Price as a Quality Promise

Dynamic Pricing when Reviews Judge Value for Money

João Ramos and Peiran Xiao

July 2026 — SAET, Rio de Janeiro

Motivation

- Reviews drive demand — and reviews judge value for money, not quality alone: "worth it at \$20, not at \$60"; hotel sites even publish a *value-for-money* score
- The same product gets **worse reviews when it costs more** (identical rooms rated lower in peak season; common review for online markets: "Bought it on sale for \$30, worth it! So soft!")
- Price is **the bar the product must clear**: by pricing, the firm chooses *the difficulty of the exam the product takes* — a price is a promise of a quality
- But raising the bar is harder for good *and* bad products alike
  $\Rightarrow$ the firm cannot tilt the exam in its favor — only make it more or less revealing

This paper

A monopolist with **privately known quality** faces a demand; a review is "good" iff the quality experienced clears the price paid; buyers read the review *at the price that generated it*. What pricing does review management produce?

Questions and preview

Questions. What can a firm do to its reputation *through promises of quality via price*? What does ratings management look like when the price selects the informativeness of the review?

1. Price is an informativeness dial, not a bias dial: it cannot tilt the review's direction, only sharpen or blur the test
2. No level distortion: mean price = myopic price for both types.
   The reputational motive is **mean-preserving** — it acts on price *dispersion* following the demand shock
3. Quality is not always more expensive: Price–quality ordering **reverses** across demand regimes
   in booms, expensive = good quality; in slumps cheap = good

Model: environment

- Two periods $t = 1, 2$; discount $\beta \in (0, 1)$; monopolist with fixed private quality $\theta \in \{\theta_L, \theta_H\}$ and prior $\mu = \Pr(\theta_H)$; perceived quality $\hat{\theta}(m) = \theta_L + m\Delta_\theta$
- Demand is random, but depends on price and expected quality in a simple way:

$$D = 2\eta - \frac{p}{\hat{\theta}(\text{belief})},$$

η i.i.d. and known by the firm, $\mathbb{E}[\eta] = 1$, $\zeta := \mathbb{E}[\eta^2]$

- Myopic price $p^{\mathrm{m}}(\eta) = \hat{\theta}\,\eta$; **expected myopic price** $\mathbb{E}_\eta[p^{\mathrm{m}}] = \hat{\theta}$
  $=$ the *ex-ante fair-value price*: product offered at exactly its perceived value
- Timing: firm sees quality and η_1, posts p and sells according to D. $\rightarrow$ Review $s \in \{\text{good, bad}\} \rightarrow$ buyers update $\mu_2 \rightarrow$ terminal period, myopic pricing at $\hat{\theta}(\mu_2)$

Model: reviews as value-for-money verdicts

Buyer experience of the good depends on quality and taste, adaptation, fit, etc

$$q = \theta + \varepsilon, \qquad \varepsilon \sim U[-A, A];$$

Buyer posts **good** review iff experienced quality is higher than the promise (price):

Review is positive iff $q \ \geq \ p$

Model: reviews as value-for-money verdicts

Buyer experience of the good depends on quality and taste, adaptation, fit, etc

$$q = \theta + \varepsilon, \qquad \varepsilon \sim U[-A, A];$$

Buyer posts **good** review iff experienced quality is higher than the promise (price):

Review is positive iff $q \;\geq\; p$

The probability of a positive review depends both on type and price.

$$L_k(p) = \frac{\theta_k + A - p}{2A}, \qquad \textit{with} \quad \bar{L}(p) = \mu L_H + (1-\mu) L_L = \frac{\hat{\theta} + A - p}{2A}$$

Higher quality makes it easier to pass any bar; but the level of the bar is chosen by the price.

Updating with reviews as value-for-money verdicts

Bayes update depends on price level:

$$\mu_2^+(p) = \frac{\mu L_H(p)}{\bar{L}(p)}, \qquad \mu_2^-(p) = \frac{\mu(1 - L_H(p))}{1 - \bar{L}(p)}$$

good review at a demanding price = strong praise; bad review at a demanding price = forgivable

Both are increasing in price, but price controls how far apart the two posteriors sit.

$$\operatorname{sgn}\left(\frac{\partial\left(\mu_2^+ - \mu_2^-\right)}{\partial p}\right) = \operatorname{sgn}\left(p - \hat{\theta}\right)$$

Let the (expected) **reputational gain** from making a promise be:

$$\mathcal{R}_k(p) := \mathbb{E}_k\left[\mu_2(p)\right] - \mu = L_k(p)\mu_2^+(p) + (1 - L_k(p))\mu_2^-(p) - \mu,$$

Reputation gain losses depend on type and price.

