Platform's Surprise Discounts with Salient Consumers

Abstract

Retail platforms increasingly deploy surprise discounts — vouchers randomly allocated to a subset of consumers upon visiting a website or opening an app. When consumers have limited attention in evaluating vertically differentiated products, surprise discounts can create heterogeneous choