpha > 0$ is the CARA coefficient on numeraire consumption, $a, b > 0$ govern the Melitz–Ottaviano quadratic subutility v, and $\tilde{\beta} \equiv \beta(1-\delta_{\text{death}})$ is the effective discount factor under perpetual-youth death at rate $\delta_{\text{death}} \in [0,1)$ (Blanchard, 1985; Yaari, 1965). Death is realized at the end of each period, after consumption and savings; the dying entrepreneur is replaced at the start of the next period by a newborn with wealth $a_0 = 0$ and a fresh lag-shock draw $\boldsymbol{\theta}_{-1}^{\text{new}} \sim \pi^z \otimes \pi^\xi$ from the ergodic marginals. Bond holdings of dying agents are rebated lump-sum to surviving entrepreneurs via an actuarially fair annuity.

The CARA + quasilinear structure delivers two analytic simplifications that make the dynamic problem tractable. First, the variety FOC $\partial_y v(y,\xi) = p$ pins variety demand $y = \max(a + \xi - bp, 0)$ *independent of wealth*, so variety expenditure depends on $\boldsymbol{\theta}$ but not on entrepreneurial bond holdings. Second, CARA over numeraire consumption isolates the savings problem on the numeraire margin, while the variety-pricing problem inherits the static risk adjustment from Theorem 1 with absolute risk aversion equal to the CARA coefficient α. The pricing policy is therefore two-dimensional in (z_{-1}, ξ_{-1}) — independent of wealth — which exactly nests the analytic framework of Corollary 3.

Labor supply, bonds, and budget. Each entrepreneur supplies one unit of labor inelastically and earns wage income w_t. Labor is perfectly mobile between the two sectors. Entrepreneurs trade a one-period bond at price q_t in zero net supply, subject to a borrowing limit $a_{it} \geq \underline{a}$. Total resources after production are

$$\mathcal{W}(p_{it}; \boldsymbol{\theta}_{it}) \;=\; a_{it} + w_t + \pi(p_{it}; \boldsymbol{\theta}_{it}),$$

and the budget constraint is

$$c_{0it} + \int_0^1 p_{jt}\, c_{jit}\, dj + q_t\, a_{it+1} = \mathcal{W}(p_{it}; \boldsymbol{\theta}_{it}).$$

Quasilinearity pins the variety bundle from the variety FOC and leaves the residual budget — net of variety expenditure and savings — to be consumed as numeraire.

Timing. Each period proceeds in three stages:

1. Entrepreneurs enter with bond holdings a_{it} and observe lagged shocks $\boldsymbol{\theta}_{it-1}$. Each firm posts a price p_{it}, taking aggregate prices as given.
2. Idiosyncratic shocks $\boldsymbol{\theta}_{it}$ realize. Production and trade occur.
3. Entrepreneurs choose numeraire consumption c_{0it}, variety consumption $\{c_{jit}\}$, and bond holdings a_{it+1}.

Aggregate prices are determined in equilibrium to clear all markets.

4.2 Dynamic program

Let $\boldsymbol{s}_{it} = (a_{it}, \boldsymbol{\theta}_{it-1})$ denote the individual state. The entrepreneur's problem has two stages within each period: a pricing decision before shocks realize, and a consumption–savings decision after. Because firms commit to supply the quantity demanded at their posted price, the ex-ante pricing decision must internalize how p_{it} exposes profits to the full distribution of realized shocks.

Inner problem (intratemporal allocation and savings). After shocks $\boldsymbol{\theta}_{it}$ realize, the entrepreneur chooses variety consumption, numeraire consumption, and savings. Quasilinearity over the numeraire pins variety demand at the static FOC $v_y(c_{jit}, \xi_{jt}) = p_{jt}$, giving the wealth-independent linear schedule $c_{jit} = \max(a + \xi_{jt} - b\, p_{jt}, 0)$. Variety expenditure $\mathrm{VE}(\boldsymbol{\theta}_{it})$ therefore depends only on the realized shocks. The remaining numeraire-savings problem solves

$$\hat{V}(\mathcal{W}_{it};\boldsymbol{\theta}_{it}) = \max_{a_{it+1}\geq \underline{a}} \Big[-\exp\!\Big(-\alpha\,(\mathcal{W}_{it} - \mathrm{VE}(\boldsymbol{\theta}_{it}) - q_t\,a_{it+1})\Big) + \mathcal{V}^{*}(\boldsymbol{\theta}_{it}) + \tilde{\beta}\,\mathbb{E}[V(a_{it+1},\boldsymbol{\theta}_{it})]\Big],$$

where $c_{0it} = \mathcal{W}_{it} - \mathrm{VE}(\boldsymbol{\theta}_{it}) - q_t\,a_{it+1}$ is residual numeraire consumption and $\mathcal{V}^{*}(\boldsymbol{\theta}_{it})$ is the optimized variety subutility — wealth-independent under quasilinearity. The savings FOC is solved by an endogenous-grid step (**?**) on the inverse marginal utility, with the post-savings marginal utility of resources $\hat{V}_{\mathcal{W}}(\mathcal{W};\boldsymbol{\theta}) = \alpha\,\exp(-\alpha\,c_0^{*}(\mathcal{W};\boldsymbol{\theta}))$ inheriting the CARA exponential structure.

Outer problem (pricing). At the start of the period, before shocks realize, the firm posts a price that maximizes the expected continuation:

$$V(a_{it},\boldsymbol{\theta}_{it-1}) = \max_{p_{it}>0}\,\mathbb{E}_{\boldsymbol{\theta}_{it}|\boldsymbol{\theta}_{it-1}}\Big[\hat{V}\Big(\mathcal{W}_{it}(p_{it};\boldsymbol{\theta}_{it}),\boldsymbol{\theta}_{it}\Big)\Big]. \tag{2}$$

The pricing first-order condition is the dynamic analogue of Theorem 1:

$$\mathbb{E}_{\boldsymbol{\theta}_{it}|\boldsymbol{\theta}_{it-1}}\Big[\,\hat{V}_{\mathcal{W}}\Big(\mathcal{W}_{it}(p^{*}_{it};\boldsymbol{\theta}_{it}),\boldsymbol{\theta}_{it}\Big)\cdot\pi_p(p^{*}_{it};\boldsymbol{\theta}_{it})\Big] = 0.$$

The CARA + quasilinear structure delivers a closed form for $\hat{V}_{\mathcal{W}}$ that we exploit below: the inner value function inherits the CARA exponential form in total resources $\mathcal{W}$, with effective absolute risk aversion attenuated by the savings response.

4.3 Savings and effective risk aversion

The pricing first-order condition (2) has the same structure as Theorem 1, with $\hat{V}_{\mathcal{W}}$ replacing U'. The static decomposition by source (Proposition 1) carries through to the dynamic environment with the static absolute risk aversion $-U''/U'$ replaced by the dynamic effective curvature $-\hat{V}_{\mathcal{W}\mathcal{W}}/\hat{V}_{\mathcal{W}}$. Under CARA + quasilinear preferences, this dynamic curvature has a closed-form expression in terms of the primitive α and the equilibrium bond price q: savings substitute partially for pricing as a risk-management device, attenuating effective risk aversion by the fraction of a wealth shock not absorbed in the current period.

Proposition 6 (Effective risk aversion with savings under CARA + quasilinear preferences)**.** *Consider preferences* (1) *with CARA period utility* $U(c_0) = -\exp(-\alpha\,c_0)$*,* $\alpha > 0$*, over the numeraire and the quasilinear MO subutility over the variety bundle. Consider the inner problem in steady state with constant aggregate prices and a bond price* $q \in (0,1)$*, and assume* *the borrowing constraint is slack. Then the inner value function inherits the CARA form in total resources:*

$$\hat{V}(\mathcal{W};\boldsymbol{\theta}) = -B(\boldsymbol{\theta})\exp(-\alpha_{\text{eff}}\mathcal{W}),$$

with $B(\boldsymbol{\theta}) > 0$*, and the effective absolute risk aversion with respect to total resources is*

$$\alpha_{\text{eff}} = \alpha\,(1-q) < \alpha.$$

Proof. Quasilinearity pins variety consumption independent of wealth, so variety expenditure $\text{VE}(\boldsymbol{\theta})$ depends only on $\boldsymbol{\theta}$ and the numeraire budget identity reads $c_0 = \mathcal{W} - \text{VE}(\boldsymbol{\theta}) - q\,a'$. Make the joint guess

$$\hat{V}(\mathcal{W};\boldsymbol{\theta}) = -B(\boldsymbol{\theta})\exp(-\alpha_{\text{eff}}\mathcal{W}), \qquad V(a;\boldsymbol{\theta}_{-1}) = -D(\boldsymbol{\theta}_{-1})\exp(-\alpha_{\text{eff}}\,a),$$

with $B, D > 0$ to be determined. The inner problem becomes

$$\hat{V}(\mathcal{W};\boldsymbol{\theta}) = \max_{a'}\Big\{-\exp\Big(-\alpha(\mathcal{W} - \text{VE}(\boldsymbol{\theta}) - qa')\Big) + \mathcal{V}^{*}(\boldsymbol{\theta}) - \tilde{\beta}\,\widetilde{D}(\boldsymbol{\theta})\exp(-\alpha_{\text{eff}}\,a')\Big\},$$

where $\mathcal{V}^{*}(\boldsymbol{\theta})$ is the optimal variety-subutility value (independent of $\mathcal{W}$ and a') and $\widetilde{D}(\boldsymbol{\theta}) := \mathbb{E}\Big[D(\boldsymbol{\theta}') \mid \boldsymbol{\theta}\Big]$. The first-order condition in a' is

$$\alpha\,q\,\exp\Big(-\alpha(\mathcal{W} - \text{VE}(\boldsymbol{\theta}) - qa'^{*})\Big) = \tilde{\beta}\,\alpha_{\text{eff}}\,\widetilde{D}(\boldsymbol{\theta})\exp(-\alpha_{\text{eff}}\,a'^{*}).$$

Taking logarithms and rearranging gives a linear-in-$\mathcal{W}$ savings policy,

$$a'^{*}(\mathcal{W};\boldsymbol{\theta}) = \frac{\alpha}{\alpha_{\text{eff}} + \alpha q}\Big(\mathcal{W} - \text{VE}(\boldsymbol{\theta})\Big) + \frac{1}{\alpha_{\text{eff}} + \alpha q}\log\!\left(\frac{\tilde{\beta}\,\alpha_{\text{eff}}\,\widetilde{D}(\boldsymbol{\theta})}{\alpha q}\right).$$

Substituting into the budget identity gives numeraire consumption linear in $\mathcal{W}$,

$$c_0^{*}(\mathcal{W};\boldsymbol{\theta}) = \frac{\alpha_{\text{eff}}}{\alpha_{\text{eff}} + \alpha q}\Big(\mathcal{W} - \text{VE}(\boldsymbol{\theta})\Big) - \frac{q}{\alpha_{\text{eff}} + \alpha q}\log\!\left(\frac{\tilde{\beta}\,\alpha_{\text{eff}}\,\widetilde{D}(\boldsymbol{\theta})}{\alpha q}\right).$$

Applying the envelope theorem to the inner problem and using the first-order condition (so the $\partial a'^{*}/\partial\mathcal{W}$ terms cancel) gives $\hat{V}_{\mathcal{W}}(\mathcal{W};\boldsymbol{\theta}) = U'(c_0^{*}) = \alpha\,\exp(-\alpha\,c_0^{*})$. Matching this to the derivative of the guess, $\alpha_{\text{eff}}B(\boldsymbol{\theta})\exp(-\alpha_{\text{eff}}\mathcal{W})$, requires the exponential rates in $\mathcal{W}$ to coincide:

$$\alpha\,\frac{\alpha_{\text{eff}}}{\alpha_{\text{eff}} + \alpha q} = \alpha_{\text{eff}} \quad\Longleftrightarrow\quad \alpha_{\text{eff}} + \alpha q = \alpha \quad\Longleftrightarrow\quad \alpha_{\text{eff}} = \alpha(1-q).$$

Given this value of α_{eff}, substituting c_0^* and the FOC back into the inner problem yields the closed-form level

$$B(\boldsymbol{\theta}) \;=\; \frac{1}{1-q}\,\exp\!\Big(\alpha_{\text{eff}}\,\text{VE}(\boldsymbol{\theta})\Big)\left[\frac{\tilde{\beta}\,\alpha_{\text{eff}}\,\widetilde{D}(\boldsymbol{\theta})}{\alpha q}\right]^{q} \;>\; 0,$$

and the additive variety-subutility constant $\mathcal{V}^*(\boldsymbol{\theta})$, being $\mathcal{W}$-independent, contributes to the level of $\hat{V}$ but not to its curvature in $\mathcal{W}$ and hence does not enter $\hat{V}_{\mathcal{W}}$ or the pricing FOC. Finally, since total resources are affine in bond holdings, $\mathcal{W} = a + w + \pi(p^*;\boldsymbol{\theta})$, the guess factorizes as $\hat{V}(\mathcal{W};\boldsymbol{\theta}) = -B(\boldsymbol{\theta})\exp(-\alpha_{\text{eff}}(w+\pi(p^*;\boldsymbol{\theta})))\cdot\exp(-\alpha_{\text{eff}}\,a)$. Taking expectations over $\boldsymbol{\theta}\mid\boldsymbol{\theta}_{-1}$ yields $V(a;\boldsymbol{\theta}_{-1}) = -D(\boldsymbol{\theta}_{-1})\exp(-\alpha_{\text{eff}}\,a)$ with $D(\boldsymbol{\theta}_{-1}) := \mathbb{E}[B(\boldsymbol{\theta})\exp(-\alpha_{\text{eff}}(w+\pi(p^*(\boldsymbol{\theta}_{-1});\boldsymbol{\theta})))\mid\boldsymbol{\theta}_{-1}] > 0$, closing the guess. The condition $q\in(0,1)$ ensures $\alpha_{\text{eff}}\in(0,\alpha)$ and is consistent with the standard precautionary-savings bound $\tilde{\beta}\le q<1$ in equilibrium. □

Interpretation. Without savings ($q=0$), the entrepreneur consumes everything net of variety expenditure, and effective risk aversion with respect to $\mathcal{W}$ coincides with the primitive α; with savings, $\alpha_{\text{eff}} = \alpha(1-q)$ attenuates by the fraction of a wealth shock absorbed in the current period. Combined with Theorem 1 and the CARA closed forms of Corollary 3, this implies that the static risk adjustment $\mathcal{R}^{\text{static}}$ rescales uniformly to a dynamic counterpart $\mathcal{R}^{\text{dyn}} = (1-q)\mathcal{R}^{\text{static}}$. The mapping from primitives to markups in the dynamic GE model is therefore the static formula of Corollary 3 with α replaced by $\alpha(1-q)$.

Extension to time-varying q_t. Along an MIT transition path with a time-varying sequence $\{q_t\}$, the ansatz $\hat{V}_t(\mathcal{W};\boldsymbol{\theta}) = -B_t(\boldsymbol{\theta})\exp(-\alpha_{\text{eff},t}\mathcal{W})$ closes period by period with $\alpha_{\text{eff},t} = \alpha(1-q_t)$, so the dynamic risk-adjustment relation $\mathcal{R}_t^{\text{dyn}} = (1-q_t)\mathcal{R}_t^{\text{static}}$ holds at every date.

4.4 Equilibrium

Let $\Phi_t(a,\boldsymbol{\theta}_{-1})$ denote the cross-sectional distribution of entrepreneurs over bond holdings and lagged shocks at the beginning of period t.

Definition 2 (Stationary equilibrium). *The outside good is the numeraire with price normalized to one. A stationary equilibrium consists of aggregate prices (w,q); a value function $V(a,\boldsymbol{\theta}_{-1})$; policy functions $p(\boldsymbol{\theta}_{-1})$ and $a'(\mathcal{W},\boldsymbol{\theta})$; variety demand $\{c_j(p_j;\boldsymbol{\theta})\}$; a stationary distribution $\Phi(a,\boldsymbol{\theta}_{-1})$ over entrepreneurs with total mass one; and labor allocations (L_d,L_0) across the differentiated and outside sectors such that $L_d+L_0=1$. The equilibrium satisfies:*

1. ***Optimality.*** *V, p, a', and variety demand solve the entrepreneur's problem.*
2. ***Labor market.*** *The wage w clears the labor market:*

$$\int \mathbb{E}_{\boldsymbol{\theta}|\boldsymbol{\theta}_{-1}}[\ell(p(\boldsymbol{\theta}_{-1});\boldsymbol{\theta})]\,d\Phi(a,\boldsymbol{\theta}_{-1}) + L_0 = 1.$$

3. ***Bond market.*** *Bonds are in zero net supply:*

$$\int \mathbb{E}_{\boldsymbol{\theta}|\boldsymbol{\theta}_{-1}}\Big[a'(\mathcal{W}(p;\boldsymbol{\theta});\boldsymbol{\theta})\Big]\,d\Phi(a,\boldsymbol{\theta}_{-1}) = 0.$$

4. ***Variety markets.*** *For each variety, supply equals aggregate demand.*
5. ***Outside good.*** *Aggregate numeraire consumption equals outside-sector output:*

$$\int c_{0i}\,d\Phi \;=\; A\,L_0.$$

6. ***Stationarity.*** *Φ is a fixed point of the law of motion that advances (z_{-1},ξ_{-1}) under the Rouwenhorst shock transition and the savings policy, mixing in newborns at rate δ_{death}:*

$$\Phi' \;=\; (1-\delta_{death})\cdot \mathcal{T}_{p,a'}\Phi \;+\; \delta_{death}\cdot \Phi_{newborn},$$

where $\Phi_{newborn}(a=0,z,\xi) = \pi^z(z)\,\pi^\xi(\xi)$ places newborns at wealth $a_0 = 0$ with shocks drawn from the ergodic marginals.

By Walras' law, one of the market clearing conditions is redundant: the outside-good resource constraint (condition 5) follows automatically from labor-market clearing, bond-market clearing, variety-market clearing, and the sum of entrepreneur budget constraints. We hence drop condition 5 when computing the equilibrium numerically.

4.5 From static PE to dynamic GE

The dynamic GE model of this section extends the precautionary pricing theory of Section 3 along three dimensions. First, entrepreneurs have access to a savings technology that partially substitutes for pricing as a risk-management tool. Second, the pricing decision interacts with general-equilibrium determination of wages, interest rates, and the wealth distribution. Third, the model generates a stationary wealth distribution and allows for transition dynamics following uncertainty shocks.

Despite these extensions, the pricing mechanism of Section 3 carries over intact. The dynamic pricing first-order condition has the same structure as Theorem 1, with the static marginal utility $U'(\mathcal{W})$ replaced by the marginal value of total resources $\hat{V}_{\mathcal{W}}$. Consequently, Theorem 2, Proposition 1, Proposition 2, and the closed-form corollary for CARA + linear demand (Corollary 3) continue to apply, provided the primitive risk aversion α is replaced by its dynamic counterpart $\alpha_{\text{eff}} = \alpha(1 - q)$ from Proposition 6. In particular, the decomposition of markup cyclicality by source of uncertainty, the sign conditions, and the role of cost pass-through all remain valid.

The substantive modification relative to Section 3 is quantitative: savings attenuate effective risk aversion by the factor $(1 - q)$, so the static risk adjustment $\mathcal{R}^{\text{static}}$ rescales uniformly to $\mathcal{R}^{\text{dyn}} = (1 - q)\mathcal{R}^{\text{static}}$ and the static closed-form corollaries port directly. The relative strength of the cost and demand channels is preserved: cost uncertainty still raises markups, demand uncertainty still lowers them, and their ratio is governed by the same demand-side objects as in the static theory.

This observation motivates our empirical strategy. Because the signs and relative magnitudes of the uncertainty channels are invariant across the static and dynamic models, we can test the qualitative predictions of the theory directly in the data without taking a stand on the full dynamic GE structure. Section 5 implements this test using shift-share measures of industry-specific cost and demand uncertainty. The structural calibration of the dynamic GE model that maps these reduced-form slopes into structural primitives is work in progress (Section 6).

5 Testing the Theory's Predictions: Empirical Evidence

This section tests the qualitative predictions of the precautionary pricing mechanism developed in Section 3. The theory implies that different sources of uncertainty have opposing effects on markups: cost uncertainty raises prices, while demand uncertainty lowers them. These sign predictions hold under strictly concave preferences and for a broad class of demand systems (Proposition 1). What varies across environments is the magnitude of the response, not its direction.

We proceed in two steps. First, we describe the data sources and the construction of shift-share (Bartik-style) measures of industry-specific cost and demand uncertainty. Second, we estimate industry-level panel regressions of prices on uncertainty. The structural calibration of the dynamic GE model that maps these reduced-form slopes into structural primitives is work in progress and is outlined in Section 6.

5.1 Data

Our empirical analysis combines industry-level price and cost data from the NBER-CES Manufacturing Industry Database with input-output linkages from the Bureau of Economic Analysis (BEA) Detail Use Tables. We construct Bartik-style shifters that measure each industry's exposure to upstream cost conditions and downstream demand conditions through its position in the production network. Uncertainty is measured as the rolling volatility of innovations to these shifters, providing industry-specific, time-varying measures of cost and demand uncertainty.

Industry prices and costs. We use the NBER-CES Manufacturing Industry Database, which provides annual data on output prices, input prices, and production characteristics for 6-digit NAICS manufacturing industries from 1958 to 2018.[5] We measure output prices using the shipments price deflator and construct input costs as a weighted average of materials and energy prices. Although the raw database begins in 1958, the estimation panel is restricted to years covered by the Bartik IO weights (earliest BEA detail benchmark: 1972); AR(1) pre-whitening and the three-year rolling variance window then consume the first four years, so the usable panel starts in 1976. The final estimation panel covers 335 manufacturing industries (6-digit NAICS) from 1976 to 2018 with 12,795 industry-year observations; see Table 1 below.

Bartik shifters. The theory predicts that the pricing response to uncertainty depends on industry-specific exposure to cost and demand risk. We construct Bartik-style shifters from the BEA Benchmark Detail Use Tables. We use all ten benchmark vintages available at the detail level spanning the usable NBER-CES sample—1972, 1977, 1982, 1987, 1992, 1997, 2002, 2007, 2012, and 2017—and concord them to a common NAICS 2007+ detail IO classification covering 231 manufacturing codes. Pre-NAICS vintages (1972–1992) use BEA IO codes based on successive revisions of the Standard Industrial Classification; we bridge each vintage to the 2007+ classification by first mapping its BEA IO codes to SIC-4 via the sectoring plan published with the benchmark, then chaining SIC-72 → SIC-77 → SIC-87 → NAICS-97 via the Fort-Klimek (Fort and Klimek, 2018) concordances, and finally mapping NAICS-97 to the 2007+ IO classification via BEA's official crosswalk. Agriculture (not covered by Fort-Klimek) uses a supplementary SIC → NAICS bridge constructed from the Census Bureau concordance. The 1997–2017 benchmarks are already on NAICS-based IO classifications. Between benchmark years we linearly interpolate the detail-level use shares, yielding time-varying input-output weights at annual frequency. The estimation sample is restricted to 1972 onwards to

[5]We use the 2021 vintage, which provides consistent industry definitions based on the 2012 NAICS classification. See Bartelsman and Gray (1996) for documentation.

avoid held-flat extrapolation before the earliest benchmark; the first four years are consumed by AR(1) pre-whitening, the three-year rolling variance, and one-period lagging of the uncertainty regressor. For each manufacturing industry i and year t, we compute input-output weights:

$$\omega_{ji,t}^{\text{supplier}} = \frac{\text{Use}_t(j,i)}{\sum_{j'} \text{Use}_t(j',i)}, \qquad \omega_{ij,t}^{\text{customer}} = \frac{\text{Use}_t(i,j)}{\sum_{j'} \text{Use}_t(i,j')},$$

where $\text{Use}_t(j,i)$ denotes the dollar value of commodity j used by industry i in year t, built by linear interpolation between adjacent benchmark years.

Node-level primitives. To avoid confounding demand (cost) shifters with realised supply (demand) conditions at counterparty industries, we project each industry's node-level shock onto exogenous aggregate primitives before propagating through the IO network. On the cost side, supplier j's node-level cost shock is an industry-specific Bartik over three input-price primitives — non-energy materials, energy, and wages — with shares given by industry j's variable-cost composition in NBER-CES:

$$C_{j,t}^{\text{IN}} = \text{share}_j(\text{mat}) \cdot \Delta \log P_t^{\text{mat}} + \text{share}_j(\text{ener}) \cdot \Delta \log P_t^{\text{ener}} + \text{share}_j(\text{wage}) \cdot \Delta \log P_t^{\text{wage}}.$$

The NBER-CES `matcost` variable is inclusive of energy (as in the 4-factor TFP construction); we therefore use non-energy materials `matcost` – `energy` in the numerator of $\text{share}_j(\text{mat})$ so that the three variable-cost shares are disjoint and sum to one, following the NBER-CES 5-factor TFP convention. The aggregate primitive series are P_t^{mat} (BLS PPI Industrial Commodities), P_t^{ener} (BLS PPI Energy), and P_t^{wage} (BLS Average Hourly Earnings, extrapolated pre-1964 via CPI). On the demand side, customer j's node-level demand shock is the Bartik projection onto the dominant final-demand category, personal consumption expenditures:

$$D_{j,t}^{\text{FD}} = \text{share}_{j,t}(\text{PCE}) \cdot \Delta \log \text{PCE}_t,$$

with $\text{share}_{j,t}(\text{PCE})$ the time-t share of commodity j's total output (intermediate use plus final demand) that flows to PCE. FD shares are built from all ten BEA detail benchmark Use tables (1972–2017), with native commodity codes concorded to the 2007+ IO classification via the same chain used for the supplier/customer weights, and linearly interpolated to annual frequency between benchmarks. The aggregate primitive $\Delta \log \text{PCE}_t$ is the log-growth rate of real aggregate PCE from FRED. Three choices in this construction are substantive and warrant explanation.

(i) Log-growth rather than log-level aggregates. We use $\Delta \log P_t^m$ and $\Delta \log \text{PCE}_t$ rather than $\log P_t^m$ and $\log \text{PCE}_t$. Aggregate final-demand and price series have strong secular trends over the 1958–2018 sample — real PCE grows roughly 3% per year, industrial-

commodity PPI roughly doubles each decade. AR(1) pre-whitening of first differences (next paragraph) is designed to remove *cyclical* persistence in the shifter, not a deterministic trend. When the node-level primitive is expressed in log-levels, the resulting shifter inherits a secular trend that the AR(1) residualization only partially absorbs, leaving residuals that are a mixture of trend remainder and cyclical innovations. Differencing the log aggregates before building the shifter places the node-level shock directly in growth-rate units, so the downstream AR(1) step isolates pure cyclical innovations. This choice is the single largest reason the firm-level demand channel is identified cleanly in our preferred specification; using log-level aggregates drives $\hat{\gamma}_{\text{demand}}$ toward zero or the wrong sign, consistent with residualized trend contamination.

(ii) PCE only on the demand side. Personal consumption expenditures is the final-demand category most closely aligned with the demand-uncertainty object the theory speaks to: consumption-driven demand volatility that firms cannot cost-minimize their way around. Investment, exports, and government expenditure move for reasons that are largely disconnected from the consumer-demand risk the model emphasises. Empirically, aggregating across all four final-demand categories (PCE + Investment + Exports + Government) flips the sign of $\hat{\gamma}_{\text{demand}}$ because procyclical Investment and Exports dominate the composite Bartik. We therefore restrict the demand-side primitive to PCE, which isolates the channel the theory targets and preserves the predicted sign. Results using the broader final-demand aggregation are reported as a robustness check in Appendix C.

(iii) NAICS-3 own-sector exclusion. On both sides, the IO-weighted sum below excludes counterparties j with $n_3(j) = n_3(i)$ — i.e., we drop supplier and customer industries in the firm's own 3-digit NAICS sector. This is the exclusion standard in the shift-share design of Goldsmith-Pinkham, Sorkin and Swift (2020): it ensures the Bartik shifter is predetermined with respect to the firm's own pricing decision, because a firm's NAICS-3 sector itself is subject to common latent shocks that would otherwise bias the identification. Tighter exclusion (NAICS-4 or no exclusion) mechanically loads more variation and delivers larger $\hat{\gamma}$'s, but at the cost of importing the own-sector contamination the design is meant to purge. The NAICS-3 exclusion is our baseline; alternative exclusions are reported in Appendix C.

IO propagation. The *upstream cost shifter* is the IO-weighted average of supplier node-level cost shocks, excluding suppliers in the firm's own NAICS-3 sector:

$$C_{it} = \sum_{j:\, n_3(j) \neq n_3(i)} \omega_{ji,t}^{\text{supplier}} \cdot C_{j,t}^{\text{IN}}.$$

The *downstream demand shifter* is constructed symmetrically from customer node-level demand shocks:

$$D_{it} = \sum_{j:\, n_3(j) \neq n_3(i)} \omega_{ij,t}^{\text{customer}} \cdot D_{j,t}^{\text{FD}}.$$

The exclusion of own NAICS-3 ensures that the shifters capture variation transmitted through the production network rather than own-industry conditions. Relative to a standard one-stage Bartik (which would shock the network with counterparties' realised output or unit cost), the two-stage construction is symmetric on the two sides and cleaner: both cost and demand shifters inherit only the components of their counterparties' outcomes that are traceable to exogenous factor-price and final-demand shocks, purging the demand–supply confound that would otherwise arise from using realised quantities and prices as shocks.

Uncertainty measures. We construct industry-specific, time-varying measures of uncertainty from the Bartik shifters. The shifters S_{it} are constructed from log-growth primitives and are therefore already in growth-rate (stationary) units. We remove predictable persistence in each shifter by fitting an AR(1) per IO code in the levels of the shifter and taking the residual as the innovation,

$$u_{S,it} = S_{it} - \hat{\phi}_S \, S_{i,t-1}, \qquad S \in \{C, D\}.$$

Applying AR(1) to ΔS_{it} instead would impose a second difference on a shifter that is already a growth rate, injecting spurious MA structure; levels are the correct specification given the dlog-primitive construction above. These residuals isolate shifter innovations relative to information available at $t-1$.

Before computing the rolling-window variance, we *winsorize* the innovations at the 2.5th and 97.5th percentiles (pooled across the panel). The motivation is structural rather than ad hoc: rolling-window variance is a quadratic-in-residuals object, so a single outlier observation contributes its squared residual to every window it falls in. Diagnostically, the raw residual distribution is heavily right-skewed — a handful of oil-shock and crisis-year observations account for more than 80% of the cross-sectional variance of $u_{S,it}^2$. Without winsorization, the rolling-variance series therefore reflects the timing of a small number of extreme shocks rather than the level of conditional volatility that the theory's σ^2 object is meant to proxy. Winsorization caps the leverage of these extreme points while preserving the full panel (unlike trimming, which would drop observations and bias the panel toward tranquil years). The threshold 2.5/97.5 is standard in the empirical uncertainty literature; results are qualitatively identical under 1/99 or 5/95 and are reported in Appendix C.

Denoting the winsorized residuals by $\tilde{u}_{S,it}$, we define uncertainty as the rolling-window variance,

$$\hat{\sigma}^2_{S,it} = \text{Var}(\tilde{u}_{S,i,t-W+1}, \ldots, \tilde{u}_{S,it}), \tag{3}$$

with $W = 3$ years and standardize each series by its pooled sample standard deviation so that the reduced-form slopes below are directly interpretable as the response of $\log p$ to a one-standard-deviation increase in uncertainty. In the regression specification below we follow the convention of Section 3 and denote the cost-shifter uncertainty measure $\hat{\sigma}^2_{\eta,it} \equiv \hat{\sigma}^2_{C,it}$ and the demand-shifter uncertainty measure $\hat{\sigma}^2_{\xi,it} \equiv \hat{\sigma}^2_{D,it}$.

The reduced-form slopes $\hat{\gamma}_{\text{cost}}$ and $\hat{\gamma}_{\text{demand}}$ recover the model-implied derivatives $\partial \log p/\partial\sigma^2_\eta$ and $\partial \log p/\partial\sigma^2_\xi$ up to two loadings. First, the rolling variance of a Bartik-shifter innovation is a convolution of innovations to the underlying sectoral unit-cost and real-output series, weighted by the detail IO shares; the mapping from $\hat{\sigma}^2_{S,it}$ to the variance of the structural shock (η, ξ) hitting industry i is therefore a quadratic function of those IO weights rather than an identity. Second, under the dynamic GE model of Section 4 the static pricing slope is attenuated by $1 - q$ through the savings channel (Proposition 6). The structural calibration of the dynamic GE model that absorbs both loadings into a single closed-form mapping from $(\hat{\gamma}_{\text{cost}}, \hat{\gamma}_{\text{demand}})$ to the primitive preferences and demand parameters via Corollary 3 is work in progress (Section 6).

This construction separates second-moment risk from predictable first-moment movements and yields a proxy for the conditional variance relevant for pricing decisions. The approach is similar in spirit to Jurado, Ludvigson and Ng (2015), who define uncertainty as the volatility of forecast errors rather than realized outcomes. Because the Bartik shifters embed predetermined input-output linkages, the resulting uncertainty measures vary both cross-sectionally—through differences in industry exposure—and over time—through innovations to costs and demand.

Unlike aggregate proxies such as oil price volatility or sentiment indices, these measures capture heterogeneous uncertainty exposure across industries arising from their position in the production network. Figure 4 plots the aggregate Bartik uncertainty measures over time. Both cost and demand uncertainty spike during well-known crises, including the 1970s and early-1980s oil shocks and the 2008–09 financial crisis.

Figure 4: Aggregate Bartik Uncertainty Over Time

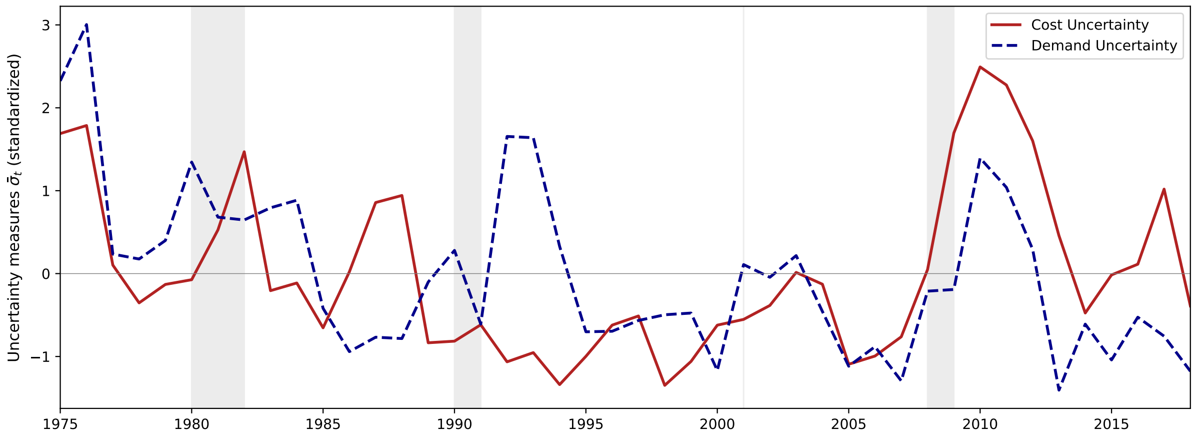

Notes: Value-weighted average of industry-specific Bartik uncertainty measures (rolling variance of AR(1) residuals, $W = 3$). Gray bands indicate NBER recession dates.

