Information acquisition, urgent production, and welfare

Pavel Koval*
University of Minnesota

July 19, 2026

Abstract

This paper studies the efficiency of information acquisition when firms hold inventories and face costly urgent production. Information reduces demand uncertainty and improves the allocation of inventories across demand states. The optimal supply function, which summarizes the firm's optimal pricing and production decisions, has constant-markup branches for regular and urgent production and a vertical segment at the firm's inventory position. Firms under-acquire information because they do not internalize its effects on consumers and other firms through the aggregate demand shifter. Evidence from management guidance and quarterly fundamentals for U.S. public firms supports the model's relationships between information, inventories, and costly urgent production. A numerical exercise quantifies the gap between private and welfare-optimal information acquisition and the resulting welfare gains.

Acknowledgments and disclosures. I thank Anmol Bhandari, David Rahman, Kjetil Storesletten, Hannes Malmberg, V. V. Chari, and participants at workshops at the University of Minnesota and at NASM 2026 for helpful comments. This research received no external financial support. I have no financial relationships or other potential conflicts of interest relevant to this research to disclose.

*Email: koval084@umn.edu.

I. Introduction

This paper studies whether firms acquire the socially optimal amount of information when they hold inventories and face costly urgent production. Firms produce in advance because meeting demand after it is realized can require overtime labor, expedited shipping, or more expensive inputs. Inventories reduce these costs, but they are expensive to maintain and depreciate over time. Information about future demand improves the allocation of inventories across states, but acquiring it is costly. The private and social values of this information differ because firms do not internalize its effects on consumers and other firms.

I develop a model with a representative household and a continuum of firms producing differentiated goods. Each firm faces uncertainty about future demand, chooses inventories, and commits to a supply function before demand is realized. It can satisfy demand using inventories net of depreciation or more costly urgent production. The firm also chooses the precision of a signal about future demand. The household has CRRA utility over a CES consumption aggregator, so each firm's demand depends on its idiosyncratic taste shock and on a common demand shifter determined in general equilibrium.

I first characterize the operating problem. The optimal quantity rule, which determines the optimal relation between production and future demand, has three regions. At low demand, the firm sells the monopoly quantity from inventories and charges a constant markup over regular marginal cost. At high demand, it uses urgent production and charges the same markup over urgent marginal cost. Between these regions, the firm sells exactly its available inventories and allows price to clear the market. The resulting supply function has two horizontal constant-markup branches and a vertical segment at the inventory position. This approach does not impose price setting or quantity setting in advance and is related to the supply-function approach of Klemperer and Meyer (1989) and Flynn et al. (2026).

The optimal inventory position increases with expected demand. Its response to posterior uncertainty depends on the location of the inaction region and therefore need not have a fixed sign. Information also changes inventories through the posterior mean: a more informative signal makes inventories more responsive to expected demand. Depreciation is central for the choice of information acquisition and inventory holdings. When depreciation is low, inventories can substitute for precise information because firms can carry a large inventory position cheaply. At intermediate depreciation, information is more useful for targeting inventories across demand states. When depreciation is high, inventories become too costly to use actively and the value of information falls.

The main result concerns the efficiency of information acquisition. At the operating optimum, greater posterior uncertainty lowers normalized operating profit, which is the firm's operating profit holding aggregate consumption and the common demand shifter fixed. An atomistic firm accounts only for this direct effect when choosing its own signal precision and treats aggregate consumption and the common demand shifter as fixed. A change in uncertainty shared by a positive mass of otherwise identical firms also changes these aggregate objects and therefore affects consumer surplus and the profits of all firms. Firm heterogeneity is not necessary for this result. In the symmetric economy, the common change applies to all firms and therefore has the strongest aggregate effect. With heterogeneous firms, a type is simply a group of

identical firms, and the aggregate effect of changing that type's uncertainty is scaled by its population mass. Aggregate operating profit rises with uncertainty when relative risk aversion is greater than one, falls when it is below one, and is locally unchanged under log utility. Nevertheless, at the decentralized equilibrium, increasing uncertainty for any positive mass of firms lowers consumer surplus and welfare for every positive value of relative risk aversion. Firms therefore under-acquire information relative to the welfare optimum.

This result also changes the welfare interpretation of urgency costs. Holding information fixed, a higher urgency cost lowers welfare. However, if it induces firms to acquire sufficiently more information, welfare evaluated along the decentralized-equilibrium information choice can increase. The sign depends on whether this induced information effect dominates the direct cost of urgent production.

I provide supporting empirical evidence using management sales guidance from I/B/E/S and quarterly fundamentals from Compustat for U.S. public firms. Inventories increase with expected sales and positive demand surprises, and firms with less precise guidance hold more inventories on average. Underestimating demand raises COGS relative to sales, consistent with costly urgent production. These results are conditional correlations and are not interpreted as causal estimates. The inventory–uncertainty correlation is positive in low-depreciation industries and negative in industries with higher depreciation. The sign reversal and moderate magnitudes support depreciation heterogeneity: unlike heterogeneity only in urgency costs or prior uncertainty, depreciation heterogeneity can generate correlations that are neither -1 nor 1.

Private information acquisition is nonmonotone in depreciation and falls to zero for sufficiently high depreciation. The planner chooses more information for every type and makes inventory policies more responsive to posterior signals. In the numerical exercise, moving to the welfare-optimal information profile raises welfare by 0.56 percent relative to welfare under the private information profile.

The policy discussion distinguishes sales and production subsidies. Both eliminate the markup and implement the same prices and quantities for a fixed uncertainty profile. However, only the sales subsidy changes the profit–revenue relation and eliminates the information wedge. A production subsidy does not eliminate information under-acquisition. The under-acquisition result is also not specific to CES demand: with linear demand and a quadratic consumption aggregator, consumer surplus remains tied to operating profit at the optimum, and firms again acquire too little information.

The paper contributes to the literature on the positive and normative role of data and information (Farboodi and Veldkamp, 2020; Jones and Tonetti, 2020; Bostanci and Ordoñez, 2025; Asriyan and Kohlhas, 2025). The mechanism here operates through state-contingent inventory and production decisions and through the aggregate effect that firms do not internalize. The paper also relates to work on optimal inventories (Arrow et al., 1951; Kahn, 1987; Khan and Thomas, 2007), the operational properties of inventory systems (Wang and Disney, 2016), and firm behavior under uncertainty (Woodford, 2003; Reis, 2006; Angeletos and La'O, 2010; Llosa and Venkateswaran, 2022).

The paper is organized as follows. Section II presents the model. Section III characterizes the optimal supply function and inventory holdings. Section IV analyzes information acquisition and welfare. Section V presents the empirical evidence and numerical welfare assessment. Section VI discusses policy implications and the linear-demand case. Section VII concludes.

II. Model

The economy consists of a representative household and a continuum of firms indexed by i. The household experiences variety-specific taste shocks $\theta_{i,t}$ and obtains utility $u\big(\int C(q_{i,t}, \theta_{i,t})di\big)$ from differentiated consumption. It faces linear disutility from supplying regular labor L^R and urgent labor L^U. The disutility of regular labor is normalized to one, while urgent labor has cost $\nu > 1$. The parameter ν captures costs such as overtime and unpredictable work schedules. Total per-period utility is

$$u\Big(\int C(q_{i,t}, \theta_{i,t})di\Big) - \int L^R_{i,t}di - \nu \int L^U_{i,t}di, \quad \nu > 1.$$

The first term is utility from aggregate consumption, while the last two terms are the labor costs required to produce it.

Each firm i produces a differentiated good $q_{i,t}$ using regular production $Y^R_{i,t}$ and urgent production $Y^U_{i,t}$. Labor is converted one-for-one into output, so $Y^R_{i,t} = L^R_{i,t}$ and $Y^U_{i,t} = L^U_{i,t}$. Because urgent labor is more costly for the household, it requires higher compensation than regular labor. Thus, urgent production is more expensive than regular production. Realized demand $\hat{q}_{i,t}$ can be satisfied only with inventories $I_{i,t-1}$ net of depreciation and urgent production:

$$\hat{q}_{i,t} \leq Y^U_{i,t} + (1-\delta)I_{i,t-1}.$$

This constraint makes urgent production the residual source of current supply when available inventories are insufficient. Regular production takes the entire period and therefore cannot meet current demand, but it replenishes inventories for the next period. The inventory-flow identity is

$$Y^R_{i,t} + Y^U_{i,t} = \hat{q}_{i,t} - (1-\delta)I_{i,t-1} + I_{i,t}.$$

It states that total production equals current sales plus ending inventories minus usable beginning inventories.

Realized demand is determined by the intersection of the household's demand schedule and the firm's previously chosen supply schedule. Firm i chooses the next-period supply function $f_{i,t}(p_{i,t+1}, q_{i,t+1}) = 0$, which specifies the price–quantity pairs it is willing to implement.

Time is discrete and infinite. At date t, firms observe all current taste shocks $\theta_{i,t}$. Each period has two stages: firms first satisfy realized demand and then choose policies for the next period.

1. Given the previously chosen supply function $f_{i,t-1}(\cdot)$, realized demand $\hat{q}_{i,t}$ is satisfied

   - from past inventories $(1-\delta)I_{i,t-1}$ and urgent production $Y^U_{i,t}$ if $\hat{q}_{i,t} > (1-\delta)I_{i,t-1}$;
   - from past inventories $(1-\delta)I_{i,t-1}$ only if $\hat{q}_{i,t} \leq (1-\delta)I_{i,t-1}$.

2a. Firm i chooses the precision of a signal about next-period demand $\hat{q}_{i,t+1}$ and then receives the signal.

2b. Given its inventory position and signal, firm i chooses the supply function $f_{i,t}(\cdot)$ for the next period and ending inventories $I_{i,t}$.

A. Household problem and log-linear demand

Suppose consumption utility is CRRA and varieties are aggregated using CES. The representative household solves

$$\max_{\{q_{i,t},\, L^R_{i,t},\, L^U_{i,t}\}_{i\in[0,1]}} u_\gamma(C_t) - \int_0^1 L^R_{i,t}\,di - \nu \int_0^1 L^U_{i,t}\,di, \qquad C_t := \left(\int_0^1 \theta_{i,t}^{1/\eta} q_{i,t}^{\frac{\eta-1}{\eta}}\,di\right)^{\frac{\eta}{\eta-1}}$$

subject to the period budget constraint

$$\int_0^1 p_{i,t} q_{i,t}\,di = \int_0^1 L^R_{i,t}\,di + \int_0^1 w_{i,t} L^U_{i,t}\,di + \int_0^1 \pi^0_{i,t}\,di.$$

CRRA utility is defined using the continuous normalization

$$u_\gamma(C) := \begin{cases} \dfrac{C^{1-\gamma}-1}{1-\gamma}, & \gamma \neq 1, \\ \log C, & \gamma = 1. \end{cases}$$

Here $\eta > 1$ is the elasticity of substitution across varieties and $\gamma > 0$ is relative risk aversion. The normalization changes utility only by an additive constant when $\gamma \neq 1$ and makes the log-utility case explicit. The CES aggregator C_t places greater weight on varieties with higher taste realizations. The budget constraint states that consumption expenditure equals labor income plus the operating profits $\pi^0_{i,t}$ rebated by firms.

The household's first-order conditions imply the inverse demand curve and urgent-labor supply correspondence:

$$p_{i,t} = C_t^{1/\eta-\gamma}\theta_{i,t}^{1/\eta} q_{i,t}^{-1/\eta}, \qquad L^{U,*}_{i,t} \in \begin{cases} \{0\}, & w_{i,t} < \nu, \\ [0,\bar{L}], & w_{i,t} = \nu, \\ \{\bar{L}\}, & w_{i,t} > \nu. \end{cases}$$

Inverse demand has elasticity η, and $C_t^{1/\eta-\gamma}$ is the common demand shifter generated by aggregate consumption and marginal utility. Urgent labor is supplied at an interior level only when its wage equals its marginal disutility ν; the regular wage is normalized to one.

