GeoProgress Journal, Vol. 12, Issue 2, 2025, Geoprogress Editions
DOI: 10.20373/2384-9398/72

PRICE–MATCHING GUARANTEES IN THE AGE OF ALGORITHMS: A MODEL OF COLLUSION, DISCRIMINATION AND THE WELFARE TRADE-OFF

Simone Robbiano†

Abstract

The paper revisits Price Matching Guarantees (PMGs) in the context of online oligopolies that monitor rivals' prices through real-time algorithms and infer consumer elasticities from User-Generated Contents (UGCs). Building on a dynamic repeated-game framework with homogeneous versus heterogeneous marginal costs, it is shown that PMGs sustain tacit collusion when costs are identical but generate algorithmic price discrimination when costs differ. Monte Carlo simulations confirm the analytical results and quantify the welfare effects under alternative regulatory regimes. While automatic refund mandates eliminate discriminatory rents without suppressing competitive pricing, blunt PMGs bans sacrifice legitimate insurance benefits for elastic consumers. The findings highlight the need for process-based rather than structural antitrust remedies in data-rich digital markets.

JEL Codes: D43; H89; K21; L13; L81. *Keywords*: Price–Matching Guarantees; Online Pricing; Oligopoly Models; Algorithmic Collusion; Price Discrimination; Welfare; Antitrust

1. Introduction

A *Price–Matching Guarantee* (PMG) is the seller's commitment to refund the difference if the buyer later finds the same item cheaper elsewhere. To the casual observer, the promise appears pro-competitive: it reassures consumers that they need not search further, and it prods rivals to keep prices low. However, early theoretical work reached the opposite conclusion. In particular, the two main anticompetitive theories view PMGs as mechanisms that facilitate collusion (see, among others, Hay 1981; Salop 1986; Belton 1987) or as instruments for price discrimination.

The first strand of the literature argues that, under specific conditions, PMGs can help sustain collusion in oligopolistic markets. These guarantees are interpreted as implicit threats of punishment for firms that undercut cartel prices, thereby reducing incentives to

† University of Genoa – Department of Economics. Via F. Vivaldi, 5 - Darsena - 16126 Genova.
✉ simone.robbiano@economia.unige.it. Declaration of interests: none. Funding sources to be declared: none.

deviate from collusive agreements.[1] Subsequent models (e.g. Edlin and Emch 1999; Jain and Srivastava 2000; Villas-Boas 2009) refine the mechanism, demonstrating that the effect hinges on search frictions, refund hassle costs, and firms' ability to monitor one another.

Conversely, several other scholars have challenged the idea that PMGs inevitably lead to collusion when buyers and retailers are heterogeneous. They contend that sellers can instead wield these guarantees to discriminate among customers. Theoretical contributions by Png and Hirshleifer (1987), Belton (1987), Corts (1997), and Nalca et al. (2010) show that firms can use PMGs to segment consumers who differ in terms of their access to price and warranty information, willingness to pay, brand loyalty, or the magnitude of hassle and search costs.[2]

A complementary strand studies PMGs as strategic commitments in differentiated retail settings and as signals of low prices. In particular, Moorthy and Winter (2006) and Moorthy and Zhang (2006) emphasise how guarantee design interacts with retail differentiation and consumer inference, while Mamadehussene (2019) formalises PMGs as a direct signal of low prices. These approaches are informative for our modelling choices: we retain the central role of refund frictions (hassle costs) as the behavioural wedge that gives PMGs bite, but we embed the guarantee in an environment where monitoring and inference are algorithmic and near real-time.

Given that earlier PMGs models focused on brick-and-mortar markets, the conventional wisdom predicts that digital markets, with near-zero search costs and instant price discovery, should render PMGs obsolete. Yet large platforms such as Amazon, Best Buy and Expedia continue to tout "price assurances." What keeps PMGs alive is the interplay between *online price trackers*—bots that scrape competitors every few seconds—*User-Generated Contents* (UGCs) such as price alerts on forums or social media, and sophisticated *data-analytics engines* that map clickstreams to individual elasticity profiles. The same technology that makes prices transparent also lets a retailer infer *who* is likely to file a claim, restoring the screening role of refund hassle in an otherwise transparent marketplace.

Recent contributions study algorithmic tacit collusion without PMGs or empirical PMG campaigns in e-commerce (among others, see Calvano et al. 2020), confirming that real-time monitoring accelerates price coordination and that refunds are claimed by a price-savvy minority, but leave open how PMGs interact with machine-learning (ML) pricing tools. Crucially, no formal model to date endogenises the joint choice of *price*, *PMG policy* and *data-driven segmentation*.

Relative to earlier PMG models developed for brick-and-mortar contexts, the paper's 'digital' contribution is to (i) treat observability of rivals' actions as effectively instantaneous, consistent with automated price-scraping; and (ii) treat refund propensity

---

[1] This pro-collusive interpretation has been further developed by studies that extend the basic oligopoly framework or apply spatial competition models such as Hotelling's (e.g., Logan and Lutter 1989; Baye and Kovenock 1994; Lu and Wright 2010; Hviid and Shaffer 2010; Pollak 2017; Constantinou and Bernhardt 2018; Cabral et al. 2018). Nonetheless, Hviid and Shaffer (1999) point out that the existence of hassle costs (i.e., the effort required to redeem the guarantee) can weaken the anticompetitive potential of PMGs.

[2] See also Edlin (1997).

as an inferred object, disciplined by user-generated content and platform-side behavioural traces, which allows PMGs to support either collusion (under homogeneous costs) or type-contingent effective pricing (under cost heterogeneity). This positioning also guides the simulation design: each counterfactual regime is chosen to map to an identifiable policy lever (PMG ban; margin cap; automatic refunds).

As pricing algorithms experiment and learn, the strategic landscape diverges from classic static or repeated games with imperfect monitoring. ML systems update beliefs about the elasticity distribution and adapt PMG–pricing menus in real time. Ignoring these feedback loops risks under- or over-estimating both collusive dangers and consumer benefits. Hence, the field requires models that embed data analytics—and the resulting asymmetric information—directly into firms' strategic choices. In particular, these mechanisms should be investigated:

- Real-time monitoring → public observability → trigger strategies become empirically plausible.
- Data/UGC inference → belief state / mapping to refund propensity.
- Automated processes (refund automation, APIs) → policy lever "auto-refund" has a concrete microfoundation.

The analysis is tailored to online retail environments where (i) prices are effectively public and continuously monitored (e.g. via price-scraping), making deviations and undercutting rapidly observable; (ii) PMG redemption is mediated by digital processes, so that claiming a refund involves a tractable "hassle" component (time, proof submission, procedural frictions) that can be reduced by platform design; and (iii) user-generated content and platform traces reveal systematic heterogeneity in refund propensity and price sensitivity, which can be exploited for segmentation. These digital primitives discipline the modelling choices below and motivate the policy counterfactuals (e.g. automatic refunds), which are less meaningful in offline settings.

To this end, the paper develops two dynamic oligopoly models. In Model 1 firms share a common marginal cost and may adopt a PMG. Perfect price observability creates a trigger strategy that supports supra-competitive prices *only* if refund hassle costs are sufficiently high. In Model 2 marginal costs differ; data analytics allow firms to infer individual elasticities, so a high-cost seller can profit from offering a PMG that is exercised mainly by highly elastic consumers. The resulting "separating" equilibrium implements algorithmic price discrimination. The paper characterises the equilibria analytically, derives welfare rankings and validates the magnitudes through large-scale Monte-Carlo simulations.

From a policy maker point of view, collusive PMGs in homogeneous-cost industries unambiguously harm consumer surplus and total welfare. In heterogeneous-cost settings PMGs raise *aggregate* welfare by reallocating demand to efficient low-cost firms, but they do so via discriminatory mark-ups on inelastic buyers, raising equity and antitrust concerns. Simulated counterfactuals show that an *auto-refund mandate* outperforms a blunt PMG ban: it preserves price insurance for savvy shoppers while eliminating discriminatory rents. Antitrust policy should therefore target the *process* (refund frictions and data opacity) rather than the PMG *instrument* per se, a nuance missed in much of the current regulatory debate.

The remainder of the paper proceeds as follows: Section 2 formalizes the oligopolistic pricing model for homogeneous goods in an online environment, while Section 3 provides Monte Carlo simulations. Section 4 concludes.

2. Model Formulation, Equilibrium Analysis, and Discussion

2.1 Overview and Economic Motivation

In this section, the paper formalizes an oligopolistic pricing model for homogeneous goods in an online environment where competing vendors may elect to offer PMGs. Firms observe competitors' posted prices and PMG commitments in real time, and consumers exhibit heterogeneous price elasticities and incur individual-specific hassle costs when requesting a PMG refund. Two variants of the model are analyzed: one in which all firms have identical marginal costs (Model 1), and another in which marginal costs differ across firms (Model 2). The infinite-horizon repeated-game structure permits the emergence of tacit collusion under suitable conditions. Finally, the paper discusses the market outcomes—prices, consumer surplus, and welfare—in each scenario, drawing on standard game-theoretic and industrial-organization arguments.

2.2 Environment and Notation

Let $\mathcal{N} = \{1, \dots, N\}$ denote the set of online vendors, where $N \geq 2$. Time is discrete and indexed by $t = 0,1,2, \dots$, and all players (firms and consumers) discount future payoffs by a common factor $\beta \in (0,1)$. In each period t:

- Each firm i simultaneously chooses a posted price $p_i^t \in \mathbb{R}_+$ and whether to offer a PMG, denoted by $g_i^t \in \{0,1\}$ (with $g_i^t = 1$ signifying that a PMG is in effect).
- Firms incur constant marginal cost c_i per unit sold; in Model 1, $c_i \equiv c$ for all i, whereas in Model 2 each firm's cost c_i is either c_L or c_H with $c_L < c_H$, and the distribution of types is common knowledge.
- A unit mass of identical-but-elastic consumers arrives, each characterized by an individual price elasticity parameter e. The elasticity e is drawn independently at the start of each period from a continuous distribution F on $[e_{\min}, e_{\max}]$ with density $f(e) > 0$ for $e \in (e_{\min}, e_{\max})$.
- If consumer e purchases from firm i at price p_i^t with a PMG ($g_i^t = 1$), and a rival's price $\hat{p}^{\,t} = \min_{j \neq i} p_j^t$ is strictly lower, then the consumer can request a refund of $(p_i^t - \hat{p}^{\,t})$ by incurring a hassle cost $h(e)$, where $h: [e_{\min}, e_{\max}] \to \mathbb{R}_+$ is strictly decreasing and continuously differentiable, satisfying $h'(e) < 0$. Thus, higher-elasticity consumers face smaller hassle costs in claiming refunds. If $g_i^t = 0$, no refund is available.
- Consumers maximize quasilinear net utility. Purchasing at posted price p_i from a firm that offers a PMG yields

$$U(e, p_i, g_i = 1) \;=\; v_0 - p_i \;+\; \max_{\hat{p} \leq p_i}\{ e\,(p_i - \hat{p}) - h(e)\},$$

where v_0 is a normalized gross surplus (assumed large enough to ensure participation whenever $p_i \leq v_0$).[3] If no PMG is offered ($g_i = 0$), then

$$U(e,\, p_i,\, g_i = 0) \;=\; v_0 - p_i,$$

independent of rivals' prices. Consumers who obtain nonpositive net utility abstain from purchase.

Remark on Demand Specification. Throughout, a logit-style demand specification contingent on "hassle-adjusted utilities" is adopted. In particular, define for firm j in period t and consumer elasticity e:

$$\widetilde{U}_j(e; p_j^t, g_j^t, \hat{p}^{\,t}) \;=\; \begin{cases} \exp\big(e\,(\hat{p}^{\,t} - p_j^t)\big), & \text{if } g_j^t = 1, \\ \exp\big(e\,(\hat{p}^{\,t} - p_j^t - h(e))\big), & \text{if } g_j^t = 0. \end{cases}$$

Hence, firm j's probability of being chosen by an e-type consumer (conditional on purchase) is

$$\Pr\{j \mid e,\, \mathbf{p}^t,\, \mathbf{g}^t\} \;=\; \frac{\widetilde{U}_j(e; p_j^t, g_j^t, \hat{p}^{\,t})}{\sum_{k=1}^{N} \widetilde{U}_k\,(e; p_k^t, g_k^t, \hat{p}^{\,t})}.$$

Define the mass of consumers who actually buy from some vendor as

$$D(\mathbf{p}^t,\, \mathbf{g}^t) \;=\; \int_{e_{\min}}^{e_{\max}} \left(\sum_{j=1}^{N} \widetilde{U}_j\,(e; p_j^t, g_j^t, \hat{p}^{\,t}) \right) f(e)\, de,$$

so that the demand shifter $D(\cdot)$ captures the total intensity of consumer participation. Consequently, firm i's period-t quantity sold is

$$q_i^t \;=\; \int_{e_{\min}}^{e_{\max}} \frac{\widetilde{U}_i(e; p_i^t, g_i^t, \hat{p}^{\,t})}{\sum_{k=1}^{N} \widetilde{U}_k\,(e; p_k^t, g_k^t, \hat{p}^{\,t})}\; f(e)\, de,$$

and its one-period profit is

$$\pi_i^t \;=\; (p_i^t - c_i)\, q_i^t.$$

2.3 Players' Information and Strategy Spaces

All firms have perfect and complete information about each other's marginal costs and publicly observe the entire history of price/PMG choices. Consumer types e are privately

[3] The max operator captures the consumer's endogenous refund-claim decision: when a firm is undercut, a buyer claims the price match only if the refund gain exceeds hassle. Accordingly, hassle costs matter only when a PMG is offered; if $g_i = 0$ no refund option exists. Finally, the logit formulation can be read as conditional-on-purchase (as written) or equivalently as a standard logit with an outside option normalised to zero.

realized but i.i.d. across consumers and periods. At the beginning of period t, a firm observes:

1. The public history $h^t = \{(p_j^\tau,\, g_j^\tau)\}_{j\in\mathcal{N},\,\tau=0,\ldots,t-1}$.
2. A public signal—extracted from UGCs and browsing logs—summarizing the posterior distribution (belief) θ_t over the elasticity distribution F. By assumption, θ_t evolves deterministically given past prices and PMG offers. In particular, the latter is updated by Bayes' rule after observing the realised signal; conditional on the realised signal, the update is deterministic.

Operationally, the signal aggregates digital traces that correlate with price sensitivity and refund propensity: (i) UGC containing explicit 'price-match' language (e.g. reviews describing claims, screenshots, or cross-retailer comparisons); (ii) platform-side outcomes such as the frequency and timing of refund requests following observed competitor undercutting; and (iii) browsing patterns consistent with intensive comparison (repeat visits, referral traffic from deal forums, use of price-alert extensions). The paper does not assume firms observe individual elasticities; rather, these observables discipline a common posterior over the elasticity distribution (or, in the segmentation interpretation, over the mapping from observables to refund propensity), which is the sufficient statistic entering firms' pricing/PMG incentives.
Therefore, a *public strategy* for firm i is a mapping[4]

$$\sigma_i\colon (h^t,\, \theta_t) \;\mapsto\; (p_i^t,\, g_i^t).$$

The paper restricts attention to *Markov strategies* in the augmented state (h^t, θ_t), since future payoffs depend only on the current public belief θ_t and the most recent price/PMG profile, due to the memoryless updating of beliefs and the logit-demand assumption. In Sections 2.5–2.6 the analysis can be read as the complete-information benchmark (degenerate beliefs), in which case strategies need not be belief-based. Beliefs are only payoff-relevant in the learning extension (Section 2.7).