Summary statistics. After concordance, matching to NBER-CES, and the NAICS-3 own-sector exclusion in the Bartik construction, the panel retains 209 of the 231 manufacturing detail IO codes, which we match to 335 NAICS-6 industries over 1976–2018, yielding 12,795 industry-year observations. Table 1 reports summary statistics. There is substantial cross-industry heterogeneity in both cost and demand uncertainty, driven by differences in IO linkages.

Table 1: Summary Statistics

Variable	Mean	SD	P10	P50	P90
Log price	-0.3303	0.4872	-0.8510	-0.3052	0.0315
Cost uncertainty	0.0254	1.0021	-0.6249	-0.4150	1.5744
Demand uncertainty	0.0695	1.0801	-0.4154	-0.3773	1.0497
Upstream cost level	-0.3980	0.2850	-0.8117	-0.3307	-0.0798
Downstream demand level	3.0206	2.6775	0.0698	2.3157	7.0868
Industries (NAICS-6)			335		
IO codes			209		
Years			1976–2018		
Observations			12,795		

Notes: Panel of manufacturing industries from NBER-CES. Uncertainty measures are Bartik shift-share (detail IO weights × IO-code AR(1) innovations, winsorized at 2.5/97.5, rolling variance with window $W = 3$), z-scored by the pooled in-sample standard deviation. Level controls are IO-weighted log upstream cost and log downstream demand, with own NAICS-3 excluded. All regressors enter the main regression lagged one year.

5.2 Panel Evidence

Our main specification relates industry-level log prices to lagged uncertainty (in variance units) and lagged level controls:

$$\log p_{it} = \gamma_{\text{cost}}\,\hat{\sigma}^2_{\eta,i,t-1} + \gamma_{\text{demand}}\,\hat{\sigma}^2_{\xi,i,t-1} + \beta_C\,C_{i,t-1} + \beta_D\,D_{i,t-1} + \alpha_i + \varepsilon_{it}. \tag{4}$$

Here, $\hat{\sigma}^2_{\eta,i,t-1}$ and $\hat{\sigma}^2_{\xi,i,t-1}$ are the lagged rolling-window variance measures defined in equation (3) (using the identification $\eta \equiv C$, $\xi \equiv D$), $C_{i,t-1}$ and $D_{i,t-1}$ denote the lagged levels of the Bartik cost and demand shifters, and α_i are industry fixed effects. Because the regressor is in variance units, the estimated slopes $\hat{\gamma}_{\text{cost}}$ and $\hat{\gamma}_{\text{demand}}$ match the functional form of $\partial \log p/\partial\sigma^2$ in Proposition 1.

Uncertainty is lagged to match the model's timing: firms set prices before shocks are realized, so pricing depends on the conditional variance known at the time of decision. Lagging also mitigates mechanical contemporaneous feedback between prices and realized shocks.

We include lagged Bartik *levels* $C_{i,t-1}$ and $D_{i,t-1}$ as controls because prices respond to both the level and the volatility of cost and demand conditions. Because level and volatility measures are constructed from the same underlying shocks, omitting level controls would conflate first-moment pass-through with second-moment effects. By partialling out responses to cost and demand levels, we ensure that γ_{cost} and γ_{demand} isolate the precautionary pricing channel.

This specification is a direct empirical analogue of Proposition 1, which shows that optimal prices respond linearly to uncertainty. The theory predicts $\gamma_{\text{cost}} > 0$: when cost uncertainty rises, firms raise prices to limit exposure to high-cost realizations; and $\gamma_{\text{demand}} < 0$: when demand uncertainty rises, firms lower prices to reduce exposure to weak demand realizations. Under risk neutrality, prices respond only to cost and demand levels, with no additional response to volatility, implying $\gamma_{\text{cost}} = \gamma_{\text{demand}} = 0$.

Identification exploits the Bartik structure in the sense of Goldsmith-Pinkham, Sorkin and Swift (2020). Input-output weights are taken from the ten BEA benchmark detail tables (1972–2017) concorded to a common NAICS 2007+ classification and linearly interpolated between benchmark years; they move slowly, primarily reflecting long-run shifts in the production network rather than short-run pricing conditions. Shocks are constructed at a coarser NAICS-3 level to ensure reliable aggregates for unit cost and output; finer classifications yield noisy series due to thin cells. Cross-sectional variation in the shifters comes from the detail IO weights, which differ across industries even within the same NAICS-3 sector. Own-sector shocks are excluded, ensuring that the shifters are predetermined with respect to the industry's own pricing decisions. Because the Bartik measures combine cross-sectional variation in IO exposure with time-series variation in sectoral shocks, year fixed effects would absorb the common shock component and leave only cross-sectional differences in weights for identification; we omit them in the baseline to exploit both sources of variation.

Table 2 presents the results. Both coefficients have the predicted signs and are highly significant: $\hat{\gamma}_{\text{cost}} = 0.058^{***}$ (SE 0.006) and $\hat{\gamma}_{\text{demand}}$ =-0.061*** (SE 0.012). Standard errors are exposure-robust in the sense of Borusyak, Hull and Jaravel (2022), clustered at the IO code (the unit at which the underlying shocks are defined) for the levels columns; first-difference columns instead report two-way clustering by IO code × year, because differencing amplifies within-year common shocks and strains the single-IO-code independence assumption. First-stage partial F-statistics for the cost-level and demand-level instruments are 1,461 and 701, respectively, so weak-instrument concerns are immaterial. The sign pattern is robust to alternative specifications in columns (1)–(5), including variants with year fixed effects. In one-standard-deviation units, the baseline pooled 2SLS coefficients imply that a 1-SD increase in cost uncertainty raises prices by ≈ 5.8% while a 1-SD increase in demand uncertainty lowers prices by ≈ 6.6%. We read the baseline in column (3) as joint evidence that both channels are operating: cost uncertainty raises prices and demand uncertainty lowers them, with the demand channel quantitatively larger in magnitude (standardized units).

Table 2: Main Reduced-Form Results: Uncertainty and Prices

	(1) Cost only (OLS)	(2) Demand only (OLS)	(3) **Baseline** **(2SLS)**	(4) + Industry FE (2SLS)	(5) **FD baseline** **(2SLS)**	(6) OLS pooled	(7) OLS FD
Dependent variable: $\log p_{it}$ *(cols 1–4, 6),* $\Delta \log p_{it}$ *(cols 5, 7)*							
Cost uncertainty (γ_{cost})	0.0548***		0.0578***	0.0338***	0.0018	0.0617***	−0.0009
	(0.0071)		(0.0059)	(0.0056)	(0.0067)	(0.0075)	(0.0054)
Implied 1-SD $\Delta \log p$	[5.495%]		[5.795%]	[3.092%]	[0.125%]	[6.187%]	[−0.061%]
Demand uncertainty (γ_{demand})		−0.0443***	−0.0613***	−0.0300***	−0.0025	−0.0521***	−0.0028
		(0.0110)	(0.0121)	(0.0099)	(0.0054)	(0.0110)	(0.0036)
Implied 1-SD $\Delta \log p$		[−4.788%]	[−6.618%]	[−2.705%]	[−0.208%]	[−5.623%]	[−0.234%]
Cost level (β_C)	0.6459***	0.5538***	0.5816***	1.2290***	0.7593***	0.6166***	0.2127***
	(0.0435)	(0.0502)	(0.0716)	(0.0637)	(0.1725)	(0.0452)	(0.0731)
Demand level (β_D)	−0.0028	0.0050	0.0129	0.0206	0.0261	0.0047	0.0146***
	(0.0069)	(0.0070)	(0.0096)	(0.0270)	(0.0687)	(0.0067)	(0.0047)
Industry FE	No	No	No	Yes	(differenced)	No	(differenced)
Year FE	No	No	No	No	No	No	No
Observations	12,795	12,795	12,649	12,649	12,322	12,795	12,460
R^2	0.129	0.126				0.140	0.041
Estimator	OLS	OLS	2SLS	2SLS	2SLS	OLS	OLS
Predicted sign (γ)	(+)	(−)					

Notes: Column (3) is the baseline: pooled 2SLS (no industry FE) with Bartik level controls C, D instrumented by their 1972-benchmark-weight analogs; the uncertainty regressors are predetermined by construction (AR(1) forecast errors, one-period lagged) and enter as their own instruments. Pooled identification exploits both cross-industry differences in persistent risk exposure and within-industry time variation in σ^2; the 1972-weight IVs absorb confounding from time-invariant industry traits. Stars on γ rows denote *one-sided* significance in the direction predicted by the model ($\gamma_{\text{cost}} > 0$, $\gamma_{\text{demand}} < 0$): * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Stars on β rows (level controls) denote *two-sided* significance at the conventional thresholds.

6 Identification and Calibration

This section will report the structural calibration of the dynamic GE model once the empirical-side identification is settled. The CARA + quasilinear + linear-MO structure delivers a closed-form mapping from the reduced-form coefficients $(\hat{\gamma}_{\text{cost}}, \hat{\gamma}_{\text{demand}})$ to the structural primitives (α, a, b) via Corollary 3, providing one identification route; additional moment conditions and a full simulated method of moments exercise are work in progress. The calibrated parameter values, the discussion of empirical fit, and any auxiliary moment targets will appear in a subsequent draft.

7 Quantitative Results

The quantitative exercises based on the calibrated dynamic GE model — mirror impulse responses to cost- and demand-uncertainty shocks, the welfare cost of uncertainty (consumption-equivalent variation), the historical decomposition of markup cyclicality (1976–2018), sensitivity analysis, and the insurance-ladder comparison — are work in progress and will appear in a subsequent draft.

8 Conclusion

This paper develops a theory of pricing under risk that provides a unified framework for understanding markup behavior over the business cycle. By modeling firms that set prices ex ante under uncertainty and are owned by risk-averse agents, the analysis shows how markups arise endogenously as part of firm's precautionary pricing responses to uncertainty in costs and demand. A central implication is that markup cyclicality is not a primitive feature of the economy, but an equilibrium outcome determined by the composition of risks firms face.

The theory delivers sharp and intuitive predictions: different sources of uncertainty distort pricing incentives in systematically different—and sometimes opposite—ways, even when overall uncertainty is similar. As a result, economies with different risk structures can exhibit markedly different markup dynamics. Empirical evidence from U.S. manufacturing supports these predictions, showing that prices respond differently to cost and demand uncertainty in a manner consistent with the theory.

Beyond explaining observed markup dynamics, the analysis highlights a pricing channel through which uncertainty and risk exposure shape the transmission of aggregate shocks. Understanding how firms price under risk is therefore essential for interpreting markup movements and their implications for macroeconomic dynamics.

References

Angeletos, G.-M. and La'O, J. (2013). Sentiments, *Econometrica* **81**(2): 739–779.

Atkeson, A. and Burstein, A. (2008). Pricing-to-market, trade costs, and international relative prices, *American Economic Review* **98**(5): 1998–2031.

Bartelsman, E. J. and Gray, W. B. (1996). The NBER manufacturing productivity database, *Technical Working Paper 205*, National Bureau of Economic Research.

Basu, S. and Fernald, J. G. (1997). Returns to scale in us production: Estimates and implications, *Journal of political economy* **105**(2): 249–283.

Bils, M. (1987). The cyclical behavior of marginal cost and price, *The American Economic Review* pp. 838–855.

Blanchard, O. J. (1985). Debt, deficits, and finite horizons, *Journal of Political Economy* **93**(2): 223–247.

Blanchard, O. J. and Kiyotaki, N. (1987). Monopolistic competition and the effects of aggregate demand, *The american economic review* pp. 647–666.

Boar, C., Gorea, D. and Midrigan, V. (2025). The aggregate costs of uninsurable business risk, *Technical report*, working paper.

Borusyak, K., Hull, P. and Jaravel, X. (2022). Quasi-experimental shift-share research designs, *The Review of economic studies* **89**(1): 181–213.

Burstein, A. T., Carvalho, V. M. and Grassi, B. (2025). Bottom-up markup fluctuations, *The Quarterly Journal of Economics* **140**(4): 2619–2684.

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

Chamberlin, E. H. (1933). *The Theory of Monopolistic Competition*, Harvard University Press, Cambridge, MA.

De Loecker, J., Eeckhout, J. and Unger, G. (2020). The rise of market power and the macroeconomic implications, *The Quarterly Journal of Economics* **135**(2): 561–644.

Di Tella, S., Malgieri, C. and Tonetti, C. (2025). Risk markups, *Technical report*, National Bureau of Economic Research.

Dixit, A. K. and Stiglitz, J. E. (1977). Monopolistic competition and optimum product diversity, *The American economic review* **67**(3): 297–308.

Edmond, C., Midrigan, V. and Xu, D. Y. (2023). How costly are markups?, *Journal of Political Economy* **131**(7): 1619–1675.

Flynn, J. P., Nikolakoudis, G. and Sastry, K. A. (2024). Prices vs. quantities: A macroeconomic analysis.

Fort, T. C. and Klimek, S. D. (2018). The effects of industry classification changes on us employment composition, *CES Working Paper 18-28*, Center for Economic Studies, US Census Bureau.

Goldsmith-Pinkham, P., Sorkin, I. and Swift, H. (2020). Bartik instruments: What, when, why, and how, *American Economic Review* **110**(8): 2586–2624.

Golosov, M. and Lucas, R. E. (2007). Menu costs and phillips curves, *Journal of Political Economy* **115**(2): 171–199.

Hall, R. E. (1988). The relation between price and marginal cost in us industry, *Journal of political Economy* **96**(5): 921–947.

Hasenzagl, T. and Pérez, L. (2023). The micro-aggregated profit share, *arXiv preprint arXiv:2309.12945* .

Jurado, K., Ludvigson, S. C. and Ng, S. (2015). Measuring uncertainty, *American Economic Review* **105**(3): 1177–1216.

Knight, F. H. (1921). *Risk, uncertainty and profit*, Vol. 31, Houghton Mifflin.

Lucas, R. E. (1978). On the size distribution of business firms, *Bell Journal of Economics* **9**(2): 508–523.

Melitz, M. J. and Polanec, S. (2015). Dynamic olley-pakes productivity decomposition with entry and exit, *RAND Journal of Economics* **46**(2): 362–375.

Menzio, G. (2024). Markups: A search-theoretic perspective, *Technical report*, National Bureau of Economic Research.

Nekarda, C. J. and Ramey, V. A. (2020). The cyclical behavior of the price-cost markup, *Journal of Money, Credit and Banking* **52**(S2): 319–353.

Yaari, M. E. (1965). Uncertain lifetime, life insurance, and the theory of the consumer, *The Review of Economic Studies* **32**(2): 137–150.

Appendices

"Precautionary Pricing and Markup Cyclicality"

Thomas Hasenzagl Luis Pérez

April 28, 2026

List of Appendices

A Microfoundations for Price Commitment

The baseline model assumes that firms commit to prices before observing demand and cost shocks. This appendix provides a microfoundation for both price commitment and demand uncertainty based on directed search. Cost uncertainty—arising from input price risk—is standard in pricing models and is therefore taken as primitive.

A.1 Environment

Time is discrete. In each period, a unit mass of buyers and a unit mass of sellers participate in a market for a homogeneous good. Each buyer demands at most one unit and has valuation v, drawn i.i.d. from a distribution F with support $[\underline{v}, \overline{v}]$ and continuous density f. Buyers observe their valuations prior to search.

Each seller produces output using one unit of an intermediate input per unit of output. The input is purchased in a competitive spot market at price

$$\bar{\kappa} + \eta,$$

where $\bar{\kappa} > 0$ and η is a mean-zero cost shock realized after output prices are posted. This captures environments in which firms commit to prices before upstream markets clear (e.g., catalog pricing prior to the realization of energy or commodity prices).

Sellers are risk-averse entrepreneurs with CARA preferences

$$U(\pi) = -\exp(-\alpha\pi),$$

where $\alpha > 0$.

The timing within period is as follows:

1. Sellers simultaneously post prices.
2. Buyers observe all posted prices and direct their search toward sellers.
3. Each buyer visits exactly one seller.
4. Input prices $\bar{\kappa} + \eta$ are realized.
5. Trade occurs at posted prices; sellers produce and deliver.

Thus, prices are set before sellers know either how many buyers will arrive or the realization of input costs.

A.2 Buyer Search and Demand

Consider a seller posting price p. In a symmetric directed-search equilibrium, buyers randomize across sellers posting the same price. Let $\lambda(p)$ denote the expected number of buyers visiting a seller that posts price p. Standard directed search arguments imply that $\lambda(p)$ is decreasing in p, reflecting the fact that higher prices attract fewer buyers. We assume

$$\lambda'(p) < 0.$$

Given a unit mass of buyers and independent search decisions, the realized number of arrivals n at a seller posting price p follows a Poisson distribution:

$$n \sim \text{Poisson}(\lambda(p)).$$

Each arriving buyer purchases one unit, so realized output equals arrivals:

$$y = n.$$

A.3 Seller Profits and Risk Decomposition

A seller posting price p earns profits

$$\pi(p; n, \eta) = (p - \bar{\kappa} - \eta)n.$$

Taking expectations over arrivals and cost shocks (which are independent),

$$\mathbb{E}[\pi(p)] = (p - \bar{\kappa})\lambda(p).$$

Writing realized demand as the sum of its mean and a deviation:

$$n = \lambda(p) + \xi, \qquad \xi := n - \lambda(p),$$

where $\mathbb{E}[\xi] = 0$ and, under the Poisson specification,

$$\text{Var}(\xi) = \lambda(p).$$

Substituting yields

$$\pi(p;\xi,\eta) = (p-\bar{\kappa}-\eta)[\lambda(p)+\xi]$$
$$= \underbrace{(p-\bar{\kappa})\lambda(p)}_{\text{deterministic}} + \underbrace{(p-\bar{\kappa})\xi - \lambda(p)\eta - \eta\xi}_{\text{stochastic}}.$$

This expression coincides with the profit representation in Section 2, with expected demand $y(p) = \lambda(p)$. Demand uncertainty arises endogenously from buyer arrival risk, while cost uncertainty originates from spot-market input prices.