Two assumptions

A1 (Noisy reviews) $A > \theta_H$
Experience noise exceeds the quality scale: reviews are weak signals.

Technical importance: keeps the problem well-posed.

A2 (No learning directly from prices)
Buyers draw **no inference regarding quality from the price choice** — but uses the posted price to *interpret the review*. Isolates the (review manipulation) channel.

Map from quality to prices is governed by firm-specific variables, like market size, patience, or liquidity shocks, that may not be known by the consumer. Observing price does not infer quality.

Key lemma: price is an informativeness dial

Lemma 1

(i) The likelihood **spread is price-independent**: $L_H(p) - L_L(p) = \Delta_L$ for all p. Price moves only the *base rate* $\bar{L}(p)$.

(ii) The reputational gains depend on p **only through the predictive variance** $V(p) = \bar{L}(1 - \bar{L})$:

$$\mathcal{R}_H = \frac{\mu(1-\mu)^2 \Delta_L^2}{V(p)}, \qquad \mathcal{R}_L = -\frac{\mu^2(1-\mu)\Delta_L^2}{V(p)}$$

Raising p raises the hurdle *for both types*.

The firm can only choose how sharp a test it will face.

Reputational gains: geometry

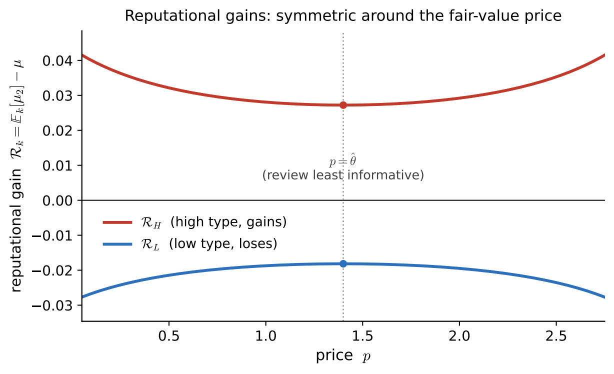

$$\mathcal{R}_H = \frac{\mu(1-\mu)^2\Delta_\theta^2}{A^2 - d^2}$$

$$\mathcal{R}_L = -\frac{\mu^2(1-\mu)\Delta_\theta^2}{A^2 - d^2}$$

$d := p - \hat{\theta}$:

- at $p = \hat{\theta}$: $\bar{L} = \frac{1}{2}$, Variance maximal
- $\mathcal{R}_k'(\hat{\theta}) = 0$: informativeness is **flat at expected myopic value**

H wants extreme prices (sharp test) — L wants fair-value prices (blurred test)

...something is special at the expected myopic value.

Building a reputation is costly

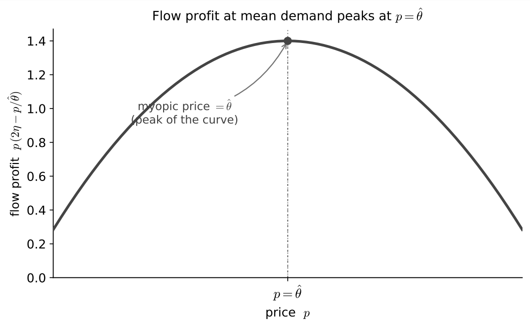

$$\Pi_k(p,\eta) \doteq \underbrace{p\left(2\eta-\frac{p}{\hat\theta}\right)}_{\text{flow profit}} + \underbrace{\beta\,\zeta\,\Delta_\theta\,\mathcal{R}_k(p)}_{\text{reputational payoff}}$$

Flow profit is maximized at the myopic price. Deviating from it to build a reputation means reduced profits in period 1.

$$p^{\mathrm{m}}(\eta)=\hat\theta\,\eta,\quad D^{\mathrm{m}}=\eta,\quad \pi^{\mathrm{m}}=\hat\theta\,\eta^2$$

Making a quality promise towards developing a reputation may be in conflict with optimally serving the demand. The cost depends on the demand shock.