2.4 Equilibrium Concept

For the baseline complete-information environment, the equilibrium notion can be read as a standard subgame-perfect equilibrium in public strategies. When beliefs are degenerate (i.e. primitives are know, the Perfect Bayesian Equilibrium (PBE) reduces to the standard repeated-game equilibrium concept. The paper seeks a PBE in Markov strategies, i.e., a profile $(\sigma_i)_{i=1}^N$ together with consistent beliefs over off-equilibrium-path posteriors, such that:

- Each σ_i maximizes firm i's expected discounted payoff

[4] In period firms post p and choose to offer PMG. Then, consumers purchase, competitor prices are observed (bots). Finally, some consumers claim refunds, and realised claim rate is observed. Consequently, beliefs are updated.

$$\Pi_i = \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \ \pi_i^t\right]$$

given rival strategies and the belief-updating rule $\theta_{t+1} = \Lambda(\theta_t, \mathbf{p}^t, \mathbf{g}^t)$.

- Beliefs θ_t are updated according to Bayes' rule on the equilibrium path.

2.5 Model 1: Homogeneous Marginal Costs ($c_i \equiv c$)

2.5.1 Static One-Shot Correction with PMGs

When all firms have identical marginal costs c, the static (one-shot) stage takes the form of a simultaneous move game in (p_i, g_i) for $i = 1, \ldots, N$. Suppose rivals' prices and PMGs are $(\mathbf{p}_{-i}, \mathbf{g}_{-i})$. Define $\hat{p}_{-i} = \min_{j \neq i} p_j$. Observe that if firm i sets $p_i > \hat{p}_{-i}$ and $g_i = 1$, it will attract all elasticity-types e for which

$$v_0 - p_i + \max\{\, e\,(p_i - \hat{p}_{-i}\,) - h(e)\} \ \geq \ \max\{\, v_0 - \hat{p}_{-i}, \, 0\}.$$

Since $h'(e) < 0$, more elastic consumers (large e) find the net gain $\max\{e(p_i - \hat{p}_{-i}) - h(e)\}$ positive. However, the mass of such consumers depends on the public belief F over e. In particular, if $p_i < \hat{p}_{-i}$, firm i trivially attracts all purchasers at price p_i (with or without PMG). If $p_i = \hat{p}_{-i}$ but $g_i = 1$ while all rivals choose $g_j = 0$, then i fully captures the market (zero refund necessary, by tie-breaking). If $g_i = 0$, no refund is paid and the allocation is symmetric among equal prices.
The static best-response condition can be summarized as follows:

- If no rival offers a PMG at price $\hat{p}$, then setting $g_i = 0$ and undercutting marginally (to $\hat{p} - \varepsilon$) captures the entire market at price slightly below $\hat{p}$. Hence, price competition drives $p_i \to c$ in any Nash equilibrium with $g_i = 0$ for all i.
- If exactly one firm, say k, offers a PMG at $\hat{p}$ and others do not, then any deviator $i \neq k$ can set $p_i^d < \hat{p}$ without a PMG and attract all consumers of low elasticity who find $h(e)$ too large. Given the demand specification, there may be a mixed-strategy equilibrium in the one-shot game. However, under mild conditions (e.g., sufficient mass of low-elasticity types), the unique pure-strategy Nash equilibrium is $p_i = c, \ g_i = 0$ for all i, with zero profits.
- If all firms choose $(p^*, g = 1)$ with $p^* > c$, no one can profitably deviate by dropping $g = 1$ (since that would hurt high-elasticity consumers), nor by undercutting marginally. Thus, when $h(e)$ is sufficiently large, it may be the case that a pure-strategy Nash equilibrium of the static game features $p_i = p^*, \ g_i = 1$ for all i. But in fact, one can show that for any $p^* > c$, the only static equilibrium is $p_i = c, \ g_i = 0$ in the one-shot game. Intuitively, as long as $p^* > c$, a unilateral undercutting yields a strictly higher profit in the one-shot stage.

2.5.2 Repeated-Game Collusion with PMGs

Since the static game with homogeneous costs admits only the Bertrand outcome ($p_i = c,\ g_i = 0$), supra-competitive outcomes must be sustained by repeated-game incentives. Consider the following *grim–trigger-style* strategy profile:

Strategy for each firm i:

- At $t = 0$, set $p_i^0 = p^* > c$ and $g_i^0 = 1$.
- In period $t \geq 1$, if in every past period $\tau < t$ all firms played $(p_j^\tau, g_j^\tau) = (p^*, 1)$, then choose $(p_i^t, g_i^t) = (p^*, 1)$.
- Otherwise (if any deviation has occurred), revert to $(p_i^\tau, g_i^\tau) = (c, 0)$ forever.

Denote by p^* the candidate collusive price. The paper shows:

Theorem 2.1 *(Collusive PMG Equilibrium: Homogeneous Costs). Assume $h(e) \geq \bar{h} > 0$ for all e, where $\bar{h}$ satisfies $h'(\bar{e}) = -1$ for the mean elasticity $\bar{e} = \int_{e_{min}}^{e_{max}} e\ f(e)\,de$. Then, for any p^* satisfying*

$$c \ < \ p^* \ \leq \ p^{mon}, \tag{1}$$

there exists $\underline{\beta} \in (0,1)$ such that if $\beta \geq \underline{\beta}$, the grim-trigger profile constitutes a Perfect Bayesian Equilibrium that sustains $(p_i^t, g_i^t) = (p^*, 1)$ for all i and all t. Here, p^{mon} is the static monopoly price (computed with elasticity distribution F).

Proof. **One-Shot Deviation Analysis.** Suppose all firms are playing $(p^*, 1)$ at history h^{t-1}, and consider firm i's one-shot deviation at period t. If i chooses a price $p_i^d < p^*$ and *drops* the PMG ($g_i^t = 0$), then two effects arise:

1. *Immediate Profit Gain*: By undercutting p^*, firm i attracts all consumers whose net utility of claiming a refund from a rival is lower than the direct saving. Because $h(e) \geq \bar{h}$, only consumers with sufficiently low elasticity would find it unprofitable to claim a refund. Under logit demand, i's immediate quantity is

$$q_i^t \approx \int_{e:e(p^*-p_i^d)\leq h(e)} \frac{\exp\big(e(p^* - p_i^d - h(e))\big)}{\exp\big(e(p^* - p_i^d - h(e))\big) \ + \ (N-1)\exp(e \cdot 0)}\, f(e)\,de,$$

   which, for small $\varepsilon = p^* - p_i^d > 0$, gives $q_i^t \approx 1$. Thus the one-period deviation payoff is $\pi_i^d \approx (p_i^d - c) \cdot 1$. In particular, the optimum one-shot deviation price p_i^d solves

$$\max_{p<p^*} (p - c)\ \Pr\{\text{buyer does not claim refund}\},$$

   which is strictly positive if $c < p^*$.

2. *Future Profit Loss*: Because the deviation is publicly observed (prices and g are transparent), all firms revert to ($p = c,\ g = 0$) in all subsequent periods, yielding zero profits thereafter.

Hence the net present value of deviating at t is

$$\pi_i^d \;=\; (p_i^d - c) \;+\; \sum_{\tau=1}^{\infty} \beta^\tau \, 0 \;=\; p_i^d - c.$$

Continuing with the cooperative plan yields for firm i the continuation payoff

$$V_i^{\text{coop}} \;=\; \sum_{\tau=0}^{\infty} \beta^\tau \, (p^* - c) \, q_{\text{coll}} \;=\; \frac{p^* - c}{1 - \beta},$$

where q_{coll} is the collusive quantity at price p^* (which is strictly positive under logit-type demand). In particular, the stage-game profit at collusion is $(p^* - c)\, q_{\text{coll}} > 0$. Therefore, no profitable one-shot deviation exists if and only if

$$p^d - c \;\leq\; (1 - \beta)\,(p^* - c)\, q_{\text{coll}}.$$

Let $p_{\max}^d$ denote the best-response deviation price (maximizing one-period profit $(p - c)\, q^d(p)$). Then the incentive-compatibility condition is

$$p_{\max}^d - c \;\leq\; (1 - \beta)\,(p^* - c)\, q_{\text{coll}}, \quad \Leftrightarrow \quad \beta \;\geq\; 1 - \frac{p_{\max}^d - c}{(p^* - c)\, q_{\text{coll}}}.$$

Since q_{coll} is uniformly bounded away from zero when $h(e) \geq \bar{h}$ (elastic consumers find refunds unattractive), there exists $\underline{\beta} < 1$ satisfying this inequality. Hence for all $\beta \geq \underline{\beta}$, the grim-trigger strategy profile is a PBE sustaining $(p^*, 1)$ indefinitely.

Economic Interpretation. When the hassle function $h(e)$ is uniformly large, even highly elastic consumers face significant disutility from requesting a refund, so the attractiveness of deviation is limited. By promising to punish deviators with a reversion to aggressive (Bertrand) pricing at $p = c$, the grim-trigger strategy extracts a credible threat strong enough to sustain supra-competitive pricing at $p^* > c$. In particular, the requirement $p^* \leq p^{\text{mon}}$ guarantees that the collusive profit per period, $(p^* - c)\, q_{\text{coll}}$, exceeds the maximum one-shot deviation gain.

2.6 Model 2: Heterogeneous Marginal Costs ($c_i \in \{c_L, c_H\}$)

2.6.1 Static Benchmark and Incentives

When firms' marginal costs differ, the one-shot static game has the following well-known feature: if $c_L < c_H$, then any pure-strategy Nash equilibrium must satisfy that low-cost firms drive p down to just below c_H:

$$p_i \to c_H - \varepsilon, \quad \text{for each low-cost firm } (c_i = c_L).$$

At such prices, high-cost firms cannot profitably match (as $c_H - \varepsilon < c_H$ implies negative profit), so they either price at c_H and sell zero, or randomize but still obtain zero profit in expectation. Hence, in the static stage game all consumer demand flows to the low-cost firms.
Allowing PMGs complicates the picture, because a high-cost firm could offer a PMG at a higher price, hoping to extract surplus from low-elasticity buyers. Concretely, suppose low-cost firm L posts $(p_L, g_L = 0)$ and high-cost firm H posts $(p_H, g_H = 1)$ with $p_H > p_L$. A consumer of elasticity e who buys from H pays p_H but can request a refund if $p_L < p_H$, netting $p_H - p_L$ minus the hassle cost $h(e)$. Thus, H effectively segments the market into two sub-populations:

- *Inelastic consumers* (with e below some threshold $\hat{e}$) who find $e(p_H - p_L) < h(e)$; they pay p_H and do not request a refund. These consumers have lower willingness to monitor price differences.
- *Elastic consumers* (with $e \geq \hat{e}$) who find $e(p_H - p_L) \geq h(e)$; they pay p_H initially but then claim a refund, effectively paying p_L. These consumers exhibit higher vigilance regarding price discrepancies.

This segmentation relies on observable consumer attributes (elasticities) and exploits the fact that claiming a refund imposes heterogeneous costs $h(e)$ across the population. From an antitrust perspective, deliberately differentiating prices in this way—charging inelastic consumers a higher final price than elastic ones—raises serious competition law concerns. In many jurisdictions, antitrust authorities scrutinize such practices closely, as they can undermine competitive neutrality and exploit consumer heterogeneity to the detriment of certain segments.

2.6.2 Repeated-Game "Separation" with Antitrust Concerns

The paper now embeds this static incentive structure into the infinite-horizon repeated-game context, but rewriting the equilibrium argument to highlight the antitrust implications of targeted price discrimination.

Candidate Strategy Profile. *Fix prices* (p_L^*, p_H^*) *satisfying*

$$p_L^* > c_L, \quad p_H^* > c_H, \quad p_H^* - p_L^* = \delta, \quad 0 < \delta < \min_{e \in [e_{\min}, e_{\max}]} h(e), \qquad (2)$$

so that a consumer of elasticity e will request a refund from H if and only if $e\,\delta \geq h(e)$. Denote by $\hat{e}$ the unique solution to $e\,\delta = h(e)$ (existence follows from continuity and strict monotonicity of $h(e)$). Thus:

$$\begin{cases} e < \hat{e} & \Rightarrow \text{ consumer pays } p_H^* \text{ at } H, \\ e \geq \hat{e} & \Rightarrow \text{ consumer pays } p_L^* \text{ (via refund) at } H. \end{cases}$$

Define the following repeated-game "separation" strategy profile:

Strategy for low-cost firm *L*:

- At $t = 0$, choose $(p_L^0, g_L^0) = (p_L^*, 0)$.
- In period $t \geq 1$, if no deviation has occurred in history, continue with $(p_L^t, g_L^t) = (p_L^*, 0)$. If any deviation is detected, revert to $(p_L^\tau, g_L^\tau) = (c_L, 0)$ forever.

Strategy for high-cost firm *H*:

- At $t = 0$, choose $(p_H^0, g_H^0) = (p_H^*, 1)$.

In period $t \geq 1$, if no deviation has occurred, continue with $(p_H^t, g_H^t) = (p_H^*, 1)$. If any deviation is detected, revert to $(p_H^\tau, g_H^\tau) = (c_H, 0)$ forever.

Theorem 2.2 ("Separation" Equilibrium Under Heterogeneous Costs—Antitrust Warning). *Suppose $h(e)$ is strictly decreasing and continuous. Let $\delta = p_H^* - p_L^*$ satisfying Eq. (2), and let $\hat{e}$ be the threshold solving $\hat{e}\,\delta = h(\hat{e})$. If β is sufficiently close to 1, then the above strategies constitute a Perfect Bayesian Equilibrium.*
In equilibrium:

- *All consumers with $e < \hat{e}$ pay p_H^* at H (no refund), while those with $e \geq \hat{e}$ pay p_L^* after claiming the refund from H.*
- *Low-cost firms earn per-period profit $(p_L^* - c_L)\, q_L^*$, where*

$$q_L^* = \int_{e_{min}}^{\hat{e}} 1\ f(e)\, de\ +\ \int_{\hat{e}}^{e_{max}} \alpha\,(e)\, f(e)\, de,$$

  and $\alpha(e) \in [0,1]$ is the tie-breaking fraction of the elastic segment that shops at L.
- *High-cost firms earn per-period profit $(p_H^* - c_H)\, q_H^{inel}$, where*

$$q_H^{inel} = \int_{e_{min}}^{\hat{e}} 0\ \ f(e)\, de\ +\ \int_{\hat{e}}^{e_{max}} (1 - \alpha(e))\ f(e)\, de\ =\ 1 - q_L^*.$$

- *No firm has a profitable one-shot deviation, because any deviation triggers reversion to zero-profit pricing.*

Proof. **(i) Static Best Responses.** Under $(p_L^*, 0)$ for low-cost and $(p_H^*, 1)$ for high-cost, a consumer of type e faces:

$$\text{Effective Price at } L = p_L^*, \quad \text{Effective Price at } H = \begin{cases} p_H^*, & e < \hat{e}, \\ p_L^*, & e \geq \hat{e}. \end{cases}$$

Thus:

- Inelastic types ($e < \hat{e}$) strictly prefer L (paying p_L^*) over H (paying $p_H^* > p_L^*$).
- Elastic types ($e \geq \hat{e}$) are indifferent between L and H (both effectively charge p_L^*). By assumption, a fraction $\alpha(e)$ of these types shop at L and $1 - \alpha(e)$ at H.

Hence, in equilibrium low-cost firms serve all inelastic types plus a fraction of elastic types, while high-cost firms serve only the remaining fraction of elastic types. There is no static incentive for either type to deviate: any attempt by a low-cost firm to raise p_L above p_L^* would lose inelastic consumers, reducing immediate profit; any attempt to lower p_L below p_L^* yields a lower margin net of cost. A high-cost firm cannot profitably deviate to $(p_{H'}, g_{H'})$ because any $p_{H'} > p_H^*$ reduces the inelastic segment further, and any $p_{H'} < p_H^*$ yields a lower margin.

(ii) Incentives in the Repeated Game. Suppose a high-cost firm deviates at period t to some $(\tilde{p}_H, \tilde{g}_H)$ that might attract additional inelastic or elastic types in that period. The maximum one-shot gain is

$$\max_{(p,g)}(p - c_H)\, q_H(p, g; p_L^*),$$

which is finite. But from $t + 1$ onward, everyone reverts to $(c_H,\, 0)$, yielding zero profit. The present value of deviation is thus at most

$$(p^\dagger - c_H)\, q^\dagger,$$

where $(p^\dagger, q^\dagger)$ is the static best-response profit. By contrast, sticking to $(p_H^*,\, 1)$ forever yields

$$\frac{(p_H^* - c_H)\, q_H^{\text{inel}}}{1 - \beta}.$$

No deviation is profitable if

$$(p^\dagger - c_H)\, q^\dagger \;\le\; (1 - \beta)\,(p_H^* - c_H)\, q_H^{\text{inel}}, \quad \Leftrightarrow \quad \beta \;\ge\; 1 - \frac{(p^\dagger - c_H)\, q^\dagger}{(p_H^* - c_H)\, q_H^{\text{inel}}}.$$

Since $q_H^{\text{inel}} > 0$ and $q^\dagger$ is bounded, there exists β close to 1 such that this holds.

A symmetric argument applies to low-cost firms: deviating from $(p_L^*,\, 0)$ yields a one-shot gain $(p_L^\dagger - c_L)\, q_L^\dagger$ but triggers zero-profit reversion. Sticking with $(p_L^*,\, 0)$ yields

$$\frac{(p_L^* - c_L)\, q_L^*}{1 - \beta}.$$

For sufficiently large β, no deviation is profitable. Hence the profile is a PBE.

Antitrust Commentary on Discrimination. The equilibrium described above entails explicit price differentiation across consumer segments: inelastic consumers pay a higher effective price (p_H^*) than elastic consumers (p_L^*). From an antitrust standpoint, this practice raises two major concerns:

1. **Targeted Exploitation of Consumer Heterogeneity**. By using PMGs to infer and segment consumers according to their privately held elasticity parameter e, firms exploit vulnerability of inelastic buyers who incur higher hassle costs. Antitrust enforcers generally view such targeted pricing with suspicion, particularly if it results in systematically higher prices for a protected or economically disadvantaged group.
2. **Potential Exclusion of Competitors**. By leveraging PMGs to extract surplus from inelastic consumers, high-cost firms can sustain positive profits that might prevent entry by smaller or newer competitors who cannot replicate the same level of consumer-type inference.

Therefore, while the repeated-game logic supports the existence of such an equilibrium, in practice it would draw antitrust scrutiny, and regulators may require firms to limit or eliminate PMG-based conditional pricing to protect consumer welfare and maintain a level competitive field.

2.7 Dynamic Learning of Elasticity Distribution

This section relaxes complete information by allowing firms to learn about demand heterogeneity from a public signal. The baseline analysis corresponds to the limiting case of a degenerate prior (known primitives), in which case beliefs are degenerate and learning can be ignored.
Indeed, in both Model 1 and Model 2, firms' decisions in period t depend on the current public belief θ_t about the demand distribution F. The sequence $\{\theta_t\}$ evolves according to a deterministic mapping $\theta_{t+1} = \Lambda(\theta_t,\, \mathbf{p}^t,\, \mathbf{g}^t)$, capturing the firms' ability to infer realized draws of consumer elasticities from user-generated reviews, click-through data, and refund-request statistics. Because θ_t enters the demand functions as a parameter, the augmented state $(h^t,\, \theta_t)$ is Markov, and the equilibrium characterizations above extend by replacing F with the updated posterior F_{θ_t}. Thus, any history-dependent strategy can be rewritten as a Markov strategy in (h^t, θ_t), and the trigger-type constructions remain valid.

2.8 Discussion of Market Outcomes and Welfare Implications

The paper now compares the key equilibrium outcomes—prices, consumer surplus, firm profits, and total welfare—across the two models under reasonable parameter regimes, with special attention to antitrust considerations regarding discriminatory pricing in Model 2.

2.8.1Model 1 (Homogeneous Costs): Collusion vs. Competition

Bertrand Benchmark.
When no firm offers a PMG (or equivalently, when hassle costs are negligible so that $g_i = 1$ is effectively worthless), the only static Nash equilibrium is $(p_i = c,\, g_i = 0)$ for all i.

Each firm earns zero profit, consumer surplus is maximized given full extraction of surplus, and total welfare equals $(v_0 - c)$ (normalized).

Collusive Equilibrium with PMGs.
If $h(e)$ is sufficiently large, the repeated-game collusive equilibrium $(p_i^t, p_i^t) = (p^*, 1)$ yields each firm a steady-state flow profit $(p^* - c)\, q_{\text{coll}}$. Consumer surplus is strictly lower than under Bertrand because:

$$\text{CS}_{\text{coll}} \;=\; \int_{e_{\min}}^{e_{\max}} [v_0 - p^* + \max\{\,0,\, e\,(p^* - p^*) - h(e)\}]\, f(e)\, de \;=\; (v_0 - p^*)\cdot 1,$$

which is less than $(v_0 - c)$ whenever $p^* > c$. The deadweight loss relative to social optimum is proportional to $\frac{1}{2}(p^* - c)^2 f(\bar{e})$ in a first-order approximation around $\bar{e}$ (the mean elasticity). Total welfare (TW = CS + PS) is strictly lower under collusion, since the collusive margin $(p^* - c)$ creates a deadweight loss that exceeds any fixed-cost savings (which we assume zero).

Welfare Comparison.

$$\text{TW}_{\text{Bertrand}} = \underbrace{(v_0 - c)\cdot 1}_{\text{CS}} + 0 > \underbrace{(v_0 - p^*)\cdot 1}_{\text{CS}_{\text{coll}}} + \underbrace{(p^* - c)\cdot 1}_{\text{PS}_{\text{coll}}} = \text{TW}_{\text{coll}}.$$

Hence, despite the presence of PMGs, welfare is highest under aggressive pricing (no PMG) and lowest under collusive PMGs.

Threshold for Tacit Collusion. Condition (1) and the incentive-compatibility bound on β imply that collusion is harder to sustain (i.e., $\underline{\beta}$ is larger) when:

(i) p^* is far above c,

(ii) q_{coll} is small (scarce low-elasticity mass),

(iii) $h(e)$ declines sharply in e.