If expected arrivals are linear in price,

$$\lambda(p) = a - bp, \qquad a, b > 0,$$

then profits reduce to

$$\pi(p;\xi,\eta) = (p-\bar{\kappa})(a-bp+\xi) - (a-bp+\xi)\eta,$$

which matches exactly the profit function analyzed in Section 2 without demand-elasticity risk.

A.4 Optimal Pricing

The seller chooses p to maximize expected utility prior to observing ξ or η:

$$\max_{p>0} \mathbb{E}_{\xi,\eta}\left[-\exp\left(-\alpha\pi(p;\xi,\eta)\right)\right].$$

This is precisely the pricing problem studied in Section 2. All those results carry over: cost uncertainty raises optimal prices, demand uncertainty lowers them, and the cyclicality of markups depends on the relative importance of these two sources of risk.

Under the Poisson specification, demand variance satisfies $\sigma_\xi^2 = \lambda(p)$, so demand risk is increasing in expected sales. Evaluated at the risk-neutral optimum,

$$\sigma_\xi^2 = \lambda(p_{\text{RN}}) = a - bp_{\text{RN}} = \frac{a-b\bar{\kappa}}{2}.$$

Time variation in demand uncertainty can therefore arise from fluctuations in the mass of active buyers or in search intensity, while time variation in cost uncertainty reflects volatility in input markets.

B Omitted Proofs

B.1 Proof of Theorem 1

Theorem 1 (Optimal pricing under cost and demand uncertainty). *Suppose that:*

A1. $U : \mathbb{R} \to \mathbb{R}$ *is twice continuously differentiable, strictly increasing, and strictly concave.*

A2. *For almost all* $\boldsymbol{\theta}$, $\mathcal{D}(\cdot;\boldsymbol{\theta})$ *is twice continuously differentiable with* $\mathcal{D}_p(p;\boldsymbol{\theta}) < 0$ *for all* $p > 0$, *and* $\mathcal{C}(\cdot;\boldsymbol{\theta})$ *is twice continuously differentiable with* $\mathcal{C}_y(y;\boldsymbol{\theta}) > 0$ *for all* $y > 0$.

A3. *For almost all* $\boldsymbol{\theta}$, *the profit function* $p \mapsto \pi(p;\boldsymbol{\theta})$ *is strictly concave on* $(0,\infty)$.

A4. *There exists an open interval* $I \subset (0,\infty)$ *and an integrable function* $M(\boldsymbol{\theta})$ *such that, for all* $p \in I$,

$$\left|U\Big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\Big)\right| \le M(\boldsymbol{\theta}) \qquad \text{and} \qquad \left|U'\Big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\Big)\,\pi_p(p;\boldsymbol{\theta})\right| \le M(\boldsymbol{\theta}).$$

A5. *There exists* $0 < \underline{p} < \overline{p} < \infty$ *with* $[\underline{p},\overline{p}] \subset I$ *such that*

$$\mathbb{E}_{\boldsymbol{\theta}}\Big[U'\Big(\omega(\boldsymbol{\theta}) + \pi(\underline{p};\boldsymbol{\theta})\Big)\,\pi_p(\underline{p};\boldsymbol{\theta})\Big] > 0, \qquad \mathbb{E}_{\boldsymbol{\theta}}\Big[U'\Big(\omega(\boldsymbol{\theta}) + \pi(\overline{p};\boldsymbol{\theta})\Big)\,\pi_p(\overline{p};\boldsymbol{\theta})\Big] < 0.$$

Then there exists a unique optimal price $p^* \in (\underline{p},\overline{p})$ *characterized by*

$$\mathbb{E}_{\boldsymbol{\theta}}\Big[U'\Big(\omega(\boldsymbol{\theta}) + \pi(p^*;\boldsymbol{\theta})\Big)\,\pi_p(p^*;\boldsymbol{\theta})\Big] = 0,$$

where

$$\pi_p(p;\boldsymbol{\theta}) = y(p;\boldsymbol{\theta}) + \Big[p - \mathcal{MC}(p;\boldsymbol{\theta})\Big]y_p(p;\boldsymbol{\theta}), \qquad \mathcal{MC}(p;\boldsymbol{\theta}) := \mathcal{C}_y\Big(\mathcal{D}(p;\boldsymbol{\theta});\boldsymbol{\theta}\Big).$$

Proof. Define the objective function $F(p) := \mathbb{E}_{\boldsymbol{\theta}}[U(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta}))]$. By A3, profits are strictly concave in price almost surely with respect to $\boldsymbol{\theta}$. By A1, utility is strictly increasing and strictly concave, implying that $p \mapsto U(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta}))$ is strictly concave almost surely. Hence, F is strictly concave. By A4 and the dominated convergence theorem, differentiation under the expectation is valid, and for all $p \in I$,

$$F'(p) = \mathbb{E}_{\boldsymbol{\theta}}\Big[U'\Big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\Big)\pi_p(p;\boldsymbol{\theta})\Big].$$

By A5, $F'(\underline{p}) > 0$ and $F'(\overline{p}) < 0$ for some $0 < \underline{p} < \overline{p}$ with $[\underline{p},\overline{p}] \subset I$. Continuity of F' on $[\underline{p},\overline{p}]$ implies the existence of $p^* \in (\underline{p},\overline{p})$ such that $F'(p^*) = 0$. Strict concavity of F

guarantees uniqueness. Finally, by A2 and the chain rule,

$$\pi_p(p;\boldsymbol{\theta}) = \mathcal{D}(p;\boldsymbol{\theta}) + \Big[p - \mathcal{C}_y\big(\mathcal{D}(p;\boldsymbol{\theta});\boldsymbol{\theta}\big)\Big]\mathcal{D}_p(p;\boldsymbol{\theta}) = y(p;\boldsymbol{\theta}) + \Big[p - \mathcal{MC}(p;\boldsymbol{\theta})\Big]y_p(p;\boldsymbol{\theta}),$$

where $\mathcal{MC}(p;\boldsymbol{\theta}) := \mathcal{C}_y\big(\mathcal{D}(p;\boldsymbol{\theta});\boldsymbol{\theta}\big)$. □

B.2 Proof of Theorem 2

Theorem 2 (Local pricing formula around the risk-neutral price). *Suppose the assumptions of Theorem 1 hold. Let the risk-neutral price p_{RN} be the unique interior maximizer of expected profits, characterized by*

$$\mathbb{E}_{\boldsymbol{\theta}}[\pi_p(p_{\mathrm{RN}};\boldsymbol{\theta})] = 0, \qquad \mathbb{E}_{\boldsymbol{\theta}}[\pi_{pp}(p_{\mathrm{RN}};\boldsymbol{\theta})] < 0.$$

In addition, assume:

A6. *Profits have finite fourth moments uniformly in a neighborhood of p_{RN}:*

$$\sup_{p\in\mathcal{N}(p_{\mathrm{RN}})} \mathbb{E}_{\boldsymbol{\theta}}\Big[|\pi(p;\boldsymbol{\theta})|^4\Big] < \infty.$$

A7. *The mapping $p \mapsto \mathbb{E}_{\boldsymbol{\theta}}[\pi(p;\boldsymbol{\theta})]$ is twice continuously differentiable on $\mathcal{N}(p_{RN})$.*

A8. *Preferences belong to a one-parameter family $\{U_\alpha\}_{\alpha\geq 0}$ with $U_0(\mathcal{W}) = \mathcal{W}$ and absolute risk aversion*

$$A_\alpha(\mathcal{W}) := -\frac{U''(\mathcal{W})}{U'(\mathcal{W})} = \alpha A(\mathcal{W}),$$

where $A : \mathbb{R} \to \mathbb{R}_+$ is continuously differentiable.

Define benchmark objects evaluated at p_{RN} as

$$\mathcal{W}_{\mathrm{RN}}(\boldsymbol{\theta}) := \omega(\boldsymbol{\theta}) + \pi(p_{\mathrm{RN}};\boldsymbol{\theta}), \qquad \pi_{\mathrm{RN}}(\boldsymbol{\theta}) := \pi(p_{\mathrm{RN}};\boldsymbol{\theta}),$$
$$\pi_{p,\mathrm{RN}}(\boldsymbol{\theta}) := \pi_p(p_{\mathrm{RN}};\boldsymbol{\theta}), \qquad y_{p,\mathrm{RN}}(\boldsymbol{\theta}) := y_p(p_{\mathrm{RN}};\boldsymbol{\theta}).$$

Then, for sufficiently small α, the optimal price $p(\alpha)$ under preferences U_α satisfies

$$p(\alpha) = p_{\mathrm{RN}} + \alpha\mathcal{R} + \mathcal{O}(\alpha^2) + \alpha \cdot o_{\|\Sigma\|\to 0}(1),$$

where the first-order risk adjustment admits the covariance representation

$$\mathcal{R} = \varphi\Bigg[\underbrace{\frac{\overline{A}\cdot \mathrm{Cov}_{\boldsymbol{\theta}}\big(\pi_{\mathrm{RN}},\pi_{p,\mathrm{RN}}\big)}{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,\mathrm{RN}}]}}_{\text{pure risk term}} + \underbrace{\frac{\mathrm{Cov}_{\boldsymbol{\theta}}\big(A(\mathcal{W}_{\mathrm{RN}}),\pi_{\mathrm{RN}}\pi_{p,\mathrm{RN}}\big)}{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,\mathrm{RN}}]}}_{\text{wealth-effects term}}\Bigg],$$

with $\overline{A} := \mathbb{E}_{\boldsymbol{\theta}}[A(\mathcal{W}_{RN})]$.

Proof. For $\alpha \geq 0$, define the expected-utility objective

$$V(p,\alpha) := \mathbb{E}_{\boldsymbol{\theta}}\Big[U_\alpha\big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\big)\Big],$$

and its first-order condition

$$F(p,\alpha) = \mathbb{E}_{\boldsymbol{\theta}}\Big[U'_\alpha\big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\big)\pi_p(p;\boldsymbol{\theta})\Big].$$

By Theorem 1 applied to U_α and assumptions A6–A8, the optimal price $p(\alpha)$ is the unique interior solution to $F(p,\alpha) = 0$. At $\alpha = 0$, we have $U_0(\mathcal{W}) = \mathcal{W}$ and hence $U'_0(\mathcal{W}) \equiv 1$. Therefore, $F(p,0) = \mathbb{E}_{\boldsymbol{\theta}}[\pi_p(p;\boldsymbol{\theta})]$, and the risk-neutral price p_{RN} satisfies

$$F(p_{\mathrm{RN}},0) = \mathbb{E}_{\boldsymbol{\theta}}[\pi_p(p;\boldsymbol{\theta})] = 0.$$

By assumptions A6–A8 and dominated convergence, differentiation under the expectation is valid in a neighborhood of $(p_{\mathrm{RN}},0)$. Hence, $F_p(p,0) = \mathbb{E}_{\boldsymbol{\theta}}[\pi_{pp}(p;\boldsymbol{\theta})]$ so that

$$F_p(p_{\mathrm{RN}},0) = \mathbb{E}_{\boldsymbol{\theta}}[\pi_{pp}(p_{\mathrm{RN}};\boldsymbol{\theta})] < 0.$$

To compute $F_\alpha(p_{\mathrm{RN}},0)$, note that

$$F_\alpha(p,\alpha) = \mathbb{E}_{\boldsymbol{\theta}}\Big[\partial_\alpha U'_\alpha\big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\big)\pi_p(p;\boldsymbol{\theta})\Big].$$

Under assumption A8, absolute risk aversion satisfies

$$-\frac{U''_\alpha(\mathcal{W})}{U'_\alpha(\mathcal{W})} = \alpha A(\mathcal{W}).$$

Integrating this ODE and normalizing $U'_\alpha(\mathcal{W}_0) = 1$ at the reference point $\mathcal{W}_0 = \omega(\boldsymbol{\theta})$ gives the exact expression

$$U'_\alpha(\mathcal{W}) = \exp\left(-\alpha\int_{\mathcal{W}_0}^{\mathcal{W}} A(s)\,ds\right) = 1 - \alpha\int_{\mathcal{W}_0}^{\mathcal{W}} A(s)\,ds + \mathcal{O}(\alpha^2),$$

uniformly in $\mathcal{W}$ on any bounded interval (since A is continuous by A8). Evaluating at $\mathcal{W} = \omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})$,

$$\partial_\alpha U'_\alpha\Big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\Big)\Big|_{\alpha=0} = -\int_{\omega(\boldsymbol{\theta})}^{\omega(\boldsymbol{\theta})+\pi(p;\boldsymbol{\theta})} A(s)\,ds.$$

By the mean-value theorem for integrals there exists $\xi(\boldsymbol{\theta})$ between $\omega(\boldsymbol{\theta})$ and $\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})$ such that the above equals $-A(\xi(\boldsymbol{\theta}))\,\pi(p;\boldsymbol{\theta})$. Because A is continuously differentiable (A8), $A(\xi(\boldsymbol{\theta})) = A(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})) + \mathcal{O}(\pi(p;\boldsymbol{\theta}))$ along any sample path, so

$$\partial_\alpha U'_\alpha\Big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\Big)\Big|_{\alpha=0} = -A\Big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\Big)\pi(p;\boldsymbol{\theta}) + r(\boldsymbol{\theta}),$$

where the residual $r(\boldsymbol{\theta})$ is of order $\pi(p;\boldsymbol{\theta})^2$ and vanishes identically in the CARA case $A \equiv$ const. Multiplying by $\pi_p(p;\boldsymbol{\theta})$ and taking expectations, the residual contribution enters the implicit-function-theorem calculation below at order $\alpha \cdot \mathcal{O}(\|\boldsymbol{\theta}\|^2)$, which is subsumed in the $\mathcal{O}(\alpha^2) + o(\|\boldsymbol{\theta}\|^2)$ remainder of the statement. For the CARA specialisation of Section 3.4 (where $A \equiv \alpha_{\text{CARA}}$ is constant) the displayed formula below is exact in the first-order-in-α expansion; under variable-A preferences (CRRA and the like) it is the leading-order approximation in both α and the shock scale. Therefore,

$$F_\alpha(p_{\text{RN}}, 0) = -\mathbb{E}_{\boldsymbol{\theta}}[A(\mathcal{W}_{\text{RN}})\pi_{\text{RN}}\pi_{p,\text{RN}}],$$

where $\mathcal{W}_{\text{RN}}(\boldsymbol{\theta}) = \omega(\boldsymbol{\theta}) + \pi(p_{\text{RN}};\boldsymbol{\theta})$, $\pi_{\text{RN}}(\boldsymbol{\theta}) = \pi(p_{\text{RN}};\boldsymbol{\theta})$, and $\pi_{p,\text{RN}}(\boldsymbol{\theta}) = \pi_p(p_{\text{RN}};\boldsymbol{\theta})$.

A first-order Taylor expansion of F around $(p_{\text{RN}}, 0)$ yields

$$0 = F(p(\alpha), \alpha) = F(p_{\text{RN}}, 0) + F_p(p_{\text{RN}}, 0)\,[p(\alpha) - p_{\text{RN}}] + F_\alpha(p_{\text{RN}}, 0)\,\alpha + \mathcal{O}\Big(\alpha^2\Big).$$

Since $F(p_{\text{RN}}, 0) = 0$, rearranging gives

$$F_p(p_{\text{RN}}, 0)\,[p(\alpha) - p_{\text{RN}}] = -\alpha F_\alpha(p_{\text{RN}}, 0) + \mathcal{O}(\alpha^2).$$

Assumption A7 implies $F_p(p_{\text{RN}}, 0) \neq 0$. Dividing both sides by $F_p(p_{\text{RN}}, 0)$ gives

$$p(\alpha) = p_{\text{RN}} - \alpha \frac{F_\alpha(p_{\text{RN}}, 0)}{F_p(p_{\text{RN}}, 0)} + \mathcal{O}(\alpha^2).$$

Finally note that

$$
\begin{aligned}
-F_\alpha(p_{\mathrm{RN}},0) &= \mathbb{E}_{\boldsymbol{\theta}}[A(\mathcal{W}_{\mathrm{RN}})\pi_{\mathrm{RN}}\pi_{p,\mathrm{RN}}] \\
&= \mathrm{Cov}_{\boldsymbol{\theta}}\Big(A(\mathcal{W}_{\mathrm{RN}}),\pi_{\mathrm{RN}}\pi_{p,\mathrm{RN}}\Big) + \mathbb{E}_{\boldsymbol{\theta}}[A(\mathcal{W}_{\mathrm{RN}})]\mathbb{E}_{\boldsymbol{\theta}}[\pi_{\mathrm{RN}}\pi_{p,\mathrm{RN}}] \\
&= \mathrm{Cov}_{\boldsymbol{\theta}}\Big(A(\mathcal{W}_{\mathrm{RN}}),\pi_{\mathrm{RN}}\pi_{p,\mathrm{RN}}\Big) + \overline{A}\Big[\mathrm{Cov}_{\boldsymbol{\theta}}(\pi_{\mathrm{RN}},\pi_{p,\mathrm{RN}}) + \mathbb{E}_{\boldsymbol{\theta}}[\pi_{\mathrm{RN}}]E_{\boldsymbol{\theta}}[\pi_{p,\mathrm{RN}}\Big] \\
&= \mathrm{Cov}_{\boldsymbol{\theta}}\Big(A(\mathcal{W}_{\mathrm{RN}}),\pi_{\mathrm{RN}}\pi_{p,\mathrm{RN}}\Big) + \overline{A}\mathrm{Cov}_{\boldsymbol{\theta}}(\pi_{\mathrm{RN}},\pi_{p,\mathrm{RN}})
\end{aligned}
$$

since $\mathbb{E}_{\boldsymbol{\theta}}[\pi_{p,\mathrm{RN}}] = 0$.