B. Firm's problem and equilibrium

Each individual firm behaves as a monopolist on the demand schedule for its variety. The firm maximizes lifetime profit by choosing its supply schedule and inventories. I first solve its problem for an exogenous profile of posterior uncertainty and then study the endogenous information choice separately. Given a profile of posterior uncertainties across firms and time, $\boldsymbol{\sigma} = \{\sigma_{i,\tau}\}_{i\in[0,1],\,\tau\geq t}$, and a state $x_{i,t} = (f_{i,t-1}, I_{i,t-1}, \theta_{i,t})$, firm i's net continuation value solves

$$V_i(x_{i,t};\boldsymbol{\sigma}) = \mathbb{E}\left[\max_{f_{i,t},\, I_{i,t},\, L^R_{i,t},\, L^U_{i,t}} \mathbb{E}\left[\pi^0_{i,t} - \mathcal{I}(\sigma_{i,t}) + \beta V_i(x_{i,t+1};\boldsymbol{\sigma}) \mid \theta_{i,t}, s_{i,t}\right] \middle| \theta_{i,t}\right], \tag{1}$$

$$\pi^0_{i,t} := \hat{p}_{i,t}\hat{q}_{i,t} - L^R_{i,t} - w_{i,t}L^U_{i,t} \qquad \text{(operating profit)}$$

subject to

$$\begin{aligned}
\text{Demand function:} \quad & p_{i,t} = C_t^{1/\eta-\gamma}\theta_{i,t}^{1/\eta} q_{i,t}^{-1/\eta}, \\
\text{Supply function:} \quad & f_{i,t-1}(p_{i,t}, q_{i,t}) = 0, \\
\text{Urgent production:} \quad & L^U_{i,t} \geq \hat{q}_{i,t} - (1-\delta)I_{i,t-1}, \\
\text{Regular production:} \quad & L^R_{i,t} + L^U_{i,t} = \hat{q}_{i,t} + I_{i,t} - (1-\delta)I_{i,t-1}, \\
\text{Prior:} \quad & \log\theta_{i,t+1} \mid \log\theta_{i,t} \sim N(\mu^-_{i,t}, (\sigma^-)^2), \\
\text{Signal:} \quad & s_{i,t} = \log\theta_{i,t+1} + \varepsilon_{i,t}, \qquad \varepsilon_{i,t} \sim N(0, \sigma^2_{\varepsilon,i,t}), \\
& \sigma^2_{\varepsilon,i,t} = \frac{\sigma^2_{i,t}(\sigma^-)^2}{(\sigma^-)^2 - \sigma^2_{i,t}}, \quad 0 < \sigma_{i,t} < \sigma^-, \\
\text{Posterior:} \quad & \log\theta_{i,t+1} \mid \log\theta_{i,t}, s_{i,t} \sim N(\mu_{i,t}, \sigma^2_{i,t}), \quad 0 < \sigma_{i,t} \leq \sigma^-, \\
\text{Information cost:} \quad & \mathcal{I}(\sigma_{i,t}) \text{ is strictly decreasing and convex.}
\end{aligned}$$

Equation (1) defines the firm's continuation value net of information costs. A lower $\sigma_{i,t}$ means a more precise posterior and is therefore more costly. Operational profit is revenue net of regular and urgent labor costs. Whenever urgent labor is used in equilibrium, $w_{i,t} = \nu$. The constraints impose household demand, market clearing through the inherited supply function, inventory feasibility, and Gaussian updating from prior uncertainty σ^- to posterior uncertainty $\sigma_{i,t}$. Signal noise is independent of $\log\theta_{i,t+1}$ conditional on $\log\theta_{i,t}$. The no-information case $\sigma_{i,t} = \sigma^-$ is the limit as the signal-noise variance approaches infinity.

There are two equilibrium concepts. A $\boldsymbol{\sigma}$-equilibrium fixes the profile of posterior uncertainty $\boldsymbol{\sigma} = \{\sigma_{i,\tau}\}_{i\in[0,1],\,\tau\geq t}$. It describes all operating choices and market outcomes conditional on the information structure.

Definition 1. For a fixed profile $\boldsymbol{\sigma} = \{\sigma_{i,\tau}\}_{i\in[0,1],\,\tau\geq t}$, a recursive $\boldsymbol{\sigma}$-equilibrium consists of value functions $V_i(\cdot;\boldsymbol{\sigma})$, firm policies, prices, and laws of motion such that:

1. Given $x_{i,\tau} = (f_{i,\tau-1}, I_{i,\tau-1}, \theta_{i,\tau})$ and signal $s_{i,\tau}$, firm i chooses $f_{i,\tau}$, $I_{i,\tau}$, $L^R_{i,\tau}$, $L^U_{i,\tau}$ to maximize $V_i(x_{i,\tau};\boldsymbol{\sigma})$, taking the profile $\boldsymbol{\sigma}$ and its own information costs $\{\mathcal{I}(\sigma_{i,\tau})\}_{\tau\geq t}$ as given.

2. The household chooses variety demands and labor supplies optimally. Its demand for variety i satisfies

$$p_{i,\tau} = C_\tau^{\frac{1}{\eta}-\gamma}\theta_{i,\tau}^{1/\eta} q_{i,\tau}^{-1/\eta}.$$

3. Realized $(\hat{p}_{i,\tau}, \hat{q}_{i,\tau})$ are given by the intersection of demand and the supply function chosen in the

previous period.

4. Inventories and supply functions evolve according to firms' policy functions and realized shocks.

5. Beliefs evolve according to

$$\log \theta_{i,\tau+1} \mid \log \theta_{i,\tau}, s_{i,\tau} \sim N(\mu_{i,\tau}, \sigma^2_{i,\tau}).$$

Thus, a $\boldsymbol{\sigma}$-equilibrium determines operating outcomes while treating every firm's posterior uncertainty at every date as fixed.

The second concept is a decentralized equilibrium: a $\boldsymbol{\sigma}$-equilibrium in which each firm chooses its uncertainty profile to maximize its net continuation value. Firm i's discounted information cost is

$$\mathcal{K}_i(\boldsymbol{\sigma}) := \sum_{\tau \geq t} \beta^{\tau - t} \mathcal{I}(\sigma_{i,\tau}).$$

This expression aggregates the cost of posterior precision across all future dates.

Definition 2. A decentralized equilibrium is a $\boldsymbol{\sigma}$-equilibrium such that $\boldsymbol{\sigma}^p = \{\sigma^p_{i,\tau}\}_{i \in [0,1],\, \tau \geq t}$ satisfies, for every firm i and date τ with an interior uncertainty choice,

$$\frac{\partial V_i(x_{i,\tau}; \boldsymbol{\sigma}^p)}{\partial \sigma_{i,\tau}} = 0 \quad \Longleftrightarrow \quad \frac{\partial \Pi_i^{0,*}(x_{i,\tau}; \boldsymbol{\sigma}^p)}{\partial \sigma_{i,\tau}} = \frac{\partial \mathcal{K}_i(\boldsymbol{\sigma}^p)}{\partial \sigma_{i,\tau}}.$$

The equality balances the private marginal effect of uncertainty on operating profit against its marginal information-cost saving. These are private derivatives: because each firm is atomistic, it varies its own uncertainty while holding aggregate states fixed.

C. Welfare accounting

In this framework, the representative household's per-period utility equals consumer surplus plus firms' operating profits. Because urgent labor is paid its equilibrium marginal cost ν, firm i's flow operating profit is $\pi^0_{i,t} := \hat{p}_{i,t}\hat{q}_{i,t} - L^R_{i,t} - \nu L^U_{i,t}$. Define flow consumer surplus and total flow utility as

$$\begin{aligned} CS^{flow}_t &:= u\left(\int C(\hat{q}_{i,t}, \theta_{i,t}) di\right) - \int \hat{p}_{i,t}\hat{q}_{i,t} di, \\ U_t &:= u\left(\int C(\hat{q}_{i,t}, \theta_{i,t}) di\right) - \int \left(L^R_{i,t} + \nu L^U_{i,t}\right) di. \end{aligned}$$

Consumer surplus subtracts expenditure from consumption utility, whereas U_t subtracts the real labor costs of producing consumption. Adding and subtracting firm revenues gives the flow accounting identity

$$U_t \equiv CS^{flow}_t + \int \pi^0_{i,t} di.$$

The identity shows that expenditure is a transfer from the household to firms and therefore cancels in aggregate welfare. At the operating optimum induced by $\boldsymbol{\sigma}$, define $\Pi_i^{0,*}(x_{i,t};\boldsymbol{\sigma}) := \mathbb{E}_t \sum_{\tau\geq t} \beta^{\tau-t}\pi_{i,\tau}^{0,*}$. Then the value function in Problem (1) is operating profit under the optimal inventory and supply policies net of information costs:

$$V_i(x_{i,t};\boldsymbol{\sigma}) = \Pi_i^{0,*}(x_{i,t};\boldsymbol{\sigma}) - \mathcal{K}_i(\boldsymbol{\sigma}).$$

Thus, $\Pi_i^{0,*}$ is the gross operating value and V_i is the corresponding net private value. Let $x_t := \{x_{i,t}\}_{i\in[0,1]}$ denote the collection of individual firm states at date t. Define the present value of consumer-surplus flows as $CS^*(x_t;\boldsymbol{\sigma}) := \mathbb{E}_t \sum_{\tau\geq t} \beta^{\tau-t} CS_\tau^{flow,*}$. Lifetime welfare is then

$$W^*(x_t;\boldsymbol{\sigma}) \equiv CS^*(x_t;\boldsymbol{\sigma}) + \int V_i(x_{i,t};\boldsymbol{\sigma})\, di. \tag{2}$$

Identity (2) states that lifetime welfare net of information costs is the sum of consumer surplus and firms' net continuation values.

III. Optimal supply function and inventory holdings

In this section, I characterize the optimal inventory level and supply function. I first derive the optimal supply function. Because demand is subject to a single shock and the demand function is invertible, the firm can choose a state-contingent quantity and then translate it into a supply function. This yields the optimal supply function without imposing its shape in advance.

Problem (1) is an infinite-horizon problem written recursively. However, conditional on the posterior distribution, its operating decisions can be represented as a sequence of static problems.

Lemma 1. *Let $\boldsymbol{\sigma}_t := \{\sigma_{i,t}\}_{i\in[0,1]}$ denote the cross-sectional profile of posterior uncertainties at date t. The solution to problem (1) can be represented as a series of static choices such that firm i in period t solves*

$$V_{i,t}^0(\boldsymbol{\sigma}_t) = \mathbb{E}\left[\max_{f_{i,t}(\cdot),\, J_{i,t}} \mathbb{E}\left[(\hat{p}_{i,t+1}-1)\hat{q}_{i,t+1} - (\nu-1)(\hat{q}_{i,t+1}-J_{i,t})^+ \,\middle|\, \theta_{i,t}, s_{i,t}\right] - \left(\frac{1}{(1-\delta)\beta} - 1\right) J_{i,t} \,\middle|\, \theta_{i,t}\right],$$

subject to

$$\log\theta_{i,t+1} \mid \log\theta_{i,t}, s_{i,t} \sim N(\mu_{i,t}, \sigma_{i,t}^2), \quad f_{i,t}(\hat{p}_{i,t+1}, \hat{q}_{i,t+1}) = 0, \quad \hat{p}_{i,t+1} = C_{t+1}^{1/\eta-\gamma}\theta_{i,t+1}^{1/\eta}\hat{q}_{i,t+1}^{-1/\eta},$$

where $J_{i,t} = (1-\delta)I_{i,t}$ is after-depreciation inventory level.

Lemma 1 states that, after observing the signal, current operating decisions maximize expected next-period operating profit net of the carrying cost of inventories; $V_{i,t}^0$ averages this optimized value over signals. The effective carrying cost of one unit of after-depreciation inventory is $\lambda := 1/[(1-\delta)\beta] - 1$: it is higher when more inventory depreciates and lower when the firm discounts the future less. Expected operating profit equals revenue net of regular production cost and the urgency premium $\nu - 1$ paid on sales above the after-depreciation inventory level. The constraints specify the posterior distribution, the supply schedule, and the inverse demand curve faced by the firm. The proof is provided in Appendix A.1.

The next lemma characterizes the optimal quantity rule that maps the demand shifter into quantity for a given after-depreciation inventory level.

Lemma 2. *For fixed $J_{i,t} = (1-\delta)I_{i,t}$, the optimal next-period quantity rule is*

$$Q^*(A_{i,t+1}; J_{i,t}) = \begin{cases} \dfrac{A_{i,t+1}}{\mathcal{M}^\eta}, & A_{i,t+1} < J_{i,t}\mathcal{M}^\eta, \\ J_{i,t}, & J_{i,t}\mathcal{M}^\eta \le A_{i,t+1} \le J_{i,t}(\mathcal{M}\nu)^\eta, \\ \dfrac{A_{i,t+1}}{(\mathcal{M}\nu)^\eta}, & A_{i,t+1} > J_{i,t}(\mathcal{M}\nu)^\eta, \end{cases} \qquad A_{i,t+1} := C_{t+1}^{1-\eta\gamma}\theta_{i,t+1}.$$

Here $A_{i,t+1}$ is the next-period demand shifter in levels: inverse demand can be written as $p_{i,t+1} = A_{i,t+1}^{1/\eta} q_{i,t+1}^{-1/\eta}$. Lemma 2 divides the optimal quantity rule into three regions. At low demand, the firm serves the monopoly quantity from inventories and prices at a markup over regular marginal cost. At high demand, it uses urgent production and prices at a markup over urgent marginal cost. Between these regions, the firm sells exactly its available inventories and does not produce urgently. This quantity rule maps directly into a supply function. The proof is provided in Appendix A.2.

Proposition 1. *The optimal next-period supply function, given after-depreciation inventory $J_{i,t} = (1-\delta)I_{i,t}$, is a two-branch constant-markup function with a vertical inventory segment:*

$$p_{i,t+1} = \begin{cases} \mathcal{M}, & q_{i,t+1} < J_{i,t}, \\ \textit{vertical segment at } q_{i,t+1} = J_{i,t}, & p_{i,t+1} \in [\mathcal{M}, \mathcal{M}\nu], \\ \mathcal{M}\nu, & q_{i,t+1} > J_{i,t}. \end{cases}$$

Proposition 1 states that the supply function selected at date t has two horizontal constant-markup branches and a vertical segment at $q_{i,t+1} = J_{i,t}$. Along the vertical segment, price adjusts to clear the market while quantity remains fixed at available inventories. Outside this segment, the firm charges the constant markup $\mathcal{M}$ over the relevant marginal cost: one under regular production and ν under urgent production. Figure 1 illustrates these three regions. The proof is provided in Appendix A.3.

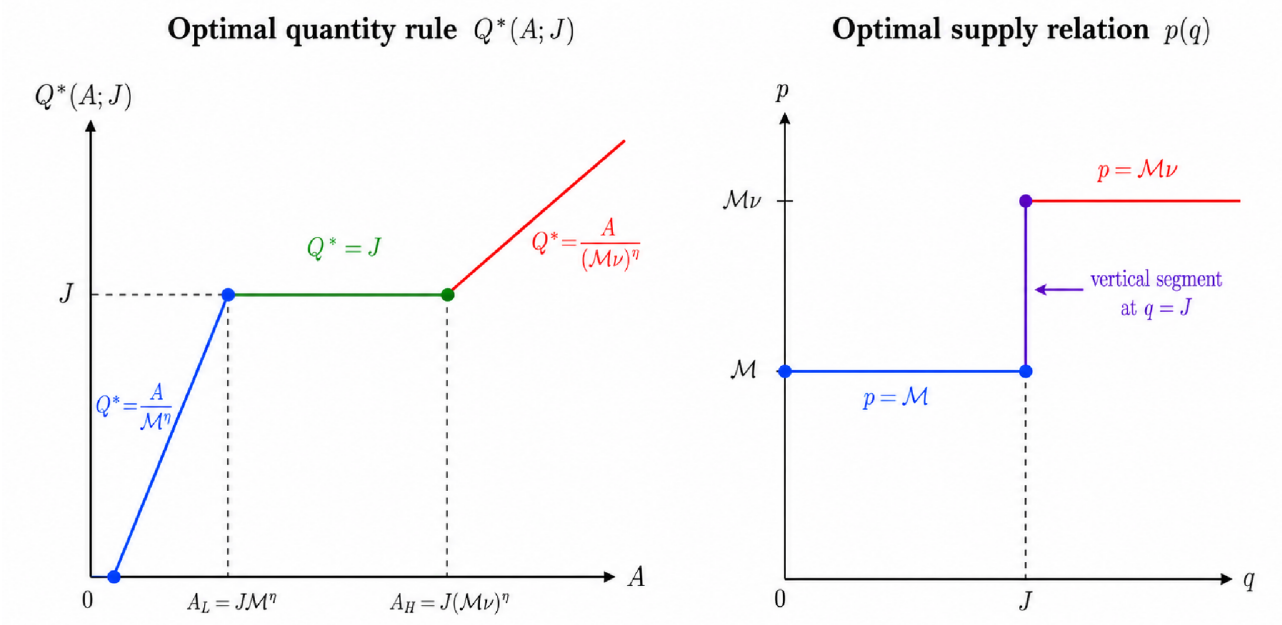


Figure 1: Optimal supply function for a given inventory position

I next characterize the optimal inventory choice. An additional unit is valuable when the firm would otherwise remain on the vertical segment or use urgent production. The firm equates this expected marginal benefit to the carrying cost generated by discounting and depreciation. Let $\mu^A_{i,t}$ denote the posterior mean of the log demand shifter, $\log A_{i,t+1} \mid \log \theta_{i,t}, s_{i,t} \sim N(\mu^A_{i,t}, \sigma^2_{i,t})$.

Optimal after-depreciation inventory holdings are

$$J^*_{i,t} = \frac{\exp(\mu^A_{i,t} + \sigma_{i,t} a^*(\sigma_{i,t}))}{\mathcal{M}^\eta}. \tag{3}$$

The factor $\exp(\mu^A_{i,t})$ scales inventories with expected demand, while $\sigma_{i,t} a^*(\sigma_{i,t})$ determines the lower boundary of the inaction region. The cutoff $a^*(\sigma_{i,t})$ solves $G_{\sigma_{i,t}}(a,b) = \lambda$[1]. This equation equates the expected marginal benefit of inventories across the inaction and urgent-production regions to their carrying cost. If $0 < \lambda < \nu - 1$, so that carrying inventories is costly but cheaper than urgent production, there is a unique interior cutoff $a^*(\sigma_{i,t})$ and hence a unique interior inventory choice $J^*_{i,t}$. The upper cutoff is $b^*(\sigma_{i,t}) = a^*(\sigma_{i,t}) + \eta \log(\nu)/\sigma_{i,t}$; the distance between the two cutoffs reflects the gap between regular and urgent marginal costs. Appendix A.4 derives the inventory first-order condition, uniqueness, and the uncertainty comparative static.

Equation 3 yields several comparative statics. First, inventory is strictly decreasing in the markup $\mathcal{M}$. A

[1]The fixed point equation is given by

$$G_{\sigma_{i,t}}(a,b) := \exp\left(-\frac{\sigma_{i,t} a}{\eta} + \frac{1}{2}\left(\frac{\sigma_{i,t}}{\eta}\right)^2\right)\left[\Phi\left(b - \frac{\sigma_{i,t}}{\eta}\right) - \Phi\left(a - \frac{\sigma_{i,t}}{\eta}\right)\right] - [\Phi(b) - \Phi(a)] + (\nu - 1)\,[1 - \Phi(b)] = \lambda.$$

higher markup reduces the profit-maximizing quantity associated with any demand realization and therefore lowers desired inventories.

Second, holding the posterior mean fixed, the response of optimal inventories to uncertainty is

$$\frac{\partial J^*_{i,t}}{\partial \sigma_{i,t}} = J^*_{i,t} \times \mathbb{E}[X | a^*(\sigma_{i,t}) \leq X \leq b^*(\sigma_{i,t})], \qquad X \sim N\left(\frac{\sigma_{i,t}}{\eta}, 1\right).$$

Here X is not an additional demand shock. Start from the standardized residual demand shock $Z = (\log A_{i,t+1} - \mu^A_{i,t})/\sigma_{i,t} \sim N(0,1)$. In the inventory derivative, each realization of Z is weighted by the marginal value of inventories along the vertical segment. This exponential tilting shifts the normal distribution's mean from zero to $\sigma_{i,t}/\eta$, which gives the variable X above. Equivalently, its conditional mean can be written as

$$\mathbb{E}[X | a \leq X \leq b] := \frac{\int_a^b e^{\tau_{i,t}(z-a)} z \phi(z) \mathrm{d}z}{\int_a^b e^{\tau_{i,t}(z-a)} \phi(z) \mathrm{d}z}, \qquad \tau_{i,t} := \frac{\sigma_{i,t}}{\eta}.$$

The exponential weight $e^{\tau_{i,t}(z-a)} = p/\mathcal{M}$ is the marginal value of inventories along the vertical segment. Hence, demand realizations for which inventories are more valuable receive greater weight in the comparative static. The variable X summarizes this weighting in a single sign statistic. If its truncated mean is positive, the marginal-value-weighted inaction states lie above the posterior mean, so greater uncertainty raises desired inventories. If it is negative, the weighted states lie below the posterior mean and greater uncertainty lowers desired inventories.

Optimal inventory $J^*_{i,t}$ is also strictly increasing in $\mu^A_{i,t}$: higher expected demand raises desired inventories. Under Gaussian updating and holding the signal realization and aggregate state fixed, a change in $\sigma_{i,t}$ also changes $\mu^A_{i,t}$ by changing the posterior weight placed on that signal. Let $\mu^{A,-}_{i,t}$ denote the corresponding prior mean of $\log A_{i,t+1}$. The total response is therefore

$$\frac{dJ^*_{i,t}}{d\sigma_{i,t}} = J^*_{i,t} \times \left(\mathbb{E}[X | a^*(\sigma_{i,t}) \leq X \leq b^*(\sigma_{i,t})] - \frac{2\sigma_{i,t}}{(\sigma^-)^2 - \sigma^2_{i,t}} (\mu^A_{i,t} - \mu^{A,-}_{i,t}) \right).$$

The first term is the direct effect of posterior uncertainty at a fixed posterior mean. The second is the indirect effect operating through the posterior mean; its sign depends on whether the signal raises or lowers expected demand relative to the prior.

Before the signal is realized, the corresponding response of expected inventories is

$$\frac{d\mathbb{E}[J^*_{i,t}]}{d\sigma_{i,t}} = \mathbb{E}[J^*_{i,t}] \Big(\mathbb{E}[X | a^*(\sigma_{i,t}) \leq X \leq b^*(\sigma_{i,t})] - \sigma_{i,t} \Big).$$

The first term again captures the operating response to uncertainty, while $-\sigma_{i,t}$ captures the ex ante response of the distribution of posterior means. Expected inventories increase if the former effect dominates the latter. Appendix A.4 derives both the signal-contingent and ex ante comparative statics.

Substituting the optimal inventory choice from Equation 3 into Proposition 1 gives the complete optimal

supply rule.

Corollary 1. *The optimal supply function is*

$$p_{i,t+1} = \begin{cases} \mathcal{M}, & q_{i,t+1} < J^*_{i,t}, \\ \text{vertical segment at } q_{i,t+1} = J^*_{i,t}, \quad p_{i,t+1} \in [\mathcal{M}, \mathcal{M}\nu], & \\ \mathcal{M}\nu, & q_{i,t+1} > J^*_{i,t}. \end{cases},$$

where $J^*_{i,t}$ is given by (3).

The corollary states that information affects the location of the vertical segment through $J^*_{i,t}$, while the two markup branches remain at $\mathcal{M}$ and $\mathcal{M}\nu$.