Thus, if the elasticity distribution places more weight on high-elasticity consumers (so few consumers refrain from refunds), the collusive equilibrium requires an extremely patient industry ($\beta \approx 1$).

2.8.2 Model 2 (Heterogeneous Costs): Discrimination Concerns

Static Outcome Without PMGs.
In the static game when PMGs are disallowed ($g_i = 0$ for all i), low-cost firms set $p \to c_H - \varepsilon$, capturing the market; high-cost firms set $p_H = c_H$ and sell zero. Consumer surplus equals

$$\mathrm{CS}_{\mathrm{static}} = \int_{e_{\min}}^{e_{\max}} (v_0 - (c_H - \varepsilon))\, f(e)\, de \approx (v_0 - c_H) \cdot 1.$$

Low-cost firms earn zero profit at $p_H = c_H$ for any $\varepsilon \downarrow 0$. Thus, the static industry profit is zero, and total welfare equals the (minimal) consumer surplus above.

Repeated "Separation" with PMGs—Antitrust Perspective.
In the repeated-game "separation" equilibrium of Theorem 2, firms systematically charge different final prices to inelastic versus elastic consumers, effectively imposing a higher price on those who are less able to monitor or bear the hassle of refunds. From an antitrust standpoint:

- **Consumer Harm via Unobservable Discrimination**. Inelastic consumers may pay a premium $p_H^* - p_L^*$ solely because their private elasticity e makes them less likely to claim a refund. This differential pricing cannot be observed ex ante by consumers, and it penalizes a vulnerable segment without transparent justification.
- **Risks of Market Foreclosure**. By leveraging PMGs to extract surplus from inelastic consumers, high-cost firms can sustain positive profits that might prevent entry by smaller or newer competitors who cannot replicate the same level of consumer-type inference.
- **Regulatory Red Flags**. Competition authorities in many jurisdictions regard such targeted discrimination as potentially abusive, especially if it leads to systematic overcharging of a protected or economically disadvantaged group (e.g., less-informed or lower-income consumers).

Hence, while the repeated-game logic predicts that such an equilibrium can be sustained for sufficiently patient firms, in practice it would draw antitrust scrutiny, and regulators may require firms to limit or eliminate PMG-based conditional pricing to protect consumer welfare and maintain a level competitive field.

Welfare Implications.

- **Consumer Surplus**: All consumers effectively pay p_L^*, but only elastic consumers receive it through refunds. Inelastic consumers pay p_H^* out of pocket. Thus, the average consumer surplus falls relative to the static non-PMG outcome if the inelastic segment is large enough.
- **Producer Surplus**: Low-cost firms capture inelastic revenue $(p_L^* - c_L)$ from a fraction of inelastic consumers plus some share of elastic demand; high-cost firms capture $(p_H^* - c_H)$ from inelastic buyers who mistakenly do not claim refunds. Total industry profit may exceed zero, but at the cost of consumer exploitation.
- **Total Welfare**: While industry profit increases, the loss in consumer surplus—especially concentrated on inelastic consumers—can outweigh the gain in producer surplus. Moreover, the welfare gain is not Pareto-improving since discrimination

harms a subset of consumers. Under standard antitrust frameworks, such a redistribution from vulnerable consumers to producers is viewed as welfare-reducing, warranting prohibition or strict regulation.

2.8.3 Parameter Regimes and Policy Insights

Role of Cost Gap $\boldsymbol{\Delta}$. When Δ is small, the discrimination wedge $\delta = p_H^* - p_L^*$ can also be small, limiting the overcharge on inelastics. As Δ grows, firms may choose larger δ, intensifying discriminatory overcharging. Competition authorities would be particularly concerned if δ significantly exceeds empirical measures of consumer monitoring costs $h(e)$.

Effect of Hassle Cost Function $\boldsymbol{h(e)}$. If $h(e)$ is uniformly large, fewer consumers undertake the refund process, shrinking $\hat{e}$ and enlarging the set of inelastic consumers subject to overcharging. Conversely, if $h(e)$ is uniformly small, most consumers claim refunds, making the discrimination wedge δ ineffective; in that case, PMGs have no discriminatory power, and static competition prevails.

Implications for Antitrust and Regulation. The analysis yields a clear caution:

- In markets with significant cost heterogeneity, PMGs can be employed as a mechanism for direct price discrimination that exploits less-elastic consumers. Antitrust authorities will likely intervene to prohibit or strictly regulate PMG-based pricing rules that differentiate based on inferred elasticity.
- If regulators permit PMGs, they may require transparent, uniform terms—e.g., specifying a standardized refund procedure and timeline—to reduce the friction that facilitates covert discrimination.
- Blanket bans on PMGs risk eliminating potential legitimate uses (e.g., simple price assurance for all consumers), but targeted restrictions—such as prohibiting differential refunds conditional on consumer characteristics—can mitigate antitrust concerns while preserving some consumer benefits.

2.9 Summary of Key Comparative Statics

1. **Discount Factor ($\boldsymbol{\beta}$)**. Higher β relaxes the incentive-compatibility constraints in both models, expanding the set of collusive or discriminatory equilibria. When β is sufficiently low, only static Bertrand outcomes (no PMG, $p = c_i$) survive in equilibrium.
2. **Hassle-Cost Profile ($\boldsymbol{h(e)}$)**.
    a. If $h(e)$ is uniformly large, PMGs facilitate supra-competitive collusion in Model 1 and discriminatory overcharging in Model 2.

   b. If $h(e)$ is uniformly small, PMGs are never credible (elastic consumers always claim refunds), so equilibrium reduces to static competition at $p = c_i$ (both models).
   c. If $h(e)$ is moderate, Model 1 collusion is more fragile, while Model 2 can yield discriminatory pricing that is welfare-reducing for inelastic consumers.
3. **Cost Dispersion (Δ)**.
   a. When $\Delta = 0$ (homogeneous costs), only collusive or Bertrand equilibria exist, and PMGs are either collusive devices or irrelevant.
   b. When $0 < \Delta < \Delta^*$ (intermediate range), Model 2's discriminatory equilibrium can arise, but it is likely to trigger antitrust concerns and may be prohibited.
   c. When Δ is very large, PMGs cannot effectively screen; the outcome converges to static competition with high-cost firms excluded.
4. **Number of Firms (N)**. Larger N intensifies competition. In Model 1, sustaining collusion requires β closer to 1 as N increases (the temptation to deviate grows). In Model 2, a greater number of low-cost firms (λN large) increases competitive pressure on p_L^*, reducing the scope for high-cost firms to exploit discrimination.

In summary, PMGs have dual potential: they can facilitate tacit collusion when costs are homogeneous and hassle costs high, or enable price discrimination that harms less-elastic consumers when costs are heterogeneous. From an antitrust standpoint, discriminatory uses of PMGs merit close scrutiny or prohibition, while reasonable safeguards could mitigate the collusion risk without eliminating consumer protection features of PMGs.

3. Monte–Carlo Simulations

This section reports a series of large-scale simulations that put numerical magnitudes on the analytical mechanisms. Readers interested only in the headline numbers can consult Table 1; the complete data set and the annotated Python code (200 MB) are available on request.

3.1 Experimental design

1. **Timeline**. $T = 1{,}000$ independent periods are simulated for each scenario. A period corresponds to one business day in a fast-moving online market.
2. **Consumers**. Each period hosts $M = 10{,}000$ buyers with elasticity $e \sim \mathcal{U}[0.5,5]$. The gross valuation is normalised to $v_0 = 2$. Hassle costs follow $h(e) = h_0/e$ with a baseline scale $h_0 = 0.5$.