Substituting yields the desired covariance representation. Finally, express the local price adjustment in terms of pass-through. By definition of cost pass-through and the risk-neutral first-order conditon, $\mathbb{E}_{\boldsymbol{\theta}}[\pi_p(p_{\mathrm{RN}}(\overline{\kappa});\boldsymbol{\theta},\overline{\kappa})] = 0$. The implicit function theorem gives

$$
\varphi := \frac{\partial p_{\mathrm{RN}}}{\partial \overline{\kappa}} = -\frac{\mathbb{E}_{\boldsymbol{\theta}}[\pi_{p\overline{\kappa},\mathrm{RN}}]}{\mathbb{E}_{\boldsymbol{\theta}}[\pi_{pp,\mathrm{RN}}]} = \frac{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,\mathrm{RN}}]}{\mathbb{E}_{\boldsymbol{\theta}}[\pi_{pp,\mathrm{RN}}]},
$$

where the last equality follows from how the cost shifter enters profits. It follows that

$$
F_p(p_{\mathrm{RN}},0) = \mathbb{E}_{\boldsymbol{\theta}}[\pi_{pp,\mathrm{RN}}] = \frac{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,\mathrm{RN}}]}{\varphi},
$$

which gives the desired result. □

B.3 Proof of Proposition 1

Proposition 1 (Decomposition of the pure risk adjustment term by source). *Let $\boldsymbol{\theta} = (\eta,\xi,\psi)$ collect a marginal-cost shifter η, a demand-level shifter ξ, and a demand-elasticity shifter ψ. Suppose the assumptions of Theorem 2 hold. In addition, assume:*

A9. *Shocks have mean zero, are mutually independent, and have finite fourth moments.*

A10. *For each p in a neighborhood of p_{RN}, the mapping $\boldsymbol{\theta} \mapsto \pi(p;\boldsymbol{\theta})$ is three times continuously differentiable in a neighborhood of $\boldsymbol{\theta} = \mathbf{0}$.*

Let $\sigma_k^2 = Var(\theta_k)$ and $\Sigma = Var(\boldsymbol{\theta})$. Let $\mathcal{R}$ denote the risk adjustment from Theorem 2, and write

$$
\mathcal{R} = \mathcal{R}_{\mathrm{risk}} + \mathcal{R}_{\mathrm{wealth}},
$$

where $\mathcal{R}_{risk}$ is the pure-risk component.

For $k,l \in \{\eta,\xi,\psi\}$, define the following exposures, evaluated at $(p,\boldsymbol{\theta}) = (p_{RN},\mathbf{0})$:

$$\pi_k := \partial_k \pi, \qquad \pi_{pk} := \partial_k \pi_p, \qquad \pi_{kl} := \partial_{kl}\pi, \qquad \pi_{p,kl} := \partial_{kl}\pi_p.$$

Then the pure-risk component admits the first-order expansion

$$\mathcal{R}_{\text{risk}} = \underbrace{\mathcal{R}^{(\eta)}(\sigma_\eta^2)}_{\text{cost risk}} + \underbrace{\mathcal{R}^{(\xi)}(\sigma_\xi^2)}_{\text{demand-level risk}} + \underbrace{\mathcal{R}^{(\psi)}(\sigma_\psi^2)}_{\text{demand-elasticity risk}} + o\left(\|\Sigma\|\right),$$

where

$$\mathcal{R}^{(k)}(\sigma_k^2) := \varphi \overline{A} \cdot \frac{\pi_k \pi_{pk}}{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,\text{RN}}]} \sigma_k^2.$$

At the next order in Σ the covariance picks up cross-risk interaction contributions

$$\mathcal{R}^{(kl)}(\sigma_k^2\sigma_l^2) := \varphi \frac{\overline{A}}{2} \frac{\pi_{kl}\pi_{p,kl}}{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,\text{RN}}]} \sigma_k^2 \sigma_l^2, \quad k \neq l,$$

with $\overline{A} := \mathbb{E}_{\boldsymbol{\theta}}[A(\mathcal{W}_{RN})]$ and $\|\cdot\|$ any matrix norm.

Proof. Fix $p = p_{\text{RN}}$. By assumption A10, the mappings $\boldsymbol{\theta} \mapsto \pi(p_{\text{RN}},\boldsymbol{\theta})$ and $\boldsymbol{\theta} \mapsto \pi_p(p_{\text{RN}},\boldsymbol{\theta})$ are three times continuously differentiable in a neighborhood of $\boldsymbol{\theta} = \mathbf{0}$. A second-order Taylor expansion yields

$$\pi(p_{\text{RN}};\boldsymbol{\theta}) = \pi(p_{\text{RN}};\mathbf{0}) + \sum_k \pi_k \theta_k + \frac{1}{2}\sum_{k,l} \pi_{kl}\theta_k\theta_l + o(\|\boldsymbol{\theta}\|^2),$$
$$\pi_p(p_{\text{RN}};\boldsymbol{\theta}) = \pi_p(p_{\text{RN}};\mathbf{0}) + \sum_k \pi_{pk}\theta_k + \frac{1}{2}\sum_{k,l} \pi_{p,kl}\theta_k\theta_l + o(\|\boldsymbol{\theta}\|^2),$$

where all derivatives are evaluated at $(p,\boldsymbol{\theta}) = (p_{\text{RN}},\mathbf{0})$ and $k,l \in \{\eta,\xi,\psi\}$.

From Theorem 2, the pure-risk component of the risk adjustment is

$$\mathcal{R}_{\text{risk}} = \varphi \frac{\overline{A} \cdot \text{Cov}_{\boldsymbol{\theta}}(\pi,\pi_p)}{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,\text{RN}}]}, \qquad \overline{A} := \mathbb{E}_{\boldsymbol{\theta}}[A(\mathcal{W}_{\text{RN}})].$$

By definition,

$$\text{Cov}_{\boldsymbol{\theta}}(\pi_{\text{RN}},\pi_{p,\text{RN}}) = \mathbb{E}\Big[\big(\pi_{\text{RN}} - \mathbb{E}[\pi_{\text{RN}}]\big)\big(\pi_{p,\text{RN}} - \mathbb{E}[\pi_{p,\text{RN}}]\big)\Big].$$

Under assumption A9 (mean zero, mutually independent shocks with finite fourth moments), the Taylor expansions imply

$$\mathbb{E}[\pi_{\text{RN}}] = \pi(p_{\text{RN}}; \mathbf{0}) + \frac{1}{2}\sum_k \pi_{kk}\mathbb{E}[\theta_k^2] + o(\|\Sigma\|),$$
$$\mathbb{E}[\pi_{p,\text{RN}}] = \pi_p(p_{\text{RN}}; \mathbf{0}) + \frac{1}{2}\sum_k \pi_{p,kk}\mathbb{E}[\theta_k^2] + o(\|\Sigma\|).$$

Hence,

$$\pi_{\text{RN}} - \mathbb{E}[\pi_{\text{RN}}] = \sum_k \pi_k\theta_k + \frac{1}{2}\sum_{k\neq l}\pi_{kl}\theta_k\theta_l + \frac{1}{2}\sum_k \pi_{kk}\left(\theta_k^2 - \mathbb{E}[\theta_k^2]\right) + o\left(\|\boldsymbol{\theta}\|^2\right)$$

and an analogous expression holds for $\pi_{p,\text{RN}} - \mathbb{E}[\pi_{p,\text{RN}}]$.

Expanding the product inside the covariance, the linear-linear term yields

$$\sum_k \pi_k\pi_{pk}\mathbb{E}[\theta_k^2] = \sum_k \pi_k\pi_{pk}\sigma_k^2.$$

All linear-quadratic terms are cubic in the shocks and therefore vanish by mean zero and independence. The off-diagonal quadratic-quadratic terms yield

$$\frac{1}{2}\sum_{k<l}\pi_{kl}\pi_{p,kl}\sigma_k^2\sigma_l^2.$$

All remaining contributions—arising from the centered diagonal quadratic terms and Taylor remainders—are of order at most $\|\Sigma\|^2$ and are absorbed into $o\left(\|\Sigma\|^2\right)$. Collecting terms,

$$\text{Cov}_{\boldsymbol{\theta}}(\pi_{\text{RN}}, \pi_{p,\text{RN}}) = \sum_k \pi_k\pi_{pk}\sigma_k^2 + \frac{1}{2}\sum_{k<l}\pi_{kl}\pi_{p,kl}\sigma_k^2\sigma_l^2 + o\left(\|\Sigma\|^2\right).$$

Multiplying by $\varphi\overline{A}/\mathbb{E}_{\boldsymbol{\theta}}[y_{p,\text{RN}}]$ yields the stated expression for $\mathcal{R}_{\text{risk}}$. □

B.4 Proof of Proposition 2

Proposition 2 (Sufficient conditions to sign each risk source in pure-risk component). *Let $\boldsymbol{\theta} = (\eta, \xi, \psi)$, where η is a marginal cost shifter, ξ is a demand-level shifter, and ψ is a demand-elasticity shifter. Suppose the assumptions of Proposition 1 hold, and that*

$$\mathbb{E}_{\boldsymbol{\theta}}\left[\pi_{pp}(p_{\text{RN}}; \boldsymbol{\theta})\right] < 0.$$

Define the second-order pure-risk contribution of source $k \in \{\eta, \xi, \psi\}$ to the price adjustment by

$$\mathcal{R}^{(k)}_{\text{risk}} := \mathcal{R}^{(k)}(\sigma_k^2) + \frac{1}{2}\sum_{l \neq k} \mathcal{R}^{(kl)}(\sigma_k^2 \sigma_l^2),$$

where $\mathcal{R}^{(k)}$ and $\mathcal{R}^{(kl)}$ are defined in Proposition 1.

Assume that, locally at $(p, \boldsymbol{\theta}) = (p_{RN}, \mathbf{0})$, the following exposure conditions hold:

$$\pi_\eta < 0, \quad \pi_{p\eta} > 0; \qquad \pi_\xi > 0, \quad \pi_{p\xi} \geq 0; \qquad \pi_\psi < 0, \quad \pi_{p\psi} < 0.$$

Assume further that interactions involving demand shifters do not overturn these signs, in the sense that

$$\pi_{\xi l}\pi_{p,\xi l} \geq 0 \ \text{ for all } \ l \neq \xi, \qquad \pi_{\psi l}\pi_{p,\psi l} \geq 0 \ \text{ for all } \ l \neq \psi.$$

Then, to second order in uncertainty:

1. *Cost-uncertainty raises prices provided the interaction contributions from demand shocks are not too large: there exist thresholds $\overline{\sigma}^2_\xi(\sigma^2_\psi), \overline{\sigma}^2_\psi(\sigma^2_\xi) > 0$ such that $\mathcal{R}^{(\eta)}_{risk} > 0$ whenever $\sigma^2_\xi \leq \overline{\sigma}^2_\xi(\sigma^2_\psi)$ and $\sigma^2_\psi \leq \overline{\sigma}^2_\psi(\sigma^2_\xi)$; equivalently, in the small-noise scaling $\Sigma = \varepsilon\Sigma_0$ the conclusion holds for all sufficiently small ε.*
2. *Demand-level uncertainty lowers prices: $\mathcal{R}^{(\xi)}_{risk} \leq 0$, with strict inequality if $\pi_{p\xi} > 0$.*
3. *Demand-elasticity uncertainty lowers prices: $\mathcal{R}^{(\psi)}_{risk} < 0$.*

Proof. By Proposition 1, the pure-risk component admits the leading-order expansion

$$\mathcal{R}_{\text{risk}} = \sum_{k \in \{\eta,\xi,\psi\}} \mathcal{R}^{(k)}(\sigma_k^2) + \sum_{k<l} \mathcal{R}^{(kl)}(\sigma_k^2\sigma_l^2) + o(\|\Sigma\|),$$

where the sourcewise contributions $\mathcal{R}^{(k)}(\sigma_k^2)$ are the first-order pure-risk channels of Proposition 1, and the cross-risk interaction terms $\mathcal{R}^{(kl)}(\sigma_k^2\sigma_l^2)$ are of order $O(\|\Sigma\|^2)$ (defined in the same proposition). Explicitly,

$$\mathcal{R}^{(k)}(\sigma_k^2) = \varphi\overline{A} \cdot \frac{\pi_k \pi_{pk}}{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,\text{RN}}]}\sigma_k^2, \qquad \mathcal{R}^{(kl)}(\sigma_k^2\sigma_l^2) = \varphi\frac{\overline{A}}{2}\frac{\pi_{kl}\pi_{p,kl}}{\mathbb{E}_{\boldsymbol{\theta}}[y_{p,\text{RN}}]}\sigma_k^2\sigma_l^2, \quad k \neq l,$$

with $\varphi, \overline{A} > 0$, and $\mathbb{E}_{\boldsymbol{\theta}}[y_{p,\text{RN}}] < 0$. All derivatives are evaluated at $(p, \boldsymbol{\theta}) = (p_{\text{RN}}, \mathbf{0})$. The $o(\|\Sigma\|)$ remainder absorbs all higher-order contributions, including the centered-quartic and third-derivative corrections discussed after Proposition 1. Under any small-noise scaling $\Sigma = \varepsilon\Sigma_0$ with $\varepsilon \to 0$, the $\mathcal{R}^{(k)}$ channels are $O(\varepsilon^2)$ and the $\mathcal{R}^{(kl)}$ interactions are $O(\varepsilon^4)$, so the latter do not overturn the signs established below in the limit; the "interactions not too large" condition in the statement is the finite-ε analogue of this asymptotic dominance.

By definition,

$$\mathcal{R}^{(k)}_{\text{risk}} = \frac{\varphi \overline{A}}{\mathbb{E}_\theta[y_{p,\text{RN}}]} \left[\pi_k \pi_{pk} \sigma_k^2 + \frac{1}{4} \sum_{l \neq k} \pi_{kl} \pi_{p,kl} \sigma_k^2 \sigma_l^2 \right].$$

Demand-level risk. For $k = \xi$, the assumed sign restrictions imply $\pi_\xi \pi_{p\xi} \geq 0$. By the interaction assumption, $\pi_{\xi l} \pi_{p,\xi l} \geq 0$ for all $l \neq \xi$, so every interaction contribution entering $\mathcal{R}^{(\xi)}_{\text{risk}}$ is weakly nonnegative. Hence the bracketed term is strictly positive for $\sigma_\xi^2 > 0$. Since $\varphi, \overline{A} > 0$ and $\mathbb{E}_\theta[y_{p,\text{RN}}] < 0$, $\mathcal{R}^{(\xi)}_{\text{risk}} \leq 0$. If $\pi_{p\xi} > 0$, the inequality is strict.

Demand-elasticity risk. For $k = \psi$, the assumed sign restrictions imply $\pi_\psi \pi_{p\psi} > 0$. The interaction assumption again implies all interaction terms entering $\mathcal{R}^{(\psi)}_{\text{risk}}$ are weakly nonnegative, so the bracketed term is strictly positive for $\sigma_\psi^2 > 0$. Since $\varphi, \overline{A} > 0$ and $\mathbb{E}_\theta[y_{p,\text{RN}}] < 0$, we obtain $\mathcal{R}^{(\psi)}_{\text{risk}} < 0$.

Cost risk. For $k = \eta$, the assumed sign restrictions imply $\pi_\eta \pi_{p\eta} < 0$, so the leading term in the bracket is strictly negative. Interaction terms take the form

$$\frac{1}{4} \sum_{l \neq \eta} \pi_{\eta l} \pi_{p,\eta l} \, \sigma_l^2,$$

which is $O(\|\Sigma\|)$ in the variances of the other shocks. Hence there exists $\overline{\sigma}_\psi^2(\sigma_\xi^2) > 0$ such that whenever $\sigma_\psi^2 \leq \overline{\sigma}_\psi^2$ the negative leading term dominates and the entire bracket remains negative. Since $\varphi, \overline{A} > 0$ and $\mathbb{E}_\theta[y_{p,\text{RN}}] < 0$, this implies $\mathcal{R}^{(\eta)}_{\text{risk}} > 0$ for all such σ_ψ^2. □

B.5 Proof of Proposition 3

Proposition 3 (Markup cyclicality). *Suppose the assumptions of Theorem 2 hold period by period for each firm i, and let the distribution of shocks vary over time through uncertainty parameters $\boldsymbol{\sigma}_{it}^2 = (\sigma_{\eta,it}^2, \sigma_{\xi,it}^2, \sigma_{\psi,it}^2)$. Then, for sufficiently small α,*

$$\mathrm{Cov}(\mu_{it}, y_{it}) = \mathrm{Cov}(\mu_{it,\mathrm{RN}}, y_{it}) + \frac{\alpha}{\overline{\kappa}_i} \sum_{k\in\{\eta,\xi,\psi\}} \underbrace{\frac{\partial \mathcal{R}_{it}}{\partial \sigma_{k,it}^2}}_{\substack{\text{markup sensitivity}\\ \text{to risk } k}} \times \underbrace{\mathrm{Cov}(\sigma_{k,it}^2, y_{it})}_{\text{cyclicality of risk } k} + \mathcal{O}(\alpha^2) + \mathcal{O}(\alpha\, \|\Delta \boldsymbol{\sigma}_{it}^2\|^2).$$

Proof. By Theorem 2, the optimal price in period t admits the expansion

$$p_{it}(\alpha) = p_{it,\mathrm{RN}} + \alpha \mathcal{R}_{it}(\boldsymbol{\sigma}_{it}^2) + \mathcal{O}(\alpha^2),$$

where $\mathcal{R}_{it}(\boldsymbol{\sigma}_{it}^2)$ depends smoothly on the vector of uncertainty parameters. Dividing by the firm-specific time-invariant deterministic marginal-cost scale $\overline{\kappa}_i > 0$ yields

$$\mu_{it}(\alpha) = \mu_{it,\mathrm{RN}} + \frac{\alpha}{\overline{\kappa}_i} \mathcal{R}_{it}(\boldsymbol{\sigma}_{it}^2) + \mathcal{O}(\alpha^2),$$

where $\mu_{it,\mathrm{RN}} := p_{it,\mathrm{RN}}/\overline{\kappa}_i$. Since $\overline{\kappa}_i$ is time-invariant per firm, it factors out of any time-series covariance $\mathrm{Cov}(\cdot, y_{it})$ as a scalar.