IV. Information acquisition and welfare

Define $\bar{V}^0_{i,t}$ as the operating problem in Lemma 1 with the common demand shifter normalized to one, $C^{1/\eta-\gamma}_{t+1} \equiv 1$, evaluated before the date-t signal is realized. I refer to the effect of uncertainty on $\bar{V}^0_{i,t}$ as the partial-equilibrium effect because it holds aggregate consumption fixed. Since $V^0_{i,t}$ is homogeneous of degree η in $C^{1/\eta-\gamma}_{t+1}$,

$$V^0_{i,t} = \left(C^{1/\eta-\gamma}_{t+1}\right)^{\eta} \bar{V}^0_{i,t}.$$

Define next-period revenue as $R_{i,t+1} := \hat{p}_{i,t+1}\hat{q}_{i,t+1}$ and let $\bar{R}_{i,t+1}$ denote the corresponding revenue in the normalized problem. Aggregate actual and normalized expected revenues are

$$R_{t+1} := \int \mathbb{E}_t[R_{i,t+1}]di, \qquad \bar{R}_{t+1} := \int \mathbb{E}_t[\bar{R}_{i,t+1}]di.$$

Lemma 3. *In a $\boldsymbol{\sigma}$-equilibrium,*

$$\bar{V}^0_{i,t}(\sigma_{i,t}) = \frac{1}{\eta}\mathbb{E}_t[\bar{R}_{i,t+1}], \qquad \frac{\partial \bar{V}^0_{i,t}(\sigma_{i,t})}{\partial \sigma_{i,t}} \leq 0.$$

Lemma 3 states that partial-equilibrium operating profit is proportional to normalized expected revenue. Greater posterior uncertainty weakly lowers normalized expected revenue and therefore the firm's partial-equilibrium operating value; the effect is strict when the vertical inventory region has positive probability. Because GE profit and GE revenue both scale by the same common demand shifter, the corresponding GE relation is also $V^0_{i,t} = (1/\eta)\mathbb{E}_t[R_{i,t+1}]$. The proof is provided in Appendix A.5.

Lemma 4. *Let*

$$\bar{V}^0_t(\boldsymbol{\sigma}_t) := \int \bar{V}^0_{i,t}(\sigma_{i,t})di.$$

Then, in a $\boldsymbol{\sigma}$-equilibrium,

$$V^0_{i,t}(\boldsymbol{\sigma}_t) = \left(\eta \bar{V}^0_t(\boldsymbol{\sigma}_t)\right)^{\frac{1-\eta\gamma}{\gamma(\eta-1)}} \bar{V}^0_{i,t}(\boldsymbol{\sigma}_t).$$

A decentralized equilibrium is a $\boldsymbol{\sigma}$-equilibrium such that, for every firm i and date t with an interior uncertainty choice,

$$\beta\left(\eta\bar{V}_t^0(\boldsymbol{\sigma}_t^p)\right)^{\frac{1-\eta\gamma}{\gamma(\eta-1)}}\frac{\partial\bar{V}_{i,t}^0(\boldsymbol{\sigma}_t^p)}{\partial\sigma_{i,t}}=\frac{\partial\mathcal{K}_i(\boldsymbol{\sigma}^p)}{\partial\sigma_{i,t}}.$$

Lemma 4 separates normalized operating profit from the common demand shifter. An atomistic firm accounts for the direct effect of its own uncertainty but treats the aggregate object $\bar{V}_t^0(\boldsymbol{\sigma}_t)$ as fixed. The proof is provided in Appendix A.6.

To allow for ex ante firm heterogeneity, suppose there are G firm types, indexed by g, with population masses $m_g > 0$ satisfying $\sum_{g=1}^{G} m_g = 1$. A type collects firms with the same model parameters, so heterogeneity across types can be generated by differences in parameters such as depreciation. The symmetric economy is the special case $G = 1$.

The general-equilibrium effect of type g's uncertainty on firm i is

$$\frac{\partial V_{i,t}^0(\boldsymbol{\sigma}_t)}{\partial\sigma_{g,t}}=V_{i,t}^0(\boldsymbol{\sigma}_t)\left[\mathbf{1}\{i\in g\}\frac{\partial_{\sigma_{i,t}}\bar{V}_{i,t}^0(\boldsymbol{\sigma}_t)}{\bar{V}_{i,t}^0(\boldsymbol{\sigma}_t)}+\frac{1-\eta\gamma}{\gamma(\eta-1)}\frac{\partial_{\sigma_{g,t}}\bar{V}_t^0(\boldsymbol{\sigma}_t)}{\bar{V}_t^0(\boldsymbol{\sigma}_t)}\right].$$

The first term is the direct response of normalized operating profit. The second captures the response of aggregate consumption, marginal utility, and the common demand shifter.

Let $V_t^0 := \int V_{i,t}^0 di$. For a type-wide change in uncertainty,

$$\frac{\partial\bar{V}_t^0}{\partial\sigma_{g,t}}=m_g\frac{\partial\bar{V}_{g,t}^0}{\partial\sigma_{g,t}}\leq 0.$$

Integrating the firm-level derivative across firms gives

$$\begin{aligned}\frac{\partial V_t^0}{\partial\sigma_{g,t}}&=\left(\eta\bar{V}_t^0\right)^{\frac{1-\eta\gamma}{\gamma(\eta-1)}}\frac{\partial\bar{V}_t^0}{\partial\sigma_{g,t}}+\frac{1-\eta\gamma}{\gamma(\eta-1)}\frac{\partial_{\sigma_{g,t}}\bar{V}_t^0}{\bar{V}_t^0}V_t^0\\&=\frac{V_t^0}{\bar{V}_t^0}\left[1+\frac{1-\eta\gamma}{\gamma(\eta-1)}\right]\frac{\partial\bar{V}_t^0}{\partial\sigma_{g,t}}\\&=\frac{V_t^0}{\bar{V}_t^0}\frac{1-\gamma}{\gamma(\eta-1)}\frac{\partial\bar{V}_t^0}{\partial\sigma_{g,t}}.\end{aligned}$$

If $\partial_{\sigma_{g,t}}\bar{V}_t^0 < 0$, aggregate operating profit rises with uncertainty when $\gamma > 1$, falls when $0 < \gamma < 1$, and is locally unchanged when $\gamma = 1$. For $\gamma > 1$, the increase in the common demand shifter more than offsets the loss in normalized operating profit. In that case,

$$\frac{\partial}{\partial\sigma_{g,t}}\int\Pi_i^{0,*}(x_{i,t};\boldsymbol{\sigma})di=\beta\frac{\partial V_t^0(\boldsymbol{\sigma}_t)}{\partial\sigma_{g,t}}>0.$$

Proposition 2. *At a decentralized equilibrium with an interior uncertainty choice for type g, the effects*

of type g's uncertainty on consumer surplus and welfare are, for every $\gamma > 0$ (with $\gamma = 1$ understood by continuity),

$$\frac{\partial CS^*(x_t;\boldsymbol{\sigma}^p)}{\partial \sigma_{g,t}} = \beta \frac{\eta}{\eta-1} \frac{V_t^0(\boldsymbol{\sigma}_t^p)}{\bar{V}_t^0(\boldsymbol{\sigma}_t^p)} \frac{\partial \bar{V}_t^0(\boldsymbol{\sigma}_t^p)}{\partial \sigma_{g,t}} < 0,$$

$$\frac{\partial W^*(x_t;\boldsymbol{\sigma}^p)}{\partial \sigma_{g,t}} = \beta \frac{1}{\gamma(\eta-1)} \frac{V_t^0(\boldsymbol{\sigma}_t^p)}{\bar{V}_t^0(\boldsymbol{\sigma}_t^p)} \frac{\partial \bar{V}_t^0(\boldsymbol{\sigma}_t^p)}{\partial \sigma_{g,t}} < 0.$$

Proposition 2 states that, at an interior decentralized information choice, a type-wide increase in uncertainty strictly lowers consumer surplus and welfare. When $\gamma > 1$, the consumer-surplus loss dominates the positive general-equilibrium profit spillover that atomistic firms do not internalize; when $0 < \gamma \leq 1$, the aggregate-profit response is nonpositive. In either case, firms under-acquire information relative to the welfare optimum. These equalities use the private first-order condition and need not hold away from the decentralized equilibrium. The proof is provided in Appendix A.7.

Implication: urgency costs and welfare

One implication of information under-acquisition is that a higher urgency cost tomorrow can raise welfare evaluated at the decentralized information choice, if the induced increase in information acquisition is sufficiently large, even though its direct effect holding information fixed is negative. The timing is important. Information chosen at date t concerns demand at date $t+1$, so it responds to ν_{t+1}, not to the current urgency cost ν_t. In particular, $\partial \sigma_{g,t}^p / \partial \nu_t = 0$.

Consider a one-time change in ν_{t+1}, holding urgency costs at all other dates fixed. Let $\boldsymbol{\sigma}^p(\nu_{t+1})$ denote the decentralized uncertainty profile. Only its date-t component responds to this change; all other components are held fixed. Along the decentralized-equilibrium path,

$$\frac{d}{d\nu_{t+1}} W^*(x_t;\boldsymbol{\sigma}^p(\nu_{t+1}),\nu_{t+1}) = \underbrace{\left.\frac{\partial W^*(x_t;\boldsymbol{\sigma},\nu_{t+1})}{\partial \nu_{t+1}}\right|_{\boldsymbol{\sigma}=\boldsymbol{\sigma}^p(\nu_{t+1})}}_{\text{direct effect at } t+1<0} + \underbrace{\sum_{g=1}^{G} \frac{\partial W^*(x_t;\boldsymbol{\sigma}^p(\nu_{t+1}),\nu_{t+1})}{\partial \sigma_{g,t}} \frac{d\sigma_{g,t}^p}{d\nu_{t+1}}}_{\text{induced date-}t\text{ information effect}}.$$

To sign the direct effect, define

$$\bar{V}_t^0(\boldsymbol{\sigma}_t,\nu_{t+1}) = \sum_{g=1}^{G} m_g \bar{V}_{g,t}^0(\sigma_{g,t},\nu_{t+1}), \qquad \rho := \frac{1-\gamma}{\gamma(\eta-1)}.$$

Holding information fixed, a higher next-period urgency cost lowers normalized operating-profit capacity:

$$\frac{\partial \bar{V}_t^0}{\partial \nu_{t+1}} = \sum_{g=1}^{G} m_g \frac{\partial \bar{V}_{g,t}^0}{\partial \nu_{t+1}} < 0.$$

This derivative contains the direct resource cost of urgent production at $t+1$. Using Lemma 4,

$$\frac{\partial V_t^0}{\partial \nu_{t+1}} = \rho \frac{V_t^0}{\bar{V}_t^0} \frac{\partial \bar{V}_t^0}{\partial \nu_{t+1}}.$$

Thus aggregate operating profit rises with ν_{t+1} when $\gamma > 1$, falls when $0 < \gamma < 1$, and is locally unchanged when $\gamma = 1$. Nevertheless, the direct effect on welfare is negative for every $\gamma > 0$:

$$\left.\frac{\partial W^*(x_t; \boldsymbol{\sigma}, \nu_{t+1})}{\partial \nu_{t+1}}\right|_{\boldsymbol{\sigma}=\boldsymbol{\sigma}^p(\nu_{t+1})} = \beta \frac{1+\gamma(\eta-1)}{\gamma(\eta-1)} \frac{V_t^0}{\bar{V}_t^0} \frac{\partial \bar{V}_t^0}{\partial \nu_{t+1}} < 0.$$

Under the interior and strictness conditions of Proposition 2, $\partial W^*/\partial \sigma_{g,t} < 0$. Consequently, if a higher urgency cost tomorrow raises date-t information acquisition, so that $d\sigma_{g,t}^p/d\nu_{t+1} \leq 0$ for every type, the induced information effect is nonnegative. Welfare rises with ν_{t+1} precisely when this effect dominates the negative direct effect at $t+1$. A change in ν_t has only its direct operating effect and cannot generate this information channel: ν_t applies to current production determined by choices made at $t-1$, whereas $\boldsymbol{\sigma}_t$ concerns demand at $t+1$.