Benchmark: No demand shocks $\Rightarrow$ No dynamic motive

Proposition: No dynamic motive without demand shocks

If $\eta \equiv 1$, then $p_H = p_L = \hat{\theta} = p^{\mathrm{m}}$.
No reputation building.

Why?

Marginal benefit vs Marginal Cost

$$\underbrace{2 - \frac{2p}{\hat{\theta}}\bigg|_{p=\hat{\theta}} = 0}_{\text{static FOC at mean demand}} \qquad \text{and} \qquad \underbrace{\mathcal{R}_k'(\hat{\theta}) = 0}_{\text{informativeness flat at fair value}}$$

Promises are just more vs less informative tests. Informativeness is costly. Naturally interpreted promise has informativeness zero.

The engine: demand shocks move the natural promise understanding

Flow profit: quadratic around the myopic anchor $p^{\mathrm{m}}(\eta) = \hat{\theta}\eta$:

- **Mean demand** $(\eta = 1)$: natural promise for the firm matches natural interpretation, $p^m = E[p^m]$, so informativeness is flat $\Rightarrow$ manipulation worthless at the margin $\Rightarrow$ no distortion
- **Strong demand** $(\eta > 1)$: natural promise for the firm is not where informativeness is zero anymore, $\mathcal{R}'_H > 0 > \mathcal{R}'_L$
  H: a further price *increase* sharpens the test — 1st-order gain, 2nd-order cost on demand not served
  L: a price *cut* toward fair value blurs it — same terms, opposite direction
- **Weak demand** $(\eta < 1)$: mirror image — now a *deeper discount* is what sharpens the test, so H **undercuts even the myopic price**; L prices back up toward $\hat{\theta}$

Demand variation is not a nuisance — it is what gives price informational power.

The engine

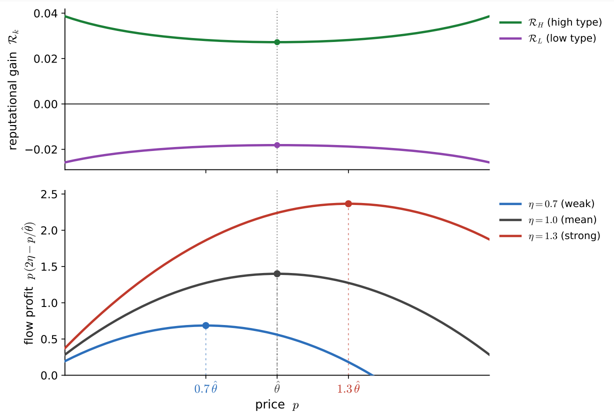

Mean demand $(\eta = 1)$: natural promise for the firm coincides with how the market understands it.

Strong demand $(\eta > 1)$: natural promise is higher than market expectation. Informativeness is obtained by overpricing.

Weak demand $(\eta < 1)$: natural promise is lower than market expects.

Making a quality promise towards developing a reputation may be in conflict with optimally serving the demand. The cost depends on the demand shock.

Solving: one function carries everything

Constants: $\kappa_H = \beta\zeta\mu(1-\mu)^2\Delta_\theta^3$, $\kappa_L = \beta\zeta\mu^2(1-\mu)\Delta_\theta^3$. Define:

$$\Psi_H(d) = d\left[1 - \frac{\hat{\theta}\kappa_H}{(A^2-d^2)^2}\right], \qquad \Psi_L(d) = d\left[1 + \frac{\hat{\theta}\kappa_L}{(A^2-d^2)^2}\right]$$

With Profit derivatives:

$$\frac{\partial \Pi_k}{\partial p} = -\frac{2}{\hat{\theta}}\left[\Psi_k(d) - \hat{\theta}(\eta-1)\right] \qquad \frac{\partial^2 \Pi_k}{\partial p^2} = -\frac{2}{\hat{\theta}}\,\Psi_k'(d)$$