Taking the covariance with output y_{it} and using the bilinearity of covaraince,

$$\mathrm{Cov}(\mu_{it}, y_{it}) = \mathrm{Cov}(\mu_{it,\mathrm{RN}}, y_{it}) + \frac{\alpha}{\overline{\kappa}_{it}} \mathrm{Cov}\Big(\mathcal{R}_{it}(\boldsymbol{\sigma}_{it}^2), y_{it}\Big) + \mathcal{O}(\alpha^2).$$

Since $\mathcal{R}_{it}(\boldsymbol{\sigma}_{it}^2)$ is continuously differentiable in the uncertainty parameters, a first-order Taylor expansion around the mean uncertainty vector implies

$$\mathcal{R}_{it}(\boldsymbol{\sigma}_{it}^2) = \mathcal{R}(\overline{\boldsymbol{\sigma}}_i^2) + \sum_k \frac{\partial \mathcal{R}_{it}}{\partial \sigma_{k,it}^2}\bigg|_{\overline{\boldsymbol{\sigma}}_i^2} (\sigma_{k,it}^2 - \overline{\sigma}_{k,i}^2) + \mathcal{O}(\|\boldsymbol{\sigma}_{it}^2 - \overline{\boldsymbol{\sigma}}_i^2\|^2).$$

The constant term drops out of the covariance, yielding

$$\mathrm{Cov}(\mathcal{R}_{it}(\boldsymbol{\sigma}_t^2), y_{it}) = \sum_k \frac{\partial \mathcal{R}_{it}}{\partial \sigma_{k,it}^2}\bigg|_{\overline{\boldsymbol{\sigma}}_i^2} \mathrm{Cov}(\sigma_{k,it}^2, y_{it}) + \mathcal{O}(\|\boldsymbol{\sigma}_{it}^2 - \overline{\boldsymbol{\sigma}}_i^2\|^2).$$

Substituting delivers the desired expression. □

B.6 Proof of Proposition 4

Proposition 4 (Sectoral markups and sectoral markup cyclicality). *Suppose the assumptions of Theorem 2 hold period by period for each firm i. Then, for sufficiently small α,*

$$\mu_{st}^{\text{hsw}} = \mu_{st,\text{RN}}^{\text{hsw}} + \alpha\left(\mu_{st,\text{RN}}^{\text{hsw}}\right)^2 \sum_{i\in\mathcal{I}_{st}} \omega_{it} \frac{1}{\overline{\kappa}_{it}} \frac{\mathcal{R}_{it}(\sigma_{it}^2)}{\mu_{it,\text{RN}}^2} + \mathcal{O}(\alpha^2).$$

Moreover, the cyclicality of the sectoral markup with sectoral output y_{st} satisfies

$$\text{Cov}(\mu_{st}^{\text{hsw}}, y_{st}) = \text{Cov}(\mu_{st,\text{RN}}^{\text{hsw}}, y_{st}) + \alpha\left(\mu_{st,\text{RN}}^{\text{hsw}}\right)^2 \sum_{i\in\mathcal{I}_{st}} \frac{\omega_{it}}{\overline{\kappa}_{it}\mu_{it,\text{RN}}^2}\left[\sum_{k\in\{\eta,\xi,\psi\}} \frac{\partial\mathcal{R}_{it}}{\partial\sigma_{k,it}^2} \times \text{Cov}(\sigma_{k,it}^2, y_{st})\right].$$

Proof. By Proposition 3, firm-level markups are

$$\mu_{it}(\alpha) = \mu_{it,\text{RN}} + \frac{\alpha}{\overline{\kappa}_{it}}\mathcal{R}_{it}(\sigma_{it}^2) + \mathcal{O}(\alpha^2).$$

Fix a sector s and a period t. Using the definition of the sectoral markup,

$$\begin{aligned}
\frac{1}{\mu_{st}^{\text{hsw}}} &= \sum_{i\in\mathcal{I}_{st}} \omega_{it} \frac{1}{\mu_{it}} \\
&= \sum_{i\in\mathcal{I}_{st}} \omega_{it} \frac{1}{\mu_{it,\text{RN}} + \frac{\alpha}{\overline{\kappa}_{it}}\mathcal{R}_{it}(\sigma_{it}^2)} \\
&= \sum_{i\in\mathcal{I}_{st}} \omega_{it} \frac{1}{\mu_{it,\text{RN}}} - \alpha \sum_{i\in\mathcal{I}_{st}} \omega_{it} \frac{\mathcal{R}_{it}(\sigma_{it}^2)}{\overline{\kappa}_{it}\mu_{it,\text{RN}}^2} + \mathcal{O}(\alpha^2),
\end{aligned}$$

where the last line follows from using the first-order identity $(x+\varepsilon)^{-1} = x^{-1} - \varepsilon x^{-2}$, where $x = \mu_{it,\text{RN}}$ and $\varepsilon = \alpha\mathcal{R}_{it}$.

Doing a first-order inversion and using the definition of sectoral markup yields

$$\mu_{st}^{\text{hsw}} = \mu_{st,\text{RN}}^{\text{hsw}} + \alpha\left(\mu_{st,\text{RN}}^{\text{hsw}}\right)^2 \sum_{i\in\mathcal{I}_{st}} \omega_{it} \frac{\mathcal{R}_{it}(\sigma_{it}^2)}{\overline{\kappa}_{it}\mu_{it,\text{RN}}^2} + \mathcal{O}(\alpha^2).$$

We now derive the cyclicality expression. Under the price-commitment setup of Assumption A0, the weights ω_{it} are constructed from p_{it} and the pre-realization output $y_{it} = \mathcal{D}(p_{it};\boldsymbol{\theta}_{it})$ projected onto $\mathcal{F}_{it}^{\text{pre}}$; similarly $\overline{\kappa}_{it} = \mathbb{E}[\kappa_{it} \mid \mathcal{F}_{it}^{\text{pre}}]$ and $\mu_{it,\text{RN}}$ are $\mathcal{F}_{it}^{\text{pre}}$-measurable. Stack these pre-period objects into the weight vector $\boldsymbol{w}_{st} = (\omega_{it}/(\overline{\kappa}_{it}\mu_{it,\text{RN}}^2))_{i\in\mathcal{I}_{st}}$, which is $\mathcal{F}_{st}^{\text{pre}}$-measurable.

Linearize in the uncertainty parameters. Write $\boldsymbol{\sigma}_{it}^2 = \bar{\boldsymbol{\sigma}}_{it}^2 + \Delta\boldsymbol{\sigma}_{it}^2$, where $\bar{\boldsymbol{\sigma}}_{it}^2 = \mathbb{E}[\boldsymbol{\sigma}_{it}^2 \mid \mathcal{F}_{st}^{\text{pre}}]$ denotes the pre-period-conditional mean. A first-order Taylor expansion of $\mathcal{R}_{it}$

around $\bar{\sigma}_{it}^2$ gives

$$\mathcal{R}_{it}(\sigma_{it}^2) = \mathcal{R}_{it}(\bar{\sigma}_{it}^2) + \sum_{k\in\{\eta,\xi,\psi\}} \frac{\partial \mathcal{R}_{it}}{\partial \sigma_{k,it}^2} \cdot \Delta\sigma_{k,it}^2 + \mathcal{O}(\|\Delta\sigma_{it}^2\|^2),$$

where the partials are evaluated at $\bar{\sigma}_{it}^2$ and are themselves $\mathcal{F}_{st}^{\text{pre}}$-measurable. Substituting into the sectoral-markup expansion and taking covariance with y_{st},

$$\text{Cov}(\mu_{st}^{\text{hsw}}, y_{st}) = \text{Cov}(\mu_{st,\text{RN}}^{\text{hsw}}, y_{st}) + \alpha\left(\mu_{st,\text{RN}}^{\text{hsw}}\right)^2 \text{Cov}\left(\sum_{i\in\mathcal{I}_{st}} w_{st,i} \sum_k \frac{\partial \mathcal{R}_{it}}{\partial \sigma_{k,it}^2} \Delta\sigma_{k,it}^2, y_{st}\right) + \mathcal{O}(\alpha^2) + \mathcal{O}(\|\Delta\sigma^2\|^2).$$

The term $\mathcal{R}_{it}(\bar{\sigma}^2)$ is $\mathcal{F}_{st}^{\text{pre}}$-measurable and so contributes to $\text{Cov}(\mu_{st,\text{RN}}^{\text{hsw}}, y_{st})$ rather than the risk-induced covariance. By the law of total covariance, splitting into the conditional component and the component through $\mathcal{F}_{st}^{\text{pre}}$ variation and using the $\mathcal{F}_{st}^{\text{pre}}$-measurability of w_{st} and $\partial\mathcal{R}_{it}/\partial\sigma_{k,it}^2$,

$$\text{Cov}\left(\sum_i w_{st,i} \sum_k \frac{\partial \mathcal{R}_{it}}{\partial \sigma_{k,it}^2} \Delta\sigma_{k,it}^2, y_{st}\right) = \sum_i w_{st,i} \sum_k \frac{\partial \mathcal{R}_{it}}{\partial \sigma_{k,it}^2} \text{Cov}(\sigma_{k,it}^2, y_{st}) + \mathcal{R}_w,$$

where the residual $\mathcal{R}_w$ collects covariance terms generated by cross-t variation in the pre-period-measurable weights $w_{st,i}$ and partials $\partial\mathcal{R}_{it}/\partial\sigma_{k,it}^2$. Two regimes suppress $\mathcal{R}_w$ to the stated remainder order: (i) a small-α, small-shock asymptotic $(\alpha, \|\Delta\sigma^2\|) \to (0,0)$, under which $\mathcal{R}_w$ is $\mathcal{O}(\alpha\|\Delta\sigma^2\|^2)$ and falls into the remainder; (ii) a stationary-weights assumption that the cross-sectional distribution of $(\omega_{it}, \overline{\kappa}_{it}, \mu_{it,\text{RN}})$ is constant in t, so that $w_{st,i}$ has no time-series variation and $\mathcal{R}_w = 0$ exactly. Either regime yields the stated cyclicality expression. □

B.7 Proof of Proposition 5

Proposition 5 (Ex-ante / ex-post aggregation equivalence). *Under the assumptions of Theorem 1 (including Assumption A0, price commitment), with $\overline{\kappa}_{it} = \mathbb{E}[\kappa_{it} \mid \mathcal{F}_{it}^{pre}]$, the continuum-firm harmonic sales-weighted aggregate satisfies $\mu_{st}^{hsw,ep} = \mu_{st}^{hsw,ea}$.*

Proof. Write the cost decomposition as $\kappa_{it} = \overline{\kappa}_{it} + \eta_{it}$ with $\mathbb{E}[\eta_{it} \mid \mathcal{F}_{it}^{\text{pre}}] = 0$, and suppose additionally that the η_{it} are conditionally independent across firms given $\{\mathcal{F}_{it}^{\text{pre}}\}_i$ with uniformly bounded second moments, $\sup_i \mathbb{E}[\eta_{it}^2 \mid \mathcal{F}_{it}^{\text{pre}}] < \infty$. The harmonic sales-weighted aggregate expands as

$$\frac{1}{\mu_{st}^{\text{hsw, ep}}} = \frac{\int y_{it}\,\kappa_{it}/p_{it}\,di}{\int y_{it}\,di} = \frac{\int y_{it}\,\overline{\kappa}_{it}/p_{it}\,di}{\int y_{it}\,di} + \frac{\int y_{it}\,\eta_{it}/p_{it}\,di}{\int y_{it}\,di}.$$

Under Assumption A0 both y_{it} and p_{it} are $\mathcal{F}_{it}^{\text{pre}}$-measurable, so the integrand $y_{it}\eta_{it}/p_{it}$ has conditional mean zero given $\mathcal{F}_{it}^{\text{pre}}$. Applying a cross-sectional law of large numbers (for conditionally-independent sequences with uniformly bounded second moments) to the continuum of firms gives

$$\int y_{it}\,\eta_{it}/p_{it}\,di \;=\; \int \frac{y_{it}}{p_{it}}\,\mathbb{E}\Big[\eta_{it}\,|\,\mathcal{F}_{it}^{\text{pre}}\Big]\,di \;=\; 0 \qquad \text{(a.s.)}.$$

Therefore

$$\frac{1}{\mu_{st}^{\text{hsw, ep}}} = \frac{\int y_{it}\,\overline{\kappa}_{it}/p_{it}\,di}{\int y_{it}\,di} = \sum_i \omega_{it}\frac{1}{\mu_{it}} = \frac{1}{\mu_{st}^{\text{hsw, ea}}},$$

using $\omega_{it} = p_{it}y_{it}/\int p_{jt}y_{jt}\,dj$ and $\mu_{it} = p_{it}/\overline{\kappa}_{it}$. This establishes $\mu_{st}^{\text{hsw, ep}} = \mu_{st}^{\text{hsw, ea}}$.

For a finite sample of N_{st} firms, under conditional independence of η_{it} across i with uniformly bounded variance and uniformly bounded y_{it}/p_{it}, the residual $\frac{1}{N_{st}}\sum_i y_{it}\,\eta_{it}/p_{it}$ is $\mathcal{O}_p(1/\sqrt{N_{st}})$ by the central limit theorem, yielding the stated finite-sample approximation.

The proof uses four ingredients: the mean-zero property of η_{it} (standard decomposition of κ); the $\mathcal{F}_{it}^{\text{pre}}$-measurability of y_{it} and p_{it} (price commitment, Assumption A0); conditional independence of η_{it} across firms; and the LLN-consistent continuum-firm aggregation. Note that under the CES/isoelastic demand used in the empirical markup literature, the aggregate price index P_t is itself $\mathcal{F}_t^{\text{pre}}$-measurable under Assumption A0 (all firms post prices before shocks realize), so y_{it} remains $\mathcal{F}_{it}^{\text{pre}}$-measurable and the equivalence continues to hold. The equivalence can fail in environments that depart from Assumption A0 — for example, if a subset of firms price flexibly on realized costs so that their p_{it} is no longer $\mathcal{F}_{it}^{\text{pre}}$-measurable, p_{it} covaries with η_{it} cross-sectionally and the residual term does not vanish. □

B.8 Proof of Corollary 1

Corollary 1 (Risk-adjusted pricing formula). *Define the risk-adjusted expectation at p by*

$$\mathbb{E}_p^{\text{RA}}[X] := \frac{\mathbb{E}\Big[U'\Big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\Big)X\Big]}{\mathbb{E}\Big[U'\Big(\omega(\boldsymbol{\theta}) + \pi(p;\boldsymbol{\theta})\Big)\Big]}.$$

Under the assumptions of Theorem 1, the optimal price satisfies

$$\mathcal{MR}_{\text{RA}}(p^*) = \mathcal{MC}_{\text{RA}}(p^*),$$

where risk-adjusted marginal revenue and marginal cost are given by

$$\mathcal{MR}_{\mathrm{RA}}(p) := p + \frac{\mathbb{E}_p^{\mathrm{RA}}[y(p;\boldsymbol{\theta})]}{\mathbb{E}_p^{\mathrm{RA}}[y_p(p;\boldsymbol{\theta})]}, \qquad \mathcal{MC}_{\mathrm{RA}}(p) := \frac{\mathbb{E}_p^{\mathrm{RA}}[\mathcal{MC}(p;\boldsymbol{\theta})y_p(p;\boldsymbol{\theta})]}{\mathbb{E}_p^{\mathrm{RA}}[y_p(p;\boldsymbol{\theta})]}.$$

Equivalently, these objects admit the decomposition

$$\mathcal{MR}_{\mathrm{RA}}(p) = \underbrace{p + \frac{\mathbb{E}[y(p;\boldsymbol{\theta})]}{\mathbb{E}[y_p(p;\boldsymbol{\theta})]}}_{\mathcal{MR}_{\mathrm{mean}}(p)} + \underbrace{\left(\frac{\mathbb{E}_p^{\mathrm{RA}}[y]}{\mathbb{E}_p^{\mathrm{RA}}[y_p]} - \frac{\mathbb{E}[y]}{\mathbb{E}[y_p]}\right)}_{\mathcal{RA}_{\mathrm{MR}}(p)}$$

$$\mathcal{MC}_{\mathrm{RA}}(p) = \underbrace{\frac{\mathbb{E}\big[\mathcal{MC}(p;\boldsymbol{\theta})\, y_p(p;\boldsymbol{\theta})\big]}{\mathbb{E}[y_p(p;\boldsymbol{\theta})]}}_{\mathcal{MC}_{\mathrm{mean}}(p)} + \underbrace{\left(\frac{\mathbb{E}_p^{\mathrm{RA}}[\mathcal{MC}y_p]}{\mathbb{E}_p^{\mathrm{RA}}[y_p]} - \frac{\mathbb{E}[\mathcal{MC}y_p]}{\mathbb{E}[y_p]}\right)}_{\mathcal{RA}_{\mathrm{MC}}(p)}.$$

Proof. The first-order condition of Theorem 1 can be written as

$$\mathbb{E}_{\boldsymbol{\theta}}\Big[U'(\omega+\pi)y\Big] + p^*\mathbb{E}_{\boldsymbol{\theta}}\Big[U'(\omega+\pi)y_p\Big] - \mathbb{E}_{\boldsymbol{\theta}}\Big[U'(\omega+\pi)\mathcal{MC}\ y_p\Big] = 0$$

using the identity $\pi_p(p;\boldsymbol{\theta}) = y(p;\boldsymbol{\theta}) + [p - \mathcal{MC}(p;\boldsymbol{\theta})]y_p(p;\boldsymbol{\theta})$, and all objects are evaluated at $(p^*, \boldsymbol{\theta})$ to simplify notation.