V. Empirical evidence and welfare assessment

This section presents empirical evidence on the relationships between information acquisition, inventory holdings, and costly urgent production. The evidence supports the model's main mechanisms and suggests an important role for heterogeneity in inventory depreciation. I then use the calibrated model to assess the welfare loss from decentralized information acquisition.

A. Data: management guidance and quarterly fundamentals

The main dataset is I/B/E/S Guidance, which contains management sales guidance for U.S. public firms and serves as a proxy for firm expectations. The sample contains 1,219 firms and 13,108 firm-quarters from 2018 to 2024. The median horizon from announcement to quarter end is 57 days.

I combine the guidance data with Compustat quarterly data for the same firms. Compustat provides sales, cost of goods sold (COGS), inventories, assets, debt, capital expenditures (CAPEX), and research and development (R&D) expenses. I also construct industry controls using NAICS, SIC, and GICS classifications.

The main empirical objects are the sales-guidance forecast error and its firm-level variance, which proxy for demand surprises and forecast imprecision:

$$FE_{i,t} := \text{realized sales}_{i,t} - \text{guided sales}_{i,t},$$
$$\text{Var}_i(FE) := \text{variance of } FE_{i,t} \text{ across time.}$$

Positive $FE_{i,t}$ indicates that the firm underestimated demand. Higher $\text{Var}_i(FE)$ indicates less precise guidance and greater demand uncertainty.

The main outcomes are inventory holdings and COGS. Controls include sales-to-assets, lagged financial positions, balance-sheet ratios, guidance characteristics, and firm, time, and industry fixed effects where appropriate.

B. Empirical evidence: information, inventories, and costly adjustment

Table 1 summarizes the main empirical patterns.

Table 1: Information, inventories, and costly adjustment

Panel A. Information and inventory holdings

Object	Regressor	Outcome	Coef.	Sample
Expected sales	log guided sales$_{t+1}$	$\log I_t$	0.214***	4,097
			(0.060)	
Demand surprise	FE_t	$\Delta \log I_t$	0.159**	4,400
			(0.063)	
Forecast imprecision	$\log \mathrm{Var}(FE)$	Avg. $\log I$	0.029*	587 firms
			(0.017)	

Panel B. Costly urgent-production margin

Object	Regressor	Outcome	Coef.	Sample
Underestimated demand	$\max(FE_t, 0)$	$\Delta(COGS/Sales)_t$	0.044*	6,792
			(0.023)	
Overestimated demand	$\max(-FE_t, 0)$	$\Delta(COGS/Sales)_t$	-0.028	6,792
			(0.021)	

Notes: Standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Panel A shows that inventories increase with expected sales and positive demand surprises, while firms with less precise guidance hold more inventories on average. Panel B shows that underestimating demand raises COGS relative to sales, consistent with costly urgent production; overestimating demand has no statistically significant effect. These conditional correlations support the model's mechanisms but are not interpreted as causal estimates.

The inventory–information relationship also varies with depreciation. Table 2 reports correlations between log forecast-error variance and average log inventories. The correlation is positive in low-depreciation industries and negative in middle- and high-depreciation industries under both industry classifications. Thus inventories appear to substitute for information when depreciation is low and to complement information when depreciation is higher. Appendix B defines effective inventory depreciation and reports both industry mappings.

Table 2: Information, inventories, and industry depreciation

Inventory–information correlations by depreciation		
Inventory depreciation	Broad NAICS2	Detailed NAICS3
Low	0.185***	0.223***
Middle or high	-0.076*	-0.081*

Notes: Entries are firm-level correlations between $\log \mathrm{Var}(FE)$ and average $\log I$. The broad-classification samples contain 399 low- and 525 middle/high-depreciation firms; the detailed-classification samples contain 420 and 504 firms, respectively. $^{*}p < 0.10$, $^{***}p < 0.01$.

The sign reversal and moderate correlation magnitudes support depreciation heterogeneity in the model. Heterogeneity only in urgency costs generates a correlation of -1 between inventories and posterior uncertainty, while heterogeneity only in prior uncertainty generates a correlation of 1. Depreciation heterogeneity instead allows the correlation to be weaker and to change sign because information is least valuable when inventories are either very cheap or too costly to use actively.

C. Numerical results

To define the welfare benchmark with G firm types, let $m_g > 0$ denote type g's mass, with $\sum_{g=1}^{G} m_g = 1$, and let $\boldsymbol{\sigma}_t = (\sigma_{1,t}, \ldots, \sigma_{G,t})$. For each candidate $\boldsymbol{\sigma}_t$, allocations and type-specific operating profits are those of the corresponding $\boldsymbol{\sigma}$-equilibrium. The welfare-maximizing uncertainty profile solves

$$\widetilde{\boldsymbol{\sigma}}_t \in \arg\max_{\boldsymbol{\sigma}_t} \left\{ \sum_{g=1}^{G} m_g \left(\beta \frac{1 + \gamma(\eta - 1)}{1 - \gamma} V_{g,t}^{0}(\boldsymbol{\sigma}_t) - \mathcal{I}(\sigma_{g,t}) \right) \right\},$$

where $V_{g,t}^{0}(\boldsymbol{\sigma}_t)$ is general-equilibrium operating profit of type g in the corresponding $\boldsymbol{\sigma}$-equilibrium. The factor β discounts the date-$t+1$ operating allocation generated by date-t information.

I use the calibrated model to quantify the gap between private and welfare-optimal information acquisition in a heterogeneous-firm economy. The benchmark sets $\eta = 8.0$, $\nu = 1.5$, $\mu = 0$, prior variance $\sigma_0^2 = 0.26$, information-cost parameter $\kappa = 5.0 \times 10^{-3}$, $\beta = 0.99$, and $\gamma = 4.5$. The economy contains 61 equally weighted depreciation types with $\delta_g \in [0, 0.30]$. Each type chooses information using its own operating problem while treating the aggregate state as given. The planner instead chooses the complete vector $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_G)$ jointly.

Figure 2 reports type-specific choices. The left panel plots informativeness, $K_g^* = 1 - (\sigma_g^*)^2/\sigma_0^2$, and the right panel plots expected pre-depreciation inventories. The private information policy is nonmonotone: informativeness initially increases with depreciation, peaks at an intermediate rate, and then falls to zero. The welfare optimum selects more information for every type, changing type-specific inventories.

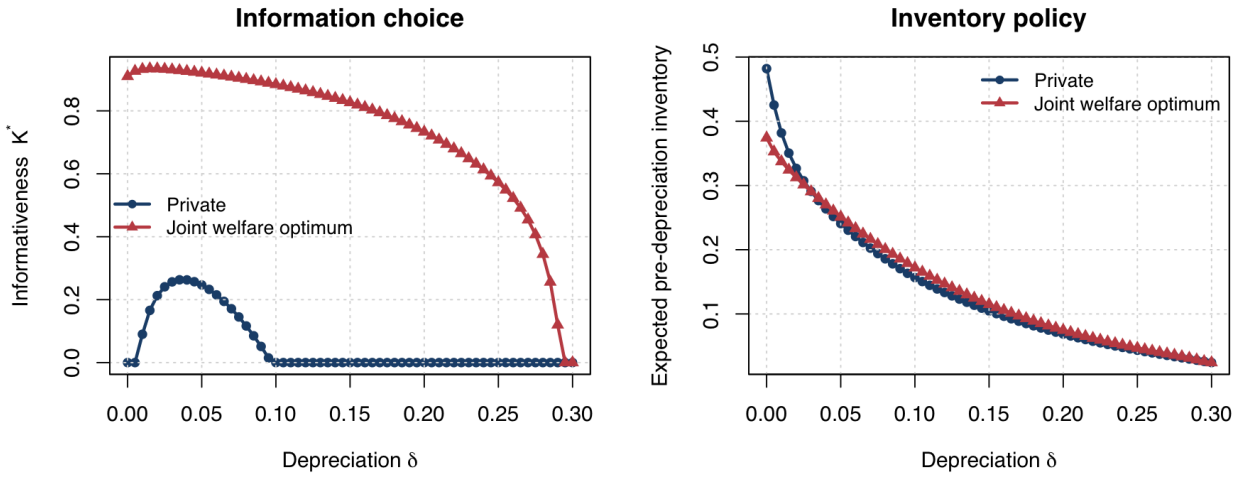


Figure 2: Information and inventory choices with heterogeneous depreciation

Figure 3 shows signal-contingent inventory rules for two types. At $\delta = 0.05$, both the private firm and the planner condition inventories on the posterior signal, but the planner chooses a more informative signal and therefore a wider inventory response. At $\delta = 0.10$, the private firm chooses no information, so the posterior mean is degenerate and private inventory is flat across signal percentiles. The welfare optimum still acquires information for this type, producing a state-contingent inventory policy.

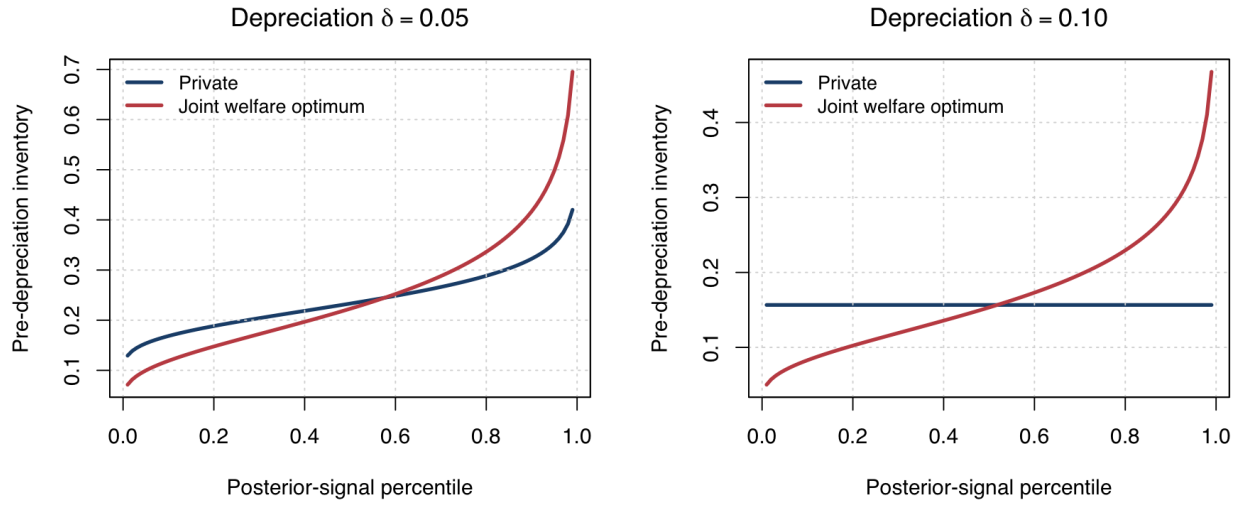


Figure 3: Signal-contingent inventory policies for selected depreciation types

Finally, the aggregate relative welfare improvement from replacing the private information profile with the welfare-optimal profile is

$$100\frac{W^*(x_t;\widetilde{\boldsymbol{\sigma}}_t) - W^*(x_t;\boldsymbol{\sigma}_t^p)}{|W^*(x_t;\boldsymbol{\sigma}_t^p)|}.$$

In the benchmark calibration, this relative welfare improvement is 0.56 percent.