- FOC $\Rightarrow \Psi_k(d) = \hat{\theta}(\eta-1) \;\Rightarrow$ pricing function = inverse of Ψ_k
- SOC $\Rightarrow$ Not globally concave but can obtain properties akin to **monotonicity** of Ψ_k
- FOC is a quintic — no closed-form root, but every property reads off the curve

Existence: single-peakedness, not concavity

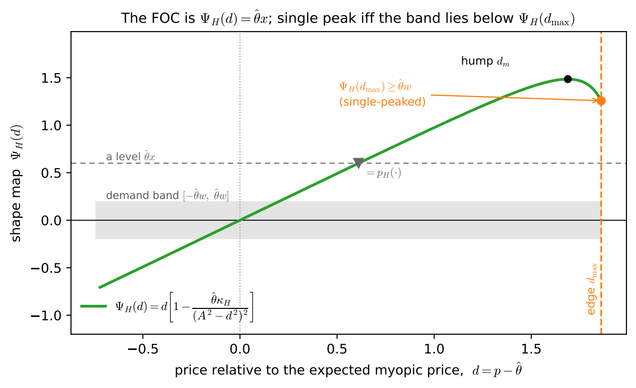

- Ψ_L increasing $\Rightarrow$ L's problem **globally concave**
- Ψ_H: rises, humps, dives at the review boundary $\Rightarrow$ Π_H *never* globally concave
- but under A1 and $\Psi_H(d_{\max}) \geq \hat\theta w$ $\Rightarrow$ level line crosses **once** $\Rightarrow$ single peak $=$ global max

The convexity at extreme prices is real but profit is lower so the firm never posts those prices.

Main theorem: the pricing functions

Theorem: Characterizing the Equilibrium Pricing Functions

(a) **Price is expected myopic at mean demand:** $p_H(1) = p_L(1) = \hat{\theta}$

(b) Both firms pricing function strictly increasing in η

(c) **Pricing sensitivity to demand ordered by quality:** $p'_H(\eta) \; > \; \hat{\theta} \; > \; p'_L(\eta)$

(d) **Direction of distortion:** $\operatorname{sgn}(p_H - \hat{\theta}\eta) = \operatorname{sgn}(\eta - 1)$, $\operatorname{sgn}(p_L - \hat{\theta}\eta) = -\operatorname{sgn}(\eta - 1)$

(e) **Ordering reversal:** $\boxed{\operatorname{sgn}\left(p_H(\eta) - p_L(\eta)\right) = \operatorname{sgn}(\eta - 1)}$

Proof discipline: Statements are consequences of the function $\Psi_k(d) = \hat{\theta}(\eta - 1)$ *and* $\Psi'_H < 1 < \Psi'_L$

Reading the theorem

(d) The deep discount is a quality flex, not a fire sale. In weak markets, the *good* firm undercuts myopic price — making the value-for-money test easy to ace, harvesting glowing reviews — while the *bad* firm holds price up to keep the test blurry.

(e) "Expensive = good" is a boom-time truth.

$$\eta > 1: \; p_H > p_L \qquad \eta = 1: \; p_H = p_L \qquad \eta < 1: \; p_H < p_L$$

The price–quality map depends on demand.
A pooled price-on-quality regression is not obvious here.

Pricing Functions

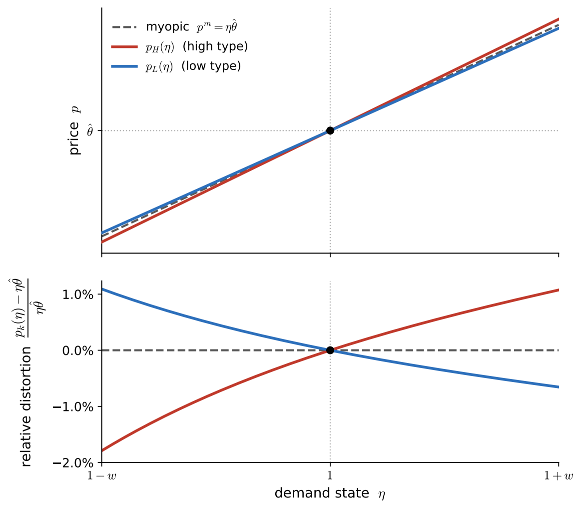

- Pricing distortion:
  a. $p_H(1) = p_L(1) = \hat{\theta}$
  b. $p'_H(\eta) > 0;\ p'_L(\eta) > 0$
  c. $p'_H(\eta) > \hat{\theta} > p'_L(\eta)$
  e. $\text{sgn}\left(p_H(\eta) - p_L(\eta)\right) = \text{sgn}(\eta - 1)$