Since utility is strictly increasing by A1, marginal utility satisfies $U'(\omega+\pi) > 0$ almost surely, and hence $\mathbb{E}_{\boldsymbol{\theta}}[U'(\omega(\boldsymbol{\theta}) + \pi(p^*;\boldsymbol{\theta}))] > 0$. Dividing the first-order condition by this term and using the definition of the risk-adjusted expectation yields

$$\mathbb{E}_{p^*}^{\mathrm{RA}}[y] + p^*\mathbb{E}_{p^*}^{\mathrm{RA}}[y_p] - \mathbb{E}_{p^*}^{\mathrm{RA}}[\mathcal{MC}\ y_p] = 0.$$

By A2, $y_p(p;\boldsymbol{\theta}) < 0$ for all $p > 0$ and almost all $\boldsymbol{\theta}$. Because the risk-adjusted expectation uses strictly positive weights, it follows that $\mathbb{E}_{p^*}^{\mathrm{RA}}[y_p] < 0$. Dividing by this term and rearranging gives

$$p^* + \frac{\mathbb{E}_{p^*}^{\mathrm{RA}}[y]}{\mathbb{E}_{p^*}^{\mathrm{RA}}[y_p]} = \frac{\mathbb{E}_{p^*}^{\mathrm{RA}}[\mathcal{MC}\ y_p]}{\mathbb{E}_{p^*}^{\mathrm{RA}}[y_p]}.$$

By definition, the left-hand side is the risk-adjusted marginal revenue $\mathcal{MR}_{\mathrm{RA}}(p^*)$ and the right-hand side is the risk-adjusted marginal cost $\mathcal{MC}_{\mathrm{RA}}(p^*)$. The decompositions into mean components and risk adjustments follow by adding and subtracting the corresponding physical-probability expressions. □

B.9 Proof of Corollary 3

Corollary 3 (Markup cyclicality with CARA utility, linear demand, and constant marginal cost). *Let preferences be $U_\alpha(\mathcal{W}_{it}) = -\exp(-\alpha\mathcal{W}_{it})$. Marginal cost is $\kappa_{it} = \overline{\kappa}_i + \eta_{it}$ with $\overline{\kappa}_i > 0$, and demand is linear*

$$\mathcal{D}_{it}(p_{it};\boldsymbol{\theta}_{it}) = a + \xi_{it} - (b + \psi_{it})p_{it},$$

where $a, b > 0$ and $a > b\overline{\kappa}_i$. Assume the shocks $\eta_{it}, \xi_{it}, \psi_{it}$ are mutually independent, mean-zero, with time-varying variances $\boldsymbol{\sigma}_{it}^2 = \sigma_{\eta,it}^2, \sigma_{\xi,it}^2, \sigma_{\psi,it}^2$ and finite fourth moments.

The risk-neutral price and markup are

$$p_{i,\text{RN}} = \frac{a + b\overline{\kappa}_i}{2b}, \qquad \mu_{i,\text{RN}} = \frac{p_{i,\text{RN}}}{\overline{\kappa}_i}.$$

Moreover, for sufficiently small α, the optimal ex-ante markup satisfies

$$\mu_{it}(\alpha) = \mu_{i,\text{RN}} + \frac{\alpha}{\overline{\kappa}_i}\mathcal{R}_{it}(\boldsymbol{\sigma}_{it}^2) + \mathcal{O}(\alpha^2),$$

where the (first-order) risk adjustment is

$$\mathcal{R}(\boldsymbol{\sigma}_{it}^2) = \underbrace{\frac{a - b\overline{\kappa}_i}{4}\sigma_{\eta,it}^2}_{\text{pure cost risk}} - \underbrace{\frac{a - b\overline{\kappa}_i}{4b^2}\sigma_{\xi,it}^2}_{\text{pure demand-level risk}} - \underbrace{\frac{a(a + b\overline{\kappa}_i)(a - b\overline{\kappa}_i)}{8b^4}\sigma_{\psi,it}^2}_{\text{pure elasticity risk}} - \underbrace{\frac{a + b\overline{\kappa}_i}{4b^2}\sigma_{\eta,it}^2\sigma_{\psi,it}^2}_{\text{cost}\times\text{elasticity risk}}.$$

Consequently, firm-level markup cyclicality is

$$\begin{aligned}\text{Cov}(\mu_{it}, y_{it}) = \frac{\alpha}{\overline{\kappa}_i}\Bigg[&\frac{a - b\overline{\kappa}_i}{4}\text{Cov}(\sigma_{\eta,it}^2, y_{it}) - \frac{a - b\overline{\kappa}_i}{4b^2}\text{Cov}(\sigma_{\xi,it}^2, y_{it}) - \\ &- \frac{a(a + b\overline{\kappa}_i)(a - b\overline{\kappa}_i)}{8b^4}\text{Cov}(\sigma_{\psi,it}^2, y_{it}) - \frac{a + b\overline{\kappa}_i}{4b^2}\text{Cov}(\sigma_{\eta,it}^2\sigma_{\psi,it}^2, y_{it})\Bigg] + \mathcal{O}(\alpha^2).\end{aligned}$$

Proof. Under the stated environment,

$$\begin{aligned}\pi_{it}(p_{it};\boldsymbol{\theta}_{it}) &= (p_{it} - \overline{\kappa}_i - \eta_{it})\big(a + \xi_{it} - (b + \psi_{it})p_{it}\big) \\ &= (p_{it} - \overline{\kappa}_i)(a - bp_{it}) + (p_{it} - \overline{\kappa}_i)\xi_{it} - (p_{it} - \overline{\kappa}_i)p_{it}\psi_{it} - (a - bp_{it})\eta_{it} - \eta_{it}\xi_{it} + p_{it}\eta_{it}\psi_{it}.\end{aligned}$$

Since $\eta_{it}, \xi_{it}, \psi_{it}$ are mutually independent with zero means,

$$\mathbb{E}_{\boldsymbol{\theta}_{it}}[\pi_{it}(p_{it};\boldsymbol{\theta}_{it})] = (p_{it} - \overline{\kappa}_i)(a - bp_{it}),$$

so the risk-neutral price solves $\mathbb{E}[\pi_{it,p}(p_{it};\boldsymbol{\theta}_{it})] = 0$, and hence

$$p_{i,\text{RN}} = \frac{a + b\overline{\kappa}_i}{2b}, \qquad \mathbb{E}_{\theta_{it}}[\pi_{it,pp}(p_{it};\boldsymbol{\theta}_{it})]\Big|_{p_{it}=p_{i,\text{RN}}} = -2b.$$

Since $\mathbb{E}_{\theta_{it}}[y_{it,p,\text{RN}}] = -b$, it follows that $\varphi = 1/2$.

Next compute $\text{Cov}_{\theta_{it}}(\pi_{it}, \pi_{it,p})$. Differentiating profits,

$$\begin{aligned}\pi_{it,p}(p_{it};\boldsymbol{\theta}_{it}) &= \Big(a + \xi_{it} - (b + \psi_{it})p_{it}\Big) + (p_{it} - \overline{\kappa}_i - \eta_{it})\Big(-(b + \psi_{it})\Big) \\ &= (a + \xi_{it} - 2bp_{it} + b\overline{\kappa}_i) - (2p_{it} - \overline{\kappa}_i)\psi_{it} + b\eta_{it} + \eta_{it}\psi_{it}.\end{aligned}$$

Using again independence and mean-zero shocks,

$$\begin{aligned}\pi_{it} - \mathbb{E}[\pi_{it}] &= (p_{it} - \overline{\kappa}_i)\xi_{it} - (p_{it} - \overline{\kappa}_i)p_{it}\psi_{it} - (a - bp_{it})\eta_{it} - \eta_{it}\xi_{it} + p_{it}\eta_{it}\psi_{it}, \\ \pi_{it,p} - \mathbb{E}[\pi_{it,p}] &= \xi_{it} - (2p_{it} - \overline{\kappa}_i)\psi_{it} + b\eta_{it} + \eta_{it}\psi_{it}.\end{aligned}$$

It follows that

$$\text{Cov}_{\theta_{it}}(\pi_{it}, \pi_{it,p}) = (p_{it} - \overline{\kappa}_i)\,\sigma^2_{\xi,it} - b(a - bp_{it})\sigma^2_{\eta,it} + (p_{it} - \overline{\kappa}_i)(2p_{it} - \overline{\kappa}_i)p_{it}\sigma^2_{\psi,it} + p_{it}\sigma^2_{\eta,it}\sigma^2_{\psi,it}.$$

Evaluating at $p_{it} = p_{i,\text{RN}}$ and applying the CARA specialization of Theorem 2,

$$p_{it}(\alpha) = p_{i,\text{RN}} + \frac{\alpha}{2}\frac{\text{Cov}_{\theta_{it}}(\pi_{it,\text{RN}}, \pi_{it,p,\text{RN}})}{\mathbb{E}_{\theta_{it}}[y_{it,p,\text{RN}}]} + \mathcal{O}(\alpha^2).$$

The cyclicality expression follows from $\mu_{it} = p_{it}/\overline{\kappa}_i$ with deterministic $\overline{\kappa}_i$, so $\text{Cov}(\mu_{it}, y_{it}) = (1/\overline{\kappa}_i)\text{Cov}(p_{it}, y_{it})$, and from the linearity of the covariance applied to $\mathcal{R}_{it}(\boldsymbol{\sigma}^2_{it})$, which is affine in $\sigma^2_{\eta,it}, \sigma^2_{\xi,it}, \sigma^2_{\psi,it}$, and $\sigma^2_{\eta,it}\sigma^2_{\psi,it}$. □

B.10 Proof of Corollary 4

Corollary 4 (Markup cyclicality with CARA utility, isoelastic demand, and constant marginal cost). *Let preferences be $U_\alpha(\mathcal{W}_{it}) = -\exp(-\alpha\mathcal{W}_{it})$. Marginal cost is $\kappa_{it} = \overline{\kappa}_i + \eta_{it}$ with $\overline{\kappa}_i > 0$, and demand is isoelastic*

$$\mathcal{D}_{it}(p_{it};\boldsymbol{\theta}_{it}) = (1 + \xi_{it})p_{it}^{-(\varepsilon+\psi_{it})},$$

where $\varepsilon > 1$ and $\boldsymbol{\theta}_{it} = (\eta_{it}, \xi_{it}, \psi_{it})$. Assume the shocks $\eta_{it}, \xi_{it}, \psi_{it}$ are mutually independent, mean-zero, with time-varying variances $\boldsymbol{\sigma}^2_{it} = (\sigma^2_{\eta,it}, \sigma^2_{\xi,it}, \sigma^2_{\psi,it})$ and finite fourth moments.

The risk-neutral price (for small elasticity uncertainty) is

$$p_{it,\mathrm{RN}} = \overline{p}_i + \frac{\overline{\kappa}_i L_i}{(\varepsilon-1)^2}\,\sigma^2_{\psi,it} + o(\sigma^2_{\psi,it}), \qquad \overline{p}_i := \frac{\varepsilon}{\varepsilon-1}\,\overline{\kappa}_i, \qquad L_i := \ln \overline{p}_i,$$

where $\overline{p}_i$ is the deterministic monopoly price, and the corresponding markup is $\mu_{it,RN} := p_{it,RN}/\,\overline{\kappa}_i$. Moreover, for sufficiently small α, the optimal ex-ante markup satisfies

$$\mu_{it}(\alpha) = \mu_{it,\mathrm{RN}} + \frac{\alpha}{\overline{\kappa}_i}\mathcal{R}_{it}(\boldsymbol{\sigma}^2_{it}) + \mathcal{O}(\alpha^2),$$

where the first-order risk adjustment decomposes as

$$\mathcal{R}_{it}(\boldsymbol{\sigma}^2_{it}) = \underbrace{C_{i\eta}\,\sigma^2_{\eta,it}}_{\text{pure cost risk}} - \underbrace{\Theta_{i\psi}\,\sigma^2_{\psi,it}}_{\text{pure elasticity risk}} + \underbrace{C_{i\eta}\sigma^2_{\eta,it}\sigma^2_{\xi,it}}_{\text{cost} \times \text{demand-level risk}} - \underbrace{\Theta_{i\eta\psi}\,\sigma^2_{\eta,it}\sigma^2_{\psi,it}}_{\text{cost} \times \text{elasticity risk}}.$$

The coefficients (evaluated at $\overline{p}_i$) are

$$C_{i\eta} := \frac{\left(\frac{\varepsilon-1}{\varepsilon}\right)^{\varepsilon-1}}{\overline{\kappa}_i^{\varepsilon}}, \qquad \Theta_{i\psi} := \frac{(\overline{p}_i - \overline{\kappa}_i)^2}{\varepsilon-1}\,\overline{p}_i^{-\varepsilon} L_i, \qquad \Theta_{i\eta\psi} := \frac{\overline{p}_i^{-\varepsilon}}{\varepsilon-1}\left(2L_i - 2\varepsilon L_i^2\right).$$

Consequently, firm-level markup cyclicality is

$$\begin{aligned}\mathrm{Cov}(\mu_{it}, y_{it}) = \frac{L_i}{(\varepsilon-1)^2}\mathrm{Cov}(\sigma^2_{\psi,it}, y_{it}) + \frac{\alpha}{\overline{\kappa}_i}\Big[C_{i\eta}\mathrm{Cov}(\sigma^2_{\eta,it}, y_{it}) \\ - \Theta_{i\psi}\mathrm{Cov}(\sigma^2_{\psi,it}, y_{it}) + C_{i\eta}\mathrm{Cov}(\sigma^2_{\eta,it}\sigma^2_{\xi,it}, y_{it}) - \Theta_{i\eta\psi}\mathrm{Cov}(\sigma^2_{\eta,it}\sigma^2_{\psi,it}, y_{it})\Big] + \mathcal{O}(\alpha^2).\end{aligned}$$

Proof. Under the stated environment,

$$\begin{aligned}\pi_{it}(p_{it};\boldsymbol{\theta}_{it}) &= (p_{it} - \overline{\kappa}_i - \eta_{it})(1+\xi_{it})p_{it}^{-(\varepsilon+\psi_{it})} \\ &= (p_{it} - \overline{\kappa}_i - \eta_{it})(1+\xi_{it})p^{-\varepsilon}e^{-\psi_{it}\ln p_{it}}.\end{aligned}$$

Under CARA utility and using the second-order certainty-equivalent approximation (neglecting higher cumulants),

$$\mathrm{CE}(p_{it};\alpha) = \mathbb{E}[\pi_{it}(p_{it})] - \frac{\alpha}{2}\mathrm{Var}(\pi_{it}(p_{it})).$$

Let $F(p_{it},\alpha) := \partial_p \mathrm{CE}(p_{it};\alpha)$. The optimal price satisfies $F(p_{it},\alpha) = 0$.

Since $\mathbb{E}[\eta_{it}] = \mathbb{E}[\xi_{it}] = 0$ and shocks are independent,

$$\mathbb{E}[\pi_{it}(p_{it})] = (p_{it} - \overline{\kappa}_i)p_{it}^{-\varepsilon}M(p_{it}), \qquad M(p_{it}) := \mathbb{E}\left[e^{-\psi_{it}\ln p_{it}}\right].$$

The risk-neutral first-order condition $F(p_{it},0)=0$ is

$$\frac{\mathrm{d}}{\mathrm{d}p_{it}}\Big[(p_{it}-\overline{\kappa}_i)p_{it}^{-\varepsilon}M(p_{it})\Big]=0 \qquad \Longleftrightarrow \qquad \frac{p_{it}}{p_{it}-\overline{\kappa}_i}=\varepsilon-J(p_{it}),$$

where

$$J(p_{it}):=\frac{\mathrm{d}\ln M(p_{it})}{\mathrm{d}\ln p_{it}}=\frac{p_{it}M'(p_{it})}{M(p_{it})}.$$

For small $\sigma^2_{\psi,it}$, the cumulant expansion of ψ_{it} gives

$$M(p_{it})=1+\frac{1}{2}\sigma^2_{\psi,it}(\ln p_{it})^2+\mathcal{O}(\sigma^2_{\psi,it}), \qquad J(p_{it})=\sigma^2_{\psi,it}\ln p_{it}+\mathcal{O}(\sigma^2_{\psi,it}).$$

Let $\overline{p}_i:=\frac{\varepsilon}{\varepsilon-1}\overline{\kappa}_i$ and $L_i:=\ln\overline{p}_i$. Implicit differentiation of the risk-neutral condition at $(p_{it},\sigma^2_{\psi,it})=(\overline{p}_i,0)$ yields

$$\left.\frac{\mathrm{d}p_{it,\mathrm{RN}}}{\mathrm{d}\sigma^2_{\psi,it}}\right|_{\sigma^2_{\psi,it}=0}=\frac{\overline{\kappa}_i}{(\varepsilon-1)^2}L_i,$$

so the risk-neutral price with small elasticity risk is

$$p_{it,\mathrm{RN}}(\sigma^2_{\psi,it})=\overline{p}_i+\frac{\overline{\kappa}_iL_i}{(\varepsilon-1)^2}\sigma^2_{\psi,it}+\mathcal{O}(\sigma^2_{\psi,it}).$$

With $X_{it}:=p_{it}-\overline{\kappa}_i-\eta_{it}$, $A:=1+\xi_{it}$, $Z:=e^{-\psi_{it}\ln p_{it}}$, and independence,

$$\mathrm{Var}(\pi_{it})=p_{it}^{-2\varepsilon}\Big(\mathbb{E}\big[X^2_{it}\big]\mathbb{E}\big[A^2_{it}\big]\mathbb{E}\big[Z^2_{it}\big]-\mathbb{E}[X_{it}]^2\,\mathbb{E}[A_{it}]^2\,\mathbb{E}[Z_{it}]^2\Big).$$