VI. Discussion

A. *Policy implications: sales versus production subsidies*

Suppose subsidies are financed with lump-sum taxes. Aggregate taxes under the sales and production subsidies are, respectively,

$$T_t^S = (\mathcal{M}-1)\int R_{i,t}di, \qquad T_t^P = \left(1-\frac{1}{\mathcal{M}}\right)\int \left(L_{i,t}^R + \nu L_{i,t}^U\right)di.$$

In either case, with T_t denoting the corresponding aggregate tax, household surplus net of taxes plus subsidized firm profits is

$$u_\gamma(C_t) - \int R_{i,t}di - T_t + \int \pi_{i,t}^{0,\text{sub}}di = u_\gamma(C_t) - \int \left(L_{i,t}^R + \nu L_{i,t}^U\right)di.$$

Thus, subsidy payments and lump-sum taxes cancel in welfare; the policies differ only through firms' private payoffs. Under the sales subsidy, the government pays firm i an additional $(\mathcal{M}-1)R_{i,t}$, so its total sales receipts are $\mathcal{M}R_{i,t}$ and its operating payoff is

$$\mathcal{M}R_{i,t} - L_{i,t}^R - \nu L_{i,t}^U.$$

Multiplying revenue by $\mathcal{M}$ makes the firm's marginal revenue equal to the consumer price and therefore eliminates the markup. In this case, Lemma 3 becomes

$$\bar{V}_{i,t}^0 = \frac{1}{\eta-1}\mathbb{E}_t[\bar{R}_{i,t+1}].$$

The corresponding GE relation is $V_{i,t}^0 = (1/(\eta-1))\mathbb{E}_t[R_{i,t+1}]$ because both sides scale by the common demand shifter. Accordingly, $\eta\bar{V}_t^0$ in Lemma 4 is replaced by $(\eta-1)\bar{V}_t^0$. The changed profit–revenue coefficient aligns the private and planner information first-order conditions, so at the private interior optimum Proposition 2 becomes

$$\left.\frac{\partial W^*(x_t;\boldsymbol{\sigma})}{\partial \sigma_{g,t}}\right|_{\boldsymbol{\sigma}^P} = 0.$$

Thus, under the usual concavity conditions, private and welfare-optimal information choices coincide.

Under the production subsidy, the government instead reimburses the fraction $1-1/\mathcal{M}$ of regular and urgent production costs. The firm pays only $(L_{i,t}^R + \nu L_{i,t}^U)/\mathcal{M}$, so its operating payoff is

$$R_{i,t} - \frac{1}{\mathcal{M}}\left(L_{i,t}^R + \nu L_{i,t}^U\right).$$

This policy also eliminates the markup because the firm's perceived marginal cost is divided by $\mathcal{M}$. However, its operating payoff is only $1/\mathcal{M}$ of the payoff under the sales subsidy. It therefore leaves Lemma 3's coefficient equal to $1/\eta$ and Lemma 4 unchanged in form. At the private interior optimum, the modified

Proposition 2 gives

$$\left.\frac{\partial W^*(x_t;\boldsymbol{\sigma})}{\partial \sigma_{g,t}}\right|_{\boldsymbol{\sigma}^p} = \beta(\mathcal{M}-1)m_g\left(\eta\bar{V}_t^0\right)^{\frac{1-\eta\gamma}{\gamma(\eta-1)}}\frac{\partial \bar{V}_{g,t}^0}{\partial \sigma_{g,t}} < 0.$$

Hence production subsidies do not eliminate information under-acquisition: implementing p equal to marginal cost is not sufficient unless the subsidy also changes the firm's revenue payoff.

In summary, the two subsidies implement the same prices and quantities for a fixed uncertainty profile, but they are not equivalent when information is endogenous. The sales subsidy changes the profit–revenue coefficient from $1/\eta$ to $1/(\eta-1)$ and eliminates the information wedge, whereas the production subsidy leaves that coefficient, and therefore under-acquisition, unchanged. Appendix A.8 derives the two subsidy coefficients.

B. Linear demand function

The information result is not specific to CES demand. With a separable quadratic consumption aggregator, individual demand is linear and the optimal supply function is piecewise linear, with an inventory segment between the regular- and urgent-production regions. At the operating optimum, consumer surplus equals one half of gross operating profit. Welfare therefore values the operating-profit gain from information by a factor of $3/2$, while the firm values it by one, so information remains under-acquired. The same logic implies that higher urgency costs can raise decentralized welfare when the induced increase in information acquisition dominates their direct cost. The linear-quadratic structure also delivers a closed-form information optimality condition and a single-peaked welfare objective.

VII. Conclusion

This paper studies information acquisition by firms that use inventories and costly urgent production to meet uncertain demand. The optimal supply function has two constant-markup branches and a vertical segment at the firm's inventory position. Information improves the allocation of inventories across demand states, but atomistic firms account only for its private profit effect and treat aggregate consumption and the common demand shifter as fixed.

As a result, firms under-acquire information relative to the welfare optimum. This result holds for every positive value of relative risk aversion and does not require firm heterogeneity. In the symmetric economy, a common change in uncertainty affects the entire mass of firms and generates the strongest aggregate effect. Depreciation heterogeneity is nevertheless important for the positive implications of the model because it allows the inventory–uncertainty relationship to vary in sign and magnitude across firms.

The empirical evidence supports the main mechanisms, and the numerical exercise quantifies the welfare gain from moving to the welfare-optimal information profile. The policy analysis shows that eliminating the markup is not by itself sufficient to correct information acquisition: a sales subsidy removes the information wedge, while a production subsidy does not. The same under-acquisition mechanism also holds under linear demand.

References

George-Marios Angeletos and Jennifer La'O. Noisy business cycles. *NBER Macroeconomics Annual*, 24 (1):319–378, 2010.

Kenneth J. Arrow, Theodore Harris, and Jacob Marschak. Optimal inventory policy. *Econometrica*, 19(3): 250–272, 1951. URL http://www.jstor.org/stable/1906813.

Vladimir Asriyan and Alexander Kohlhas. Data, forecasting, and price discrimination. Working Paper, 2025.

Kubilay Bostanci and Guillermo Ordoñez. Liquidity, information, and crises. Working Paper, 2025.

Maryam Farboodi and Laura Veldkamp. A model of the information economy. *Journal of Political Economy*, 128(8):3168–3215, 2020.

Joel P. Flynn, George Nikolakoudis, and Karthik A. Sastry. A theory of supply function choice and aggregate supply. *American Economic Review*, 116(2):710–748, 2026. doi: 10.1257/aer.20240443.

Charles I. Jones and Christopher Tonetti. Nonrivalry and the economics of data. *American Economic Review*, 110(9):2819–2858, 2020.

James A. Kahn. Inventories and the volatility of production. *The American Economic Review*, 77(4): 667–679, 1987. URL http://www.jstor.org/stable/1814538.

Aubhik Khan and Julia K. Thomas. Inventories and the business cycle: An equilibrium analysis of (s, s) policies. *American Economic Review*, 97(4):1165–1188, 2007. doi: 10.1257/aer.97.4.1165. URL https://www.aeaweb.org/articles?id=10.1257/aer.97.4.1165.

Paul D. Klemperer and Margaret A. Meyer. Supply function equilibria in oligopoly under uncertainty. *Econometrica*, 57(6):1243–1277, 1989. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1913707.

Luis Gonzalo Llosa and Venky Venkateswaran. Efficiency under endogenous information choice: A price-setting application. *Journal of the European Economic Association*, 20(4):1619–1646, 2022. doi: 10.1093/jeea/jvac010.

Ricardo Reis. Inattentive producers. *Review of Economic Studies*, 73(3):793–821, 2006.

Xun Wang and Stephen M. Disney. The bullwhip effect: Progress, trends and directions. *European Journal of Operational Research*, 250(3):691–701, 2016.

Michael Woodford. Imperfect common knowledge and the effects of monetary policy. In Philippe Aghion, Roman Frydman, Joseph E. Stiglitz, and Michael Woodford, editors, *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, pages 25–58. Princeton University Press, Princeton, 2003.

A. Proofs

A.1. Proof of Lemma 1

Using the production constraints and $w_{i,t} = \nu$ whenever urgent labor is used, flow operating profit can be written as

$$\pi^0_{i,t} = (\hat{p}_{i,t} - 1)\hat{q}_{i,t} - I_{i,t} + (1-\delta)I_{i,t-1} - (\nu - 1)\big(\hat{q}_{i,t} - (1-\delta)I_{i,t-1}\big)^+.$$

The inherited supply function and inventory determine current sales, whereas $f_{i,t}$ and $I_{i,t}$ determine next-period sales. Let $J_{i,t} = (1-\delta)I_{i,t}$. The terms depending on these current choices are therefore

$$\beta\mathbb{E}\big[(\hat{p}_{i,t+1} - 1)\hat{q}_{i,t+1} - (\nu - 1)(\hat{q}_{i,t+1} - J_{i,t})^+ + J_{i,t} \mid \theta_{i,t}, s_{i,t}\big] - \frac{J_{i,t}}{1-\delta}.$$

Dividing by the positive constant β gives the objective in the lemma, because the coefficient on $J_{i,t}$ becomes $-[1/((1-\delta)\beta) - 1]$. Taking the expectation of this optimized conditional value over $s_{i,t}$ gives $V^0_{i,t}$. The remaining continuation terms repeat the same problem at later dates, which gives the sequence of static operating problems. □

A.2. Proof of Lemma 2

For fixed $J := J_{i,t}$ and next-period demand shifter $A := A_{i,t+1}$, the firm chooses q to maximize the concave, kinked objective

$$A^{1/\eta}q^{1-1/\eta} - q - (\nu - 1)(q - J)^+.$$

For $q < J$, the first-order condition is $(1 - 1/\eta)A^{1/\eta}q^{-1/\eta} = 1$. Since $\mathcal{M} = \eta/(\eta - 1)$, it gives $q = A/\mathcal{M}^\eta$, which lies below J if $A < J\mathcal{M}^\eta$. For $q > J$, marginal cost is ν, so the first-order condition gives $q = A/(\mathcal{M}\nu)^\eta$, which lies above J if $A > J(\mathcal{M}\nu)^\eta$. For intermediate values of A, the left derivative at J is nonnegative and the right derivative is nonpositive, so the optimum is the kink $q = J$. These three cases give the stated rule. □

A.3. Proof of Proposition 1

Inverse demand is $p = A^{1/\eta}q^{-1/\eta}$. On the low-demand branch, $q = A/\mathcal{M}^\eta$ and hence $p = \mathcal{M}$. On the high-demand branch, $q = A/(\mathcal{M}\nu)^\eta$ and hence $p = \mathcal{M}\nu$. In the intermediate region, $q = J$ while A ranges from $J\mathcal{M}^\eta$ to $J(\mathcal{M}\nu)^\eta$, so the market-clearing price ranges from $\mathcal{M}$ to $\mathcal{M}\nu$. This is the vertical segment at $q = J$. □

A.4. Derivation of the optimal inventory choice

Write $h := \nu - 1$ and $\lambda := 1/[(1-\delta)\beta] - 1$. The envelope theorem applied to the quantity rule gives the marginal value of after-depreciation inventory:

$$\mathbb{E}\Big[\Big(\frac{p}{\mathcal{M}} - 1\Big)\mathbf{1}\{J\mathcal{M}^\eta \le A \le J(\mathcal{M}\nu)^\eta\} + h\mathbf{1}\{A > J(\mathcal{M}\nu)^\eta\}\Big] - \lambda.$$

Let $Z = (\log A - \mu^A)/\sigma \sim N(0,1)$ and define

$$a = \frac{\log(J\mathcal{M}^\eta) - \mu^A}{\sigma}, \qquad b = a + \frac{\eta \log \nu}{\sigma}, \qquad \tau = \frac{\sigma}{\eta}.$$

On the vertical segment, $p/\mathcal{M} = e^{\tau(Z-a)}$. The interior first-order condition is therefore

$$G_\sigma(a,b) := \int_a^b \left(e^{\tau(z-a)} - 1\right) \phi(z)\, dz + h[1 - \Phi(b)] = \lambda,$$

which is the normal-CDF expression reported below Equation (3). Holding σ fixed, boundary terms cancel and

$$\frac{\partial G_\sigma}{\partial a} = -\tau \int_a^b e^{\tau(z-a)} \phi(z)\, dz < 0.$$

Moreover, $G_\sigma \to h$ as $a \to -\infty$ and $G_\sigma \to 0$ as $a \to \infty$. Thus, $0 < \lambda < h$ implies a unique cutoff $a^*(\sigma)$. Solving the definition of a for J gives Equation (3).