No mean distortion

Proposition: Zero mean distortion

$$\mathbb{E}_\eta[p_H(\eta)] \;=\; \mathbb{E}_\eta[p_L(\eta)] \;=\; \hat{\theta} \;=\; \mathbb{E}_\eta[p^{\mathrm{m}}(\eta)]$$

$$\mathrm{Var}(p_H) \;>\; \mathrm{Var}(p^{\mathrm{m}}) \;>\; \mathrm{Var}(p_L)$$

Ratings management is a **informativeness phenomenon**:

- H goes for more extreme prices; L goes for expected price

Comparative statics: what strengthens the motive

Everything reads off the two constants in the pricing condition

$$\kappa_H = \beta\,\zeta\,\mu(1-\mu)^2\Delta_\theta^3, \qquad \kappa_L = \beta\,\zeta\,\mu^2(1-\mu)\Delta_\theta^3, \qquad \text{force} \;\propto\; \kappa_k\,/\,(A^2-d^2)^2$$

- **Patience** (β): more patient firms distort more; the slope gap $p'_H - p'_L$ widens
- **Experience noise** (A): noisier experiences blunt the dial; as $A \to \infty$ both types price myopically

  strongest in *sharp-review* categories (objective performance), weakest in taste-driven ones
- **The prior** (μ): $\kappa_H/\kappa_L = (1-\mu)/\mu$ — the *unexpected* type distorts hardest

  a rare high type strains to prove itself; a rare low type strains to hide
- **Quality gap** (Δ_θ): enters cubically — negligible in commodity categories, dominant under wide differentiation

Testable content

1. **Premise:** Good reviews become *less likely at higher prices paid*

   Well established

2. **No level effect — but variance ordered:** high- and low-quality sellers show the *same average price*, yet $\mathrm{Var}(p_H) > \mathrm{Var}(p_L)$

3. **Slope and sign reversal:** price–quality correlation *positive* in strong demand, *negative* in weak demand

4. **The dial itself:** a review moves subsequent demand *more*, the farther the price that generated it sits from "expected" price.

Data: Amazon (reviews $\times$ Keepa prices; for demand use exogenous events); Placebo: fixed-price categories (books under RPM) show nothing. Cross-category: sharp-review categories ($1/A^2$) amplify all effects.

Relation to the literature

Pricing to manage ratings (Carnehl–Stenzel–Schmidt, *Mgmt Sci* 2024; Li–Hitt, *MISQ* 2010; Crapis et al., *Mgmt Sci* 2017; Shin–Vaccari–Zeevi, *Mgmt Sci* 2022)

- the market reads a rating that *does not adjust for the price the reviewer paid*
  $\Rightarrow$ the seller can push its rating around

Price as a statement about quality (Wolinsky, *REStud* 1983; Milgrom–Roberts, *JPE* 1986; Bagwell–Riordan, *AER* 1991; Klein–Leffler, *JPE* 1981; Board–Meyer-ter-Vehn, *Ecta* 2013)

- buyers read the price *itself*

This paper: readers of reviews know what the reviewer paid

- price cannot push the review in the seller's favor
- price affects informativeness
- novel dynamics, relation price-quality depends on demand shocks

Conclusion

A model where the price is **the promise the product must keep**:

- Price cannot bias reviews — it only chooses how demanding a promise to make (an informativeness dial)
- The reputational motive is mean-preserving: no level distortion; ratings management impacts *dispersion*
- Quality is revealed through slopes and regime-contingent signs:
  $p'_H > \hat{\theta} > p'_L, \quad \text{and} \quad \operatorname{sgn}\left(p_H - p_L\right) = \operatorname{sgn}(\eta - 1)$
  expensive = good is a boom-time truth; the deep discount is a quality flex
- Testable in environments where levels regressions see nothing: equal means, ordered variances, and a sign flip

Thank you.