Using $\mathbb{E}[X_{it}]=p_{it}-\overline{\kappa}_i$, $\mathbb{E}\big[X^2_{it}\big]=(p_{it}-\overline{\kappa}_i)^2+\sigma^2_{\eta,it}$, $\mathbb{E}[A_{it}]=1$, $\mathbb{E}\big[A^2_{it}\big]=1+\sigma^2_{\xi,it}$, and

$$\mathbb{E}[Z_{it}]=1+\frac{1}{2}(\ln p_{it})^2\sigma^2_{\psi,it}+\mathcal{O}(\sigma^2_{\psi,it}), \qquad \mathbb{E}\big[Z^2_{it}\big]=1+2(\ln p_{it})^2\sigma^2_{\psi,it}+\mathcal{O}(\sigma^2_{\psi,it}),$$

we obtain, to first order in $\sigma^2_{\psi,it}$,

$$\mathrm{Var}(\pi_{it}(p_{it}))=p_{it}^{-2\varepsilon}\Big[(p_{it}-\overline{\kappa}_i)^2\sigma^2_{\xi,it}+\sigma^2_{\eta,it}(1+\sigma^2_{\xi,it})+(p_{it}-\overline{\kappa}_i)^2(\ln p_{it})^2\sigma^2_{\psi,it}+2\sigma^2_{\eta,it}(\ln p_{it})^2\sigma^2_{\psi,it}\Big].$$

At $\alpha=0$, the optimum is $p_{it,\mathrm{RN}}(\sigma^2_{\psi,it})$ and satisfies $F(p_{it,\mathrm{RN}},0)=0$. Since

$$F(p_{it},\alpha)=\mathbb{E}[\pi_{it,p}(p_{it})]-\frac{\alpha}{2}\partial_p\mathrm{Var}(\pi_{it}(p_{it})),$$

we have

$$F_\alpha(p_{it},0) = -\frac{1}{2}\partial_p \mathrm{Var}(\pi_{it}(p_{it})).$$

The implicit function theorem yields the expansion

$$p_{it}(\alpha) = p_{it,\mathrm{RN}}(\sigma^2_{\psi,it}) - \alpha\frac{F_\alpha(p_{it,\mathrm{RN}}(\sigma^2_{\psi,it}),0)}{F_p(p_{it,\mathrm{RN}}(\sigma^2_{\psi,it}),0)} + \mathcal{O}(\alpha^2).$$

To the displayed order, we may evaluate the α-correction at $(p_{it},\sigma^2_{\psi,it}) = (\overline{p}_i,0)$. At $\sigma^2_{\psi,it}=0$, $\mathbb{E}[\pi_{it}(p_{it})] = (p_{it}-\overline{\kappa}_i)p_{it}^{-\varepsilon}$ and hence

$$F_p(\overline{p}_i,0) = -(\varepsilon-1)\overline{p}_i^{-\varepsilon-1}.$$

Differentiating the variance expansion and evaluating at $p_{it}=\overline{p}_i$ gives

$$F_\alpha(\overline{p}_i,0) = \overline{p}_i^{-2\varepsilon-1}\Big[\varepsilon\sigma^2_{\eta,it}(1+\sigma^2_{\xi,it}) - (\overline{p}_i-\overline{\kappa}_i)^2 L_i\sigma^2_{\psi,it} - \sigma^2_{\eta,it}\sigma^2_{\psi,it}(2L_i - 2\varepsilon L_i^2)\Big].$$

Therefore,

$$-\alpha\frac{F_\alpha(\overline{p}_i,0)}{F_p(\overline{p}_i,0)} = \alpha\Bigg[\underbrace{\frac{\varepsilon}{\varepsilon-1}\overline{p}_i^{-\varepsilon}}_{\equiv C_{i\eta}}\sigma^2_{\eta,it}(1+\sigma^2_{\xi,it}) - \underbrace{\frac{(\overline{p}_i-\overline{\kappa}_i)^2}{\varepsilon-1}\overline{p}_i^{-\varepsilon}L_i}_{\equiv\Theta_{i\psi}}\sigma^2_{\psi,it} - \underbrace{\frac{\overline{p}_i^{-\varepsilon}}{\varepsilon-1}(2L_i-2\varepsilon L_i^2)}_{\equiv\Theta_{i\eta\psi}}\sigma^2_{\eta,it}\sigma^2_{\psi,it}\Bigg].$$

Combining this α-correction with the risk-neutral expansion for $p_{\mathrm{RN},t}$ yields

$$p_{it}(\alpha) = p_{it,\mathrm{RN}} + \alpha\mathcal{R}_{it}(\boldsymbol{\sigma}^2_{it}) + \mathcal{O}(\alpha^2),$$

with the stated decomposition of $\mathcal{R}_{it}(\boldsymbol{\sigma}^2_{it})$.

Dividing both sides by $\overline{\kappa}_i$, taking the covariance with y_{it} and using the bilinearity of covariance,

$$\mathrm{Cov}(\mu_{it},y_{it}) = \mathrm{Cov}(\mu_{it,\mathrm{RN}},y_{it}) + \frac{\alpha}{\overline{\kappa}_i}\mathrm{Cov}(\mathcal{R}_{it}(\boldsymbol{\sigma}^2_{it}),y_{it}),$$

where, from the risk-neutral expansion,

$$\mathrm{Cov}(\mu_{it,\mathrm{RN}},y_{it}) = \frac{L_i}{(\varepsilon-1)^2}\mathrm{Cov}(\sigma^2_{\psi t},Y_t).$$

Finally, substituting the decomposition of $\mathcal{R}_{it}(\boldsymbol{\sigma}_{it}^2)$ and applying the linearity of covariance yields

$$\begin{aligned}\mathrm{Cov}(\mathcal{R}_{it}(\boldsymbol{\sigma}_{it}^2), y_{it}) = C_{i\eta}\mathrm{Cov}(\sigma_{\eta,it}^2, y_{it}) - \Theta_{i\psi}\mathrm{Cov}(\sigma_{\psi,it}^2, y_{it}) \\ + C_{i\eta}\mathrm{Cov}(\sigma_{\eta,it}^2\sigma_{\xi,it}^2, y_{it}) - \Theta_{i\eta\psi}\mathrm{Cov}(\sigma_{\eta,it}^2\sigma_{\psi,it}^2, y_{it}).\end{aligned}$$

This completes the proof. □

B.11 Proof of Corollary 5

Corollary 5 (Markup cyclicality with CRRA utility, linear demand, and constant marginal cost). *Consider the same environment as in Corollary 3, except that preferences are of the CRRA form*

$$U(\mathcal{W}_{it}) = \frac{\mathcal{W}_{it}^{1-\sigma}}{1-\sigma}, \qquad \sigma > 0,$$

so that absolute risk aversion satisfies $A_\sigma(\mathcal{W}_{it}) = \sigma/\mathcal{W}_{it} = \sigma A(\mathcal{W}_{it})$ *with* $A(\mathcal{W}_{it}) = 1/\mathcal{W}_{it}$, *where* $\mathcal{W}_{it}(\boldsymbol{\theta}_{it}) = \omega_{it}(\boldsymbol{\theta}_{it}) + \pi_{it}(p_{it};\boldsymbol{\theta}_{it})$ *is total wealth. In this case, the ex-ante markup satisfies*

$$\mu_{it}(\sigma) = \mu_{i,\mathrm{RN}} + \frac{\sigma}{\overline{\kappa}_i}\mathcal{R}_{it} + \mathcal{O}(\sigma^2),$$

where $\mathcal{R}_{it} = \mathcal{R}_{risk}(\boldsymbol{\sigma}_{it}^2) + \mathcal{R}_{wealth}(\mathcal{W}_{it,RN})$, *and*

$$\mathcal{R}_{\mathrm{risk}}(\boldsymbol{\sigma}_{it}^2) = \overline{A}_{it}\mathcal{R}_{it}^{\mathrm{CARA}}(\boldsymbol{\sigma}_{it}^2), \qquad \overline{A}_{it} := \mathbb{E}_{\boldsymbol{\theta}}\left[\frac{1}{\mathcal{W}_{it,\mathrm{RN}}}\right],$$

$$\mathcal{R}_{\mathrm{wealth}}(\mathcal{W}_{it,\mathrm{RN}}) = \frac{1}{2}\cdot\frac{\mathrm{Cov}_{\boldsymbol{\theta}_{it}}\left(\frac{1}{\mathcal{W}_{it,\mathrm{RN}}}, \pi_{it,\mathrm{RN}}\pi_{p,it,\mathrm{RN}}\right)}{\mathbb{E}_{\boldsymbol{\theta}_{it}}[y_{p,it,\mathrm{RN}}]}.$$

Here, $\mathcal{R}_{it}^{CARA}(\boldsymbol{\sigma}_{it}^2)$ *denotes the risk adjustment from Corollary 3.*

Consequently, firm-level markup cyclicality is

$$\mathrm{Cov}(\mu_{it}, y_{it}) = \frac{\sigma}{\overline{\kappa}_i}\left[\mathrm{Cov}\left(\overline{A}_{it}\mathcal{R}_{it}^{\mathrm{CARA}}(\boldsymbol{\sigma}_{it}^2), y_{it}\right) + \mathrm{Cov}\left(\mathcal{R}_{\mathrm{wealth}}(\mathcal{W}_{it,\mathrm{RN}}), y_{it}\right)\right] + \mathcal{O}(\sigma^2).$$

Proof. The risk-neutral price is unchanged under CRRA preferences and coincides with that in Corollary 3. By Theorem 2, for sufficiently small σ, the optimal price satisfies:

$$p_{it}(\sigma) = p_{i,\mathrm{RN}} + \sigma\varphi\left[\frac{\overline{A}_{it}\mathrm{Cov}_{\boldsymbol{\theta}_{it}}(\pi_{it,\mathrm{RN}}, \pi_{p,it,\mathrm{RN}})}{\mathbb{E}_{\boldsymbol{\theta}_{it}}[y_{p,it,\mathrm{RN}}]} + \frac{\mathrm{Cov}_{\boldsymbol{\theta}_{it}}(A(\mathcal{W}_{it,\mathrm{RN}}), \pi_{it,\mathrm{RN}}\pi_{p,it,\mathrm{RN}})}{\mathbb{E}_{\boldsymbol{\theta}_{it}}[y_{p,it,\mathrm{RN}}]}\right] + \mathcal{O}(\sigma^2),$$

where $A(\mathcal{W}_{it}) = 1/\mathcal{W}_{it}$ under CRRA preferences and $\overline{A}_{it} = \mathbb{E}_{\theta_{it}}[1/\mathcal{W}_{it,\mathrm{RN}}]$.

In the linear-demand, constant-marginal-cost environment of Corollary 3,

$$\varphi \cdot \frac{\mathrm{Cov}_{\theta_{it}}(\pi_{it,\mathrm{RN}}, \pi_{p,it,\mathrm{RN}})}{\mathbb{E}_{\theta_{it}}[y_{p,it,\mathrm{RN}}]} = \mathcal{R}_{it}^{\mathrm{CARA}}(\sigma_{it}^2),$$

so that the pure-risk component satisfies

$$\mathcal{R}_{\mathrm{risk}}(\sigma_{it}^2) = \overline{A}_{it}\mathcal{R}_{it}^{\mathrm{CARA}}(\sigma_{it}^2).$$

The remaining term corresponds to the wealth-effects component

$$\mathcal{R}_{\mathrm{wealth}}(\mathcal{W}_{it,\mathrm{RN}}) = \varphi \cdot \frac{\mathrm{Cov}_{\theta_{it}}(A(\mathcal{W}_{it,\mathrm{RN}}), \pi_{it,\mathrm{RN}}\pi_{p,it,\mathrm{RN}})}{\mathbb{E}_{\theta_{it}}[y_{p,it,\mathrm{RN}}]},$$

where $\varphi = 1/2$. Since the markup is $\mu_{it}(\sigma) = p_{it}(\sigma)/\,\overline{\kappa}_i$ with deterministic $\overline{\kappa}_i > 0$, the expansion for $\mu_{it}(\sigma)$ follows immediately. The expression for markup cyclicality is obtained by taking covariances with y_{it} and applying linearity, as in Proposition 3. □

C Additional Tables and Figures

C.1 Correlation Structure

Table 3: Correlation Matrix of Key Variables

	$\log p$	σ^2_{cost}	$\sigma^2_{\mathrm{demand}}$	C	D
$\log p$	1.00				
σ^2_{cost}	-0.01	1.00			
$\sigma^2_{\mathrm{demand}}$	-0.17	0.21	1.00		
C	0.34	-0.33	-0.24	1.00	
D	-0.06	0.12	0.39	-0.16	1.00

Notes: Correlations computed at the industry-year level across the estimation sample. Uncertainty regressors are z-scored rolling variances of AR(1) innovations; C and D are log-level Bartik shifters.

C.2 Robustness of Main Results

Table 4: Robustness of Main Results

	(1) **Baseline (2SLS)**	(2) + Year FE (2SLS)	(3) First diff. (2SLS)
Dependent variable: $\log p_{it}$ *(cols 1–2),* $\Delta \log p_{it}$ *(col 3)*			
Cost uncertainty (γ_C)	0.0578***	−0.0115	0.0018
	(0.0059)	(0.0123)	(0.0067)
Demand uncertainty (γ_D)	−0.0613***	0.0021	−0.0025
	(0.0121)	(0.0090)	(0.0054)
Bartik level controls	Yes	Yes	Yes
Industry FE	No	No	(differenced)
Year FE	No	Yes	No
Observations	12,649	12,649	12,322

Notes: All three columns are pooled (no industry fixed effects); industry FE would absorb the cross-industry exposure heterogeneity that identifies γ in the Bartik design. Column (1) is the canonical baseline (Table 2 column 3). Column (2) adds year fixed effects, which absorb the aggregate component of σ^2; the within-year (idiosyncratic) component of the effect survives. Column (3) first-differences, for which 2-way clusters by IO code × year are the primary SE (first-differencing amplifies within-year common shocks). Stars are one-sided in the theoretically predicted direction. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

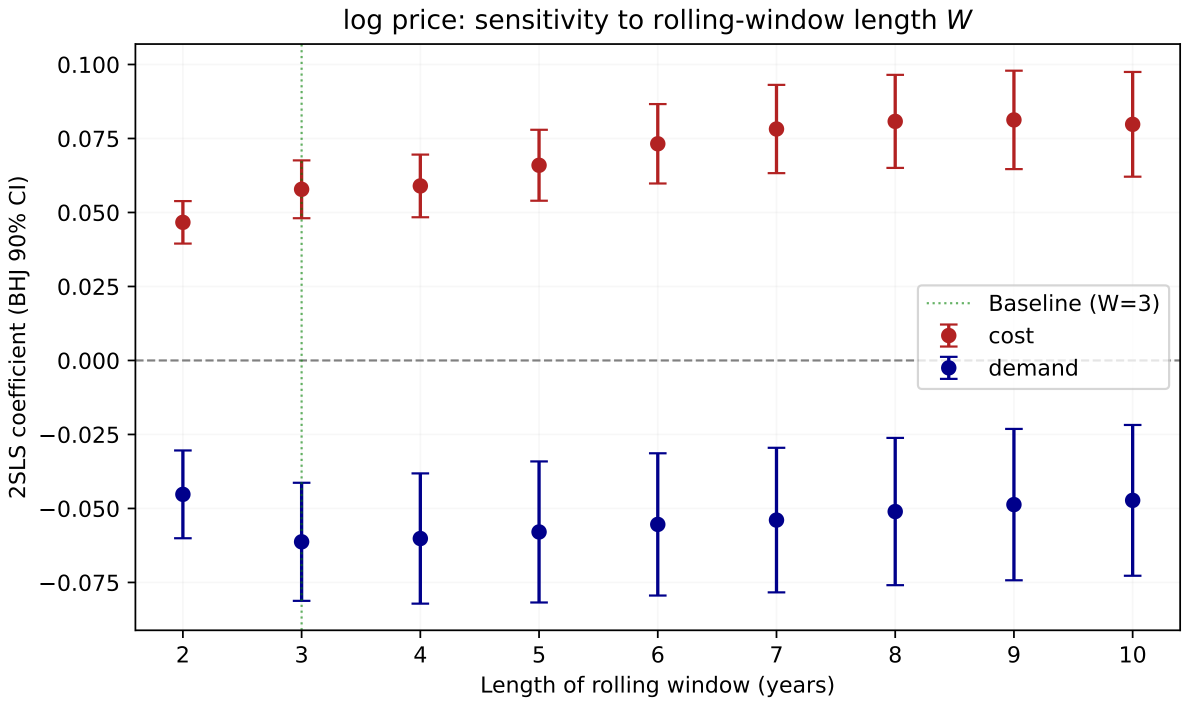

Figure 5: Sensitivity of the baseline pooled 2SLS coefficients $\hat{\gamma}_{\text{cost}}$ and $\hat{\gamma}_{\text{demand}}$ to the rolling-window length W used to construct the conditional-variance shifters $\sigma^2_{k,i,t}$. Bands are 90% confidence intervals using BHJ exposure-robust standard errors clustered at the IO code. The vertical dotted line marks the baseline $W = 3$. The signs and approximate magnitudes of both coefficients are stable across window lengths from $W = 2$ through $W = 8$.

C.3 Aggregate vs. Idiosyncratic Uncertainty

Table 5: Aggregate vs. Idiosyncratic Uncertainty

	(1) Baseline	(2) Decomposed
Dependent variable: $\log p_{it}$		
Cost uncertainty (γ_C)	0.0578***	
	(0.0059)	
Demand uncertainty (γ_D)	–0.0613***	
	(0.0121)	
Cost unc. (aggregate)		0.0802***
		(0.0117)
Cost unc. (idiosyncratic)		0.0859***
		(0.0103)
Demand unc. (aggregate)		–0.2543***
		(0.0336)
Demand unc. (idiosyncratic)		–0.0043
		(0.0096)
Bartik level controls	Yes	Yes
Industry FE	No	No
Observations	12,649	12,649

Notes: Column (2) decomposes each uncertainty regressor into an aggregate component (cross-industry year mean of σ^2) and an idiosyncratic component (deviation from that mean), entering both into the pooled 2SLS baseline. Both channels retain the theoretically predicted signs when decomposed. Stars one-sided in the predicted direction. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. BHJ exposure-robust SEs (cluster at IO code) in parentheses.