For completeness, implicit differentiation of the first-order condition gives

$$a'(\sigma) = \frac{\int_a^b (z-a) e^{\tau(z-a)} \phi(z)\, dz}{\sigma \int_a^b e^{\tau(z-a)} \phi(z)\, dz}.$$

Hence

$$a + \sigma a' = \frac{\int_a^b z e^{\tau(z-a)} \phi(z)\, dz}{\int_a^b e^{\tau(z-a)} \phi(z)\, dz},$$

which is the tilted conditional mean used in the inventory comparative static. Differentiating $J^* = \exp(\mu^A + \sigma a^*(\sigma))/\mathcal{M}^\eta$ then gives the fixed-posterior-mean formula in the text.

Under the Gaussian signal structure in Section II, the posterior mean is

$$\mu = \mu^- + \left(1 - \frac{\sigma^2}{(\sigma^-)^2}\right)(s - \mu^-).$$

Holding the signal realization fixed therefore gives

$$\frac{d\mu}{d\sigma} = -\frac{2\sigma}{(\sigma^-)^2 - \sigma^2}(\mu - \mu^-).$$

Because the common demand shifter is held fixed, $\mu^A - \mu^{A,-} = \mu - \mu^-$ and $d\mu^A/d\sigma = d\mu/d\sigma$. Hence,

$$\frac{d\mu^A}{d\sigma} = -\frac{2\sigma}{(\sigma^-)^2 - \sigma^2}(\mu^A - \mu^{A,-}).$$

Adding this term to $d \log J^*/d\sigma = a + \sigma a'$ gives the signal-contingent total derivative. Before the signal is

observed,

$$\mu^A \sim N\left(\mu^{A,-}, (\sigma^-)^2 - \sigma^2\right),$$

and therefore

$$\mathbb{E}[J^*] = \frac{1}{\mathcal{M}^\eta} \exp\left(\mu^{A,-} + \frac{1}{2}[(\sigma^-)^2 - \sigma^2] + \sigma a^*(\sigma)\right).$$

Log differentiation yields

$$\frac{d\mathbb{E}[J^*]}{d\sigma} = \mathbb{E}[J^*]\,(a + \sigma a' - \sigma),$$

which is the ex ante expression in the text after substituting the tilted conditional mean for $a + \sigma a'$. □

A.5. Proof of Lemma 3

Suppress firm and time subscripts. At the operating optimum, normalized gross profit is

$$\bar{V}^0 = \mathbb{E}[\bar{R} - Q - (\nu - 1)(Q - J)^+] - \lambda J.$$

Because $1 - 1/\eta = 1/\mathcal{M}$,

$$\bar{V}^0 - \frac{1}{\eta}\mathbb{E}[\bar{R}] = \mathbb{E}\left[\frac{\bar{R}}{\mathcal{M}} - Q - (\nu - 1)(Q - J)^+\right] - \lambda J.$$

This expression is zero in the low-demand region, where $p = \mathcal{M}$ and $Q < J$. In the middle region it equals $J(p/\mathcal{M} - 1)$, and in the high-demand region it equals $(\nu - 1)J$. Therefore

$$\bar{V}^0 - \frac{1}{\eta}\mathbb{E}[\bar{R}] = J\left\{\mathbb{E}\left[\left(\frac{p}{\mathcal{M}} - 1\right)\mathbf{1}\{J\mathcal{M}^\eta \le A \le J(\mathcal{M}\nu)^\eta\}\right] + (\nu - 1)\Pr\{A > J(\mathcal{M}\nu)^\eta\} - \lambda\right\}.$$

The product is zero either because $J = 0$ at the lower boundary or because the term in braces is the inventory first-order condition and equals zero at an interior optimum. This proves $\bar{V}^0 = (1/\eta)\mathbb{E}[\bar{R}]$ in the normalized problem. If $S := C_{t+1}^{1/\eta-\gamma}$ is the common demand shifter, homogeneity gives $V^0 = S^\eta \bar{V}^0$ and $\mathbb{E}[R] = S^\eta \mathbb{E}[\bar{R}]$. Therefore the same identity holds in general equilibrium: $V^0 = (1/\eta)\mathbb{E}[R]$.

It remains to sign the direct effect of posterior uncertainty. Let $Z \sim N(0, 1)$ denote the standardized residual demand shock and let $r_\sigma(Z)$ be normalized realized revenue after factoring out the posterior mean. Conditional expected normalized revenue is $\mathcal{R}(\sigma) := \mathbb{E}[r_\sigma(Z)]$. Since normalized profit equals $\mathcal{R}(\sigma)/\eta$, and a change in σ changes log demand by Z, the envelope theorem gives

$$\frac{1}{\eta}\mathcal{R}'(\sigma) = \frac{1}{\eta}\mathbb{E}[Z r_\sigma(Z)], \qquad \text{or equivalently} \qquad \mathcal{R}'(\sigma) = \mathbb{E}[Z r_\sigma(Z)].$$

Ex ante, the posterior mean has variance $(\sigma^-)^2 - \sigma^2$, so expected normalized revenue is

$$\mathbb{E}[\bar{R}] = \exp\left(\mu^- + \frac{1}{2}[(\sigma^-)^2 - \sigma^2]\right)\mathcal{R}(\sigma).$$

Integration by parts gives $\mathbb{E}[Zr_\sigma(Z)] = \mathbb{E}[\partial r_\sigma(Z)/\partial Z]$. Revenue is continuous at the two cutoffs, and

$$\frac{\partial r_\sigma(Z)}{\partial Z} = \begin{cases} \sigma r_\sigma(Z), & Z < a(\sigma), \\ (\sigma/\eta) r_\sigma(Z), & a(\sigma) \le Z \le b(\sigma), \\ \sigma r_\sigma(Z), & Z > b(\sigma). \end{cases}$$

Consequently,

$$\frac{\partial \mathbb{E}[\bar{R}]}{\partial \sigma} = -\exp\left(\mu^- + \frac{1}{2}[(\sigma^-)^2 - \sigma^2]\right) \sigma \left(1 - \frac{1}{\eta}\right) \mathbb{E}[r_\sigma(Z)\mathbf{1}\{a(\sigma) \le Z \le b(\sigma)\}] \le 0.$$

The inequality is strict when the vertical region has positive probability, as it does at an interior inventory optimum, and becomes an equality at the zero-inventory boundary. Because $\bar{V}^0 = (1/\eta)\mathbb{E}[\bar{R}]$, it follows directly that $\partial_\sigma \bar{V}^0 \le 0$, with the same strictness condition. □

A.6. Proof of Lemma 4

Let $S_{t+1} := C_{t+1}^{1/\eta-\gamma}$. Scaling S_{t+1} by a positive factor scales the optimal quantity and inventory rules by that factor to the power η. Consequently,

$$V_{i,t}^0 = S_{t+1}^\eta \bar{V}_{i,t}^0, \qquad V_t^0 = S_{t+1}^\eta \bar{V}_t^0.$$

The CES aggregator is homogeneous of degree one, and inverse demand implies aggregate expenditure

$$R_{t+1} = \int \mathbb{E}_t[p_{i,t+1} q_{i,t+1}]di = C_{t+1}^{1/\eta-\gamma} \int \mathbb{E}_t[\theta_{i,t+1}^{1/\eta} q_{i,t+1}^{1-1/\eta}]di = C_{t+1}^{1-\gamma}.$$

By Lemma 3, normalized aggregate revenue is $\eta \bar{V}_t^0$. Revenue has the same degree-η scaling, so

$$C_{t+1}^{1-\gamma} = S_{t+1}^\eta \eta \bar{V}_t^0 = C_{t+1}^{1-\eta\gamma} \eta \bar{V}_t^0.$$

It follows that $C_{t+1}^{\gamma(\eta-1)} = \eta \bar{V}_t^0$ and hence

$$S_{t+1}^\eta = \left(\eta \bar{V}_t^0\right)^{\frac{1-\eta\gamma}{\gamma(\eta-1)}}.$$

Substitution into $V_{i,t}^0 = S_{t+1}^\eta \bar{V}_{i,t}^0$ proves the first claim.

An atomistic firm treats $\bar{V}_t^0$, and therefore S_{t+1}, as fixed when it changes its own uncertainty. At an interior uncertainty choice, its marginal operating benefit is consequently $\beta S_{t+1}^\eta \partial_{\sigma_{i,t}} \bar{V}_{i,t}^0$. Equating this benefit to the marginal information-cost saving gives the decentralized first-order condition in the lemma. □

A.7. Proof of Proposition 2

Let $S_{t+1}^{\eta} = V_t^0/\bar{V}_t^0$. CES demand implies $R_{t+1} = C_{t+1}^{1-\gamma}$ and

$$CS_{t+1}^{flow} = u_\gamma(C_{t+1}) - C_{t+1}^{1-\gamma}.$$

From $C_{t+1}^{\gamma(\eta-1)} = \eta \bar{V}_t^0$, direct differentiation gives, for every $\gamma > 0$,

$$\frac{\partial CS_{t+1}^{flow}}{\partial \bar{V}_t^0} = \frac{\eta}{\eta - 1} S_{t+1}^{\eta}.$$

For $\gamma = 1$, the same expression follows directly from $CS_{t+1}^{flow} = \log C_{t+1} - 1$ and $C_{t+1}^{\eta-1} = \eta \bar{V}_t^0$. An interior uncertainty choice requires a vertical inventory region of positive probability; otherwise Lemma 3 gives no operating value of information and the firm chooses the no-information boundary. Hence $\partial_{\sigma_{g,t}} \bar{V}_t^0 = m_g \partial_{\sigma_{g,t}} \bar{V}_{g,t}^0 < 0$, discounting the next-period effect gives

$$\frac{\partial CS^*}{\partial \sigma_{g,t}} = \beta \frac{\eta}{\eta - 1} \frac{V_t^0}{\bar{V}_t^0} \frac{\partial \bar{V}_t^0}{\partial \sigma_{g,t}} < 0.$$

Differentiating aggregate operating profit $V_t^0 = S_{t+1}^{\eta} \bar{V}_t^0$ gives

$$\frac{\partial V_t^0}{\partial \sigma_{g,t}} = \frac{V_t^0}{\bar{V}_t^0} \frac{1 - \gamma}{\gamma(\eta - 1)} \frac{\partial \bar{V}_t^0}{\partial \sigma_{g,t}}.$$

At the interior decentralized uncertainty choice, the private first-order condition in Lemma 4, integrated over the mass m_g of type-g firms, implies

$$\frac{\partial}{\partial \sigma_{g,t}} \int \mathcal{K}_i(\boldsymbol{\sigma}^p) di = \beta \frac{V_t^0}{\bar{V}_t^0} \frac{\partial \bar{V}_t^0}{\partial \sigma_{g,t}}.$$

Using the welfare identity (2), its derivative is therefore

$$\begin{aligned} \frac{\partial W^*}{\partial \sigma_{g,t}} &= \beta \frac{V_t^0}{\bar{V}_t^0} \left[\frac{\eta}{\eta - 1} + \frac{1 - \gamma}{\gamma(\eta - 1)} - 1 \right] \frac{\partial \bar{V}_t^0}{\partial \sigma_{g,t}} \\ &= \beta \frac{1}{\gamma(\eta - 1)} \frac{V_t^0}{\bar{V}_t^0} \frac{\partial \bar{V}_t^0}{\partial \sigma_{g,t}} < 0. \end{aligned}$$

The use of the private first-order condition is why the welfare equality holds at the decentralized equilibrium, rather than at every uncertainty profile. □

A.8. Derivation of the subsidy coefficients

Suppress firm and time subscripts, write $h := \nu - 1$, and define normalized operating costs by

$$D(Q, J) := Q + h(Q - J)^+ + \lambda J.$$

Thus, absent subsidies, normalized operating profit is $\mathbb{E}[\bar{R} - D(Q, J)]$, where $\bar{R} = pQ$. The two subsidies alter this objective differently.