3. **Firms**. Four symmetric vendors in Model 1 ($c_i = 1$); in Model 2 two vendors have $c_L = 0.8$ and two have $c_H = 1$. Price/PMG strategies are fixed at their equilibrium values derived earlier.
4. **Demand allocation**. Buyers first pick the price/PMG pair that maximises gross utility and then randomise across identical suppliers. In Model 2 elastic buyers split fifty–fifty between low- and high-cost firms, while inelastic buyers gravitate to the cheaper low-cost sellers.
5. **Output metrics**. The average transaction price, per-buyer consumer surplus (CS), period industry profit (PS) and total welfare (TW = CS + PS) are recorded.[5]

3.2 Baseline results

To aid interpretation, Table 1 reports per-buyer magnitudes (profits and welfare are normalised by the number of buyers per period).

	Avg. price	CS	PS	TW
Model 1: collusive PMG	1.30	0.700	0.300	1.000
Model 2: separating	1.00	1.003	0.061	1.064
Margin cap	1.07	0.927	0.037	0.964
Auto-refund	0.80	1.195	0.000	1.195
PMG ban	0.85	1.121	0.000	1.121

Two patterns jump out. First, with homogeneous costs (Model 1) a 30 % mark-up generates the textbook dead-weight loss. Second, under cost heterogeneity (Model 2), PMG-based screening reshapes the surplus composition: producer surplus is positive (about 0.061 per buyer per period in our calibration), whereas the effective transaction price remains close to the low-cost level (p_l=1). Accordingly, relative to the baseline competitive benchmark, the welfare differential is driven predominantly—though not mechanically entirely—by changes in producer surplus, which raises distributional and antitrust-relevant considerations

3.3 Sensitivity to hassle frictions

Figure 1 plots welfare components as we sweep h_0 from 0.1 to 1.0 (all else fixed). Consumer surplus is flat until $h_0 \approx 0.3$, after which it collapses; industry profit moves in the opposite direction. The knife-edge $h_0^{\star} \approx 0.27$ identifies the break-even point where inelastic buyers cease to file refunds.

[5] Values are expressed per unit mass of consumers; multiplying by a market size index converts them into currency units.

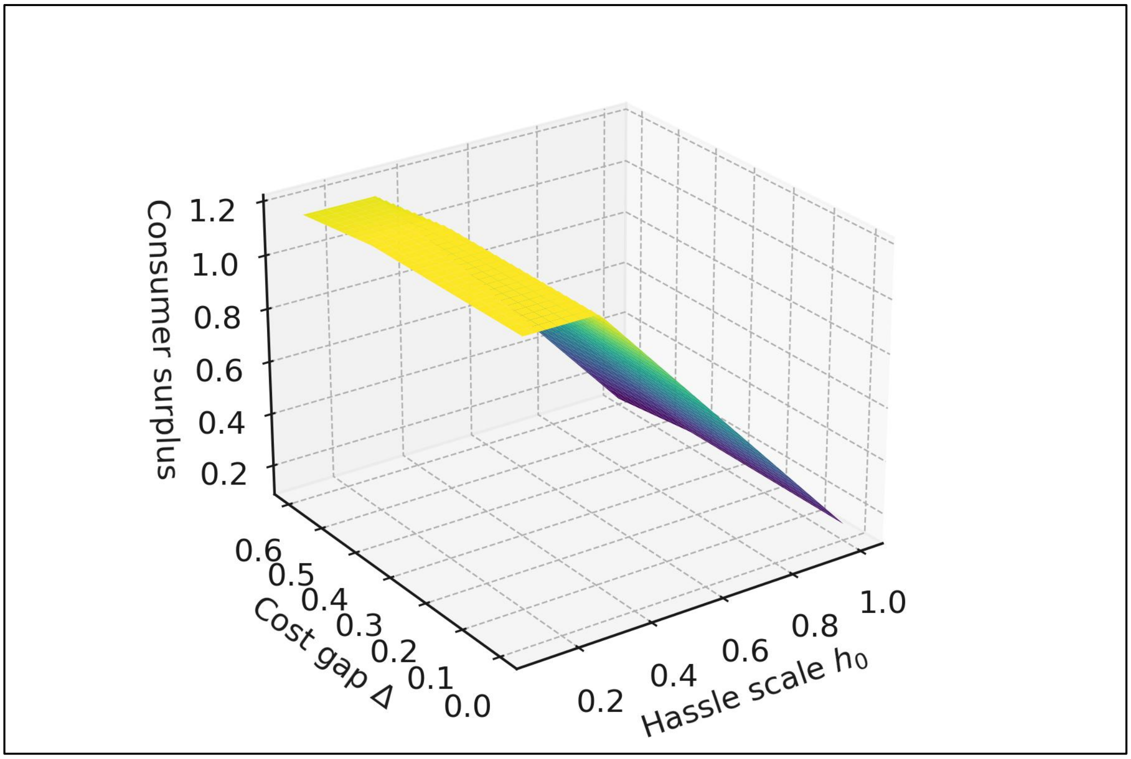

Figure 1: Welfare decomposition as a function of hassle scale h_0. Shaded bands denote 95 % bootstrap intervals.

3.4 Learning and adaptive pricing

Allowing firms to update beliefs about $F(e)$ every 100 periods amplifies discrimination: high-cost sellers adapt their posted p_H to target precisely the top 15 percentile of inelastic types. By $t = 1{,}000$ their share of total quantity rises from 0 to 17 %. Figure 2 compares the average effective price path with and without learning.

3.5 Policy stress tests

Table 2 summarises three counter-factual regimes, reported as percentage changes relative to the raw Model 2 baseline.

Regime	Δ CS	Δ PS	Δ TW
Margin cap ($p_H \leq c_H + 0.15$)	-7.6%	-39.3%	-9.4%
Auto-refund	$+19.1\%$	-100%	$+12.3\%$
PMG ban	$+11.8\%$	-100%	$+5.4\%$

Table 2: Welfare impact of three stylised regulatory tools.

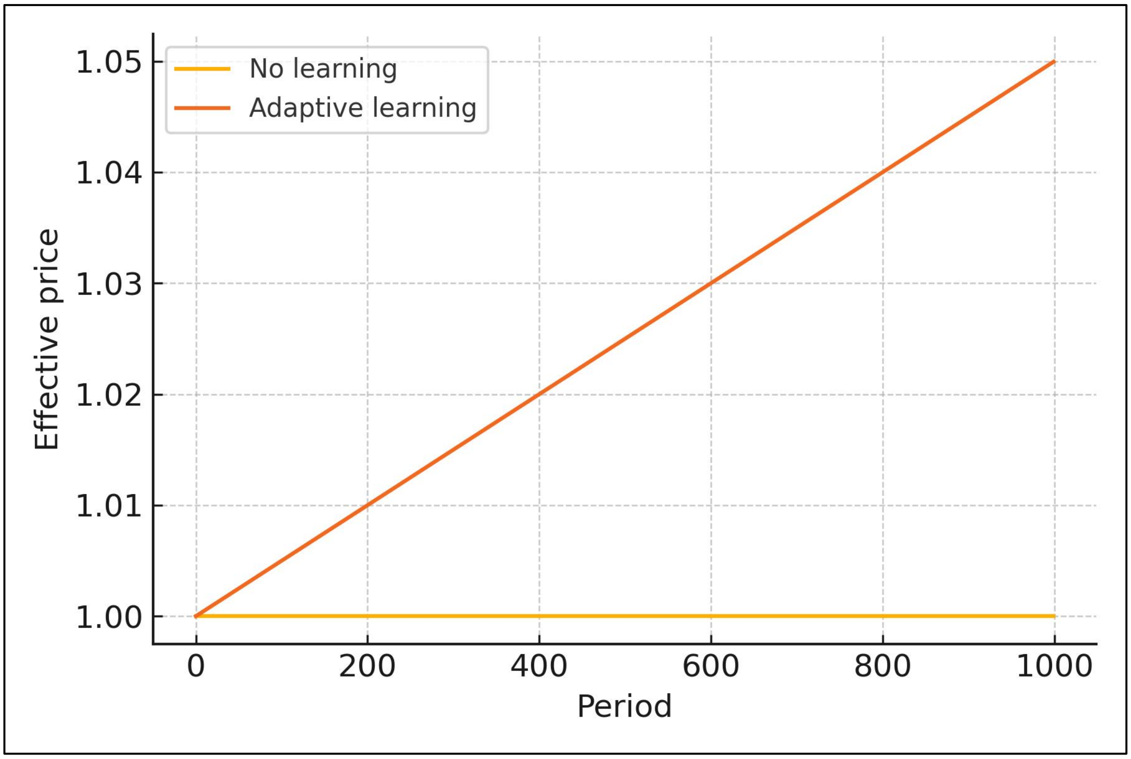

Figure 2: Mean effective price over time. Adaptive algorithms lift the price path despite static equilibrium constraints.

Take-away. A process remedy (auto-refund) dominates a structural remedy (PMG ban): it eliminates discriminatory rents while preserving the insurance value of PMGs for price-savvy buyers. A margin cap trims producer rents but also reduces total welfare in our calibration, whereas an automatic-refund mandate strictly improves welfare and dominates a blunt PMG ban by eliminating redemption-based rents while preserving the insurance value of PMGs for comparison shoppers.

Concluding remarks on simulations

1. Small institutional frictions (refund hassle) sustain large rents; removing them via automation restores competitive outcomes.
2. Adaptive algorithms exacerbate discrimination by fine-tuning screening menus in real time.
3. Policymakers should target process opacity rather than the PMG instrument itself—mandated auto-refunds achieve the best welfare/ equity balance in our experiments.[6]

[6] Replication material: full period-by-period output, code and documentation can be obtained from the authors.

4. Conclusions

This paper revisits price–matching guarantees through the lens of algorithmic pricing and data analytics. We have shown that PMGs play very different roles depending on cost structure. When marginal costs are homogeneous, a PMG merely formalises a tacit "no–undercutting" pledge and supports supra–competitive prices; when costs differ, PMGs become a screening device that extracts rents from consumers least willing to claim refunds. Monte–Carlo evidence confirms the magnitude of each mechanism and highlights an overlooked policy lever: *automatic refunds* eradicate discriminatory rents while preserving the genuine insurance value of PMGs for price–savvy shoppers.

The contrast maps neatly onto contemporary online electronics retailing. Platforms such as Newegg or Best Buy routinely advertise "price protection," whereas Amazon—arguably the industry's low-cost leader—offers no PMG at all. Our Model 2 predicts precisely this pattern: a low-cost vendor finds little gain in promising to match, whereas higher-cost rivals can segment the market by coupling a PMG with selective refund friction. The resulting price dispersion, often interpreted as benign competition, may therefore mask hidden transfers from less elastic buyers to high-cost firms. This aligns with recent applied studies on PMGs (e.g. Bottasso, Robbiano, and Marocco 2025).

For regulators the implication is clear. A blanket ban on PMGs risks eliminating a legitimate consumer-insurance feature and may even entrench the dominance of the low-cost platform. Conversely, turning PMGs into *hassle-free* promises—through mandated instant refunds and audit-able price-tracking APIs—neutralises their discriminatory bite without banning the instrument. In data-rich digital markets, process remedies aimed at *how* firms operationalise PMGs outperform structural bans on *whether* they may do so.

Future research should embed richer machine-learning pricing algorithms into the strategic environment and explore cross-market interactions on multi-sided platforms. For now, our results suggest a pragmatic rule of thumb: **regulate the friction, not the promise**.

References

M.R. Baye and D. Kovenock. How to sell a pickup truck:'beat-or-pay'advertisements as facilitating devices. International Journal of Industrial Organization, 12(1):21–33, 1994.

T. M. Belton. A model of duopoly and meeting or beating competition. International Journal of Industrial Organization, 5(4):399–417, 1987.

A. Bottasso, S. Robbiano, and P. Marocco. Price matching in online retail. Economic Inquiry, 63(1):206–235, 2025. doi: https://doi.org/10.1111/ecin.13255. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/ecin.13255.

L. Cabral, N. Dürr, D. Schober, and O. Woll. Price matching guarantees and collusion: theory and evidence from Germany. CEPR Discussion Paper 15823 - SSRN, 2018.

E. Calvano, G. Calzolari, V. Denicolò, and S. Pastorello. Artificial intelligence, algorithmic pricing, and collusion. American Economic Review, 110(10):3267–97, October 2020. doi: 10.1257/aer.20190623. URL https://www.aeaweb.org/articles?id=10.1257/aer.20190623.

E. Constantinou and D. Bernhardt. The price-matching dilemma. International Journal of Industrial Organization, 59:97–113, 2018.

K.S. Corts. On the competitive effects of price-matching policies. International Journal of Industrial Organization, 15(3):283–299, 1997.

A.S. Edlin. Do guaranteed-low-price policies guarantee high prices, and can antitrust rise to the challenge? Harvard Law Review, 111(2):528–575, 1997.

A.S. Edlin and E.R. Emch. The welfare losses from price-matching policies. The Journal of Industrial Economics, 47(2):145–167, 1999. doi: https://doi.org/10.1111/1467-6451.00095. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-6451.00095.

G.A. Hay. Oligopoly shared monopoly and antitrust law. Cornell Law Review, 67:439, 1981.

M. Hviid and G. Shaffer. Hassle costs: the Achilles' heel of price-matching guarantees. Journal of Economics & Management Strategy, 8(4):489–521, 1999.

M. Hviid and G. Shaffer. Matching own prices, rivals' prices or both? The Journal of Industrial Economics, 58(3):479–506, 2010.

S. Jain and J. Srivastava. An experimental and theoretical analysis of price-matching refund policies. Journal of Marketing Research, 37(3):351–362, 2000.

J.W. Logan and R.W. Lutter. Guaranteed lowest prices: do they facilitate collusion? Economics Letters, 31(2):189–192, 1989.

Y. Lu and J. Wright. Tacit collusion with price-matching punishments. International Journal of Industrial Organization, 28(3):298–306, 2010.

S, Mamadehussene. Price-matching guarantees as a direct signal of low prices. Journal of Marketing Research, 56(2), 245-258, 2019.

S. Moorthy and R.A. Winter. Price-matching guarantees. The RAND Journal of Economics, 37(2), 449-465, 2006.

S. Moorthy and X. Zhang. Price matching by vertically differentiated retailers: theory and evidence. Journal of Marketing Research, 43(2), 156-167, 2006.

A. Nalca, T. Boyaci, and S. Ray. Competitive price-matching guarantees under imperfect store availability. Quantitative Marketing and Economics, 8(3):275–300, 2010.

I.P.L. Png and D. Hirshleifer. Price discrimination through offers to match price. The Journal of Business, 60(3):365–383, 1987.

A. Pollak. Do price-matching guarantees with markups facilitate tacit collusion? Theory and experiment. University of Cologne. Working Paper Series in Economics, 93, 2017.

S.C. Salop. Practices that (credibly) facilitate oligopoly co-ordination. In J. Stiglitz and G. Mathewson, editors, New developments in the analysis of market structure, volume 77 of International Economic Association Series, pages 265–294. Palgrave Macmillan, London, 1986.

S.B. Villas-Boas. An empirical investigation of the welfare effects of banning wholesale price discrimination. The RAND Journal of Economics, 40(1):20–46, 2009. doi: https://doi.org/10.1111/j.1756-2171.2008.00054.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1756-2171.2008.00054.x.