Under the sales subsidy, the government adds $(\mathcal{M} - 1)\bar{R}$ to the firm's revenue. The normalized operating problem becomes

$$\bar{V}^{0,S} = \max_{Q(\cdot),J} \mathbb{E}\left[\mathcal{M}\bar{R} - D(Q, J)\right].$$

The quantity first-order condition is $\mathcal{M}(1 - 1/\eta)p = 1$ when regular production is marginal and $\mathcal{M}(1 - 1/\eta)p = \nu$ when urgent production is marginal. Since $\mathcal{M} = \eta/(\eta - 1)$, the resulting prices are $p = 1$ and $p = \nu$, respectively. The sales subsidy therefore eliminates the markup.

At an interior inventory choice, the inventory first-order condition is

$$\mathbb{E}[(p - 1)\mathbf{1}\{Q = J\} + h\mathbf{1}\{Q > J\}] - \lambda = 0.$$

To obtain the profit–revenue coefficient, subtract $(\mathcal{M} - 1)\mathbb{E}[\bar{R}]$ from the optimized payoff:

$$\begin{aligned}\bar{V}^{0,S} - (\mathcal{M} - 1)\mathbb{E}[\bar{R}] &= \mathbb{E}\left[\bar{R} - Q - h(Q - J)^+\right] - \lambda J \\ &= J\left(\mathbb{E}[(p - 1)\mathbf{1}\{Q = J\} + h\mathbf{1}\{Q > J\}] - \lambda\right) = 0.\end{aligned}$$

The second equality uses the three operating regions. The expression inside the expectation is zero when $Q < J$, equals $J(p - 1)$ when $Q = J$, and equals hJ when $Q > J$. Hence

$$\bar{V}^{0,S} = (\mathcal{M} - 1)\mathbb{E}[\bar{R}] = \frac{1}{\eta - 1}\mathbb{E}[\bar{R}].$$

Under the production subsidy, the government reimburses the fraction $1 - 1/\mathcal{M}$ of all regular and urgent production costs. Because inventories are produced with regular labor, their carrying cost is scaled by the same factor. The normalized operating problem is therefore

$$\begin{aligned}\bar{V}^{0,P} &= \max_{Q(\cdot),J} \mathbb{E}\left[\bar{R} - \frac{1}{\mathcal{M}}D(Q, J)\right] \\ &= \frac{1}{\mathcal{M}} \max_{Q(\cdot),J} \mathbb{E}\left[\mathcal{M}\bar{R} - D(Q, J)\right] = \frac{1}{\mathcal{M}}\bar{V}^{0,S}.\end{aligned}$$

Multiplication of the objective by the positive constant $1/\mathcal{M}$ does not change the optimal quantity or inventory rules. Thus, the production subsidy also implements $p = 1$ and $p = \nu$, but it scales down the

firm's entire payoff relative to the sales subsidy. Using the sales-subsidy identity gives

$$\bar{V}^{0,P} = \frac{\mathcal{M}-1}{\mathcal{M}}\mathbb{E}[\bar{R}] = \frac{1}{\eta}\mathbb{E}[\bar{R}].$$

Let $S_{t+1} := C_{t+1}^{1/\eta-\gamma}$. Profit and revenue both scale by S_{t+1}^{η}, so the corresponding GE identities are

$$V_{i,t}^{0,S} = \frac{1}{\eta-1}\mathbb{E}_t[R_{i,t+1}], \qquad V_{i,t}^{0,P} = \frac{1}{\eta}\mathbb{E}_t[R_{i,t+1}].$$

If $\bar{V}_t^{0,S} := \int \bar{V}_{i,t}^{0,S} di$ and $\bar{V}_t^{0,P} := \int \bar{V}_{i,t}^{0,P} di$, market clearing implies

$$S_{t+1}^{\eta} = \left((\eta-1)\bar{V}_t^{0,S}\right)^{\frac{1-\eta\gamma}{\gamma(\eta-1)}} = \left(\eta\bar{V}_t^{0,P}\right)^{\frac{1-\eta\gamma}{\gamma(\eta-1)}}.$$

To see the connection with operating welfare explicitly, fix the common demand shifter and a realization of normalized demand A. The gross benefit from supplying Q units is the area under inverse demand:

$$\int_0^Q A^{1/\eta} q^{-1/\eta} dq = \mathcal{M} A^{1/\eta} Q^{1-1/\eta} = \mathcal{M}\bar{R}.$$

The planner's normalized operating problem for this firm, before demand is realized, is therefore

$$\max_{Q(\cdot),J} \mathbb{E}\left[\int_0^{Q(A)} A^{1/\eta} q^{-1/\eta} dq - D(Q(A),J)\right] = \max_{Q(\cdot),J} \mathbb{E}\left[\mathcal{M}\bar{R} - D(Q,J)\right].$$

This is exactly the sales-subsidy problem defining $\bar{V}^{0,S}$, term by term: $\mathcal{M}\bar{R}$ is gross consumption benefit and $D(Q,J)$ is the resource cost of regular production, urgent production, and inventories. Hence, for a fixed common demand shifter and information profile, the sales-subsidized firm chooses the same state-contingent quantities and inventories as the planner. The production-subsidized firm implements the same operating allocation only because its entire objective is $1/\mathcal{M}$ times this expression; its private operating payoff is correspondingly only $1/\mathcal{M}$ of operating surplus.

The conditional equivalence makes the operating allocation efficient for every fixed information profile. The planner's envelope theorem therefore removes the first-order welfare effects operating through adjustments in the operating choices. Subsidy payments and the lump-sum taxes that finance them also cancel because they are transfers. Under the sales subsidy, the planner's marginal value of a type-wide change is

$$\frac{\partial W^*}{\partial \sigma_{g,t}} = \beta m_g S_{t+1}^{\eta} \frac{\partial \bar{V}_{g,t}^{0,S}}{\partial \sigma_{g,t}} - m_g \frac{\partial \mathcal{K}_g}{\partial \sigma_{g,t}}.$$

The first term is the direct change in operating surplus, with S_{t+1}^{η} converting the normalized objective into the actual allocation. The second term is the information cost. The sales-subsidized firm's private condition is

$$\frac{\partial \mathcal{K}_g}{\partial \sigma_{g,t}} = \beta S_{t+1}^{\eta} \frac{\partial \bar{V}_{g,t}^{0,S}}{\partial \sigma_{g,t}}.$$

Therefore, at the interior private optimum,

$$\left.\frac{\partial W^*(x_t;\boldsymbol{\sigma})}{\partial \sigma_{g,t}}\right|_{\boldsymbol{\sigma}^P} = 0.$$

The equality holds because the sales-subsidy payoff equals the planner's conditional operating surplus.

Under the production subsidy, the planner continues to value the full conditional operating surplus $\bar{V}^{0,S}$, whereas the firm receives only $\bar{V}^{0,P} = \bar{V}^{0,S}/\mathcal{M}$. Hence

$$\begin{aligned}\frac{\partial W^*}{\partial \sigma_{g,t}} &= \beta m_g S_{t+1}^{\eta} \frac{\partial \bar{V}_{g,t}^{0,S}}{\partial \sigma_{g,t}} - m_g \frac{\partial \mathcal{K}_g}{\partial \sigma_{g,t}} \\ &= \beta \mathcal{M} m_g S_{t+1}^{\eta} \frac{\partial \bar{V}_{g,t}^{0,P}}{\partial \sigma_{g,t}} - m_g \frac{\partial \mathcal{K}_g}{\partial \sigma_{g,t}}.\end{aligned}$$

The production-subsidized firm's private condition is

$$\frac{\partial \mathcal{K}_g}{\partial \sigma_{g,t}} = \beta S_{t+1}^{\eta} \frac{\partial \bar{V}_{g,t}^{0,P}}{\partial \sigma_{g,t}}.$$

Substituting this condition and $S_{t+1}^{\eta} = (\eta \bar{V}_t^{0,P})^{(1-\eta\gamma)/[\gamma(\eta-1)]}$ gives

$$\left.\frac{\partial W^*(x_t;\boldsymbol{\sigma})}{\partial \sigma_{g,t}}\right|_{\boldsymbol{\sigma}^P} = \beta(\mathcal{M}-1) m_g \left(\eta \bar{V}_t^{0,P}\right)^{\frac{1-\eta\gamma}{\gamma(\eta-1)}} \frac{\partial \bar{V}_{g,t}^{0,P}}{\partial \sigma_{g,t}} < 0.$$

Thus, the sales subsidy eliminates the information wedge, whereas the production subsidy leaves firms with too little information despite implementing the same operating allocation for a fixed information profile. □

B. Empirical classification by inventory depreciation

The empirical exercise assigns each firm to a depreciation group using its primary NAICS industry. Depreciation is interpreted as the effective loss in the value of goods held in inventory. It includes physical spoilage and damage, expiration, seasonal or fashion markdowns, and technological obsolescence. Price risk by itself is not treated as physical inventory depreciation. Service-heavy industries are generally classified as low because physical goods are not their principal product.

The broad classification uses two-digit NAICS sectors. Table 3 reports the complete mapping used for the Broad NAICS2 column of Table 2.

Table 3: Broad NAICS2 inventory-depreciation classification

NAICS	Sector	Group	Main source of effective depreciation
21	Mining, oil, and gas	Low	Stored commodities and materials are physically durable.
22	Utilities	Low	Inventories are mainly fuel, materials, and spare parts.
23	Construction	Middle	Work in process is project-specific and exposed to damage.
31–33	Manufacturing	Middle	Product mix combines durable inputs with obsolescent finished goods.
42	Wholesale trade	Middle	Mix of durable and faster-moving goods for resale.
44–45	Retail trade	High	Seasonality, fashion, perishability, and markdown risk.
48–49	Transportation and warehousing	Low	Inventories are mainly fuel, parts, and operating supplies.
51	Information	Low	Firms primarily sell software, media, and digital access.
52	Finance and insurance	Low	Physical goods inventory is not central to production.
53	Real estate and rental	Low	Firms hold long-lived assets rather than ordinary goods inventory.
54	Professional services	Low	Firms primarily sell services and expertise.
56	Administrative services	Low	Inventories are mainly operating supplies.
61	Educational services	Low	Physical goods are not the principal product.
62	Health care	Middle	Pharmaceuticals and medical supplies face expiration and sterility risk.
71	Arts and entertainment	Middle	Merchandise and supplies face event-specific and seasonal risk.
72	Accommodation and food	High	Food and beverages spoil and lose quality rapidly.
81	Other services	Low	Inventories are mainly parts, supplies, or small resale stocks.

The detailed classification assigns the same firms using three-digit NAICS industries. It refines broad sectors whose products have different storage properties. For example, food manufacturing (311), apparel manufacturing (315), and computer and electronic product manufacturing (334) are high, whereas primary metals (331) and fabricated metals (332) are low; the broad manufacturing sector is middle. Similarly, food retail (445) and clothing retail (448 and 458) are high, while gasoline stations and fuel dealers (457) are low. Table 4 gives the complete code mapping and representative firms from the estimation sample.

Table 4: Detailed NAICS3 inventory-depreciation classification

Group	NAICS3 industries	Representative firms
Low	213, 221, 324, 327, 331, 332, 457, 481–485, 488, 492, 511, 513, 516–519, 522–524, 531–533, 541, 561–562, 611, 721, 811–812	Halliburton, SLB, IBM, Adobe, Paychex
Middle	42, 236–238, 313, 321–323, 325–326, 333, 335–337, 339, 423, 441–442, 444, 451, 453–454, 459, 621–622, 713	Deere, Applied Materials, CVS Health, Healthcare Services Group
High	311–312, 315–316, 334, 424, 445, 448–449, 455–456, 458, 722	AMD, Analog Devices, Kroger, Ross Stores, Starbucks

Notes: A firm's group is fixed by its primary NAICS code in the merged IBES–Compustat guidance sample. The detailed mapping is judgment-based and applies the same definition of effective depreciation as the broad mapping. Representative firms illustrate the classification and are not used to assign other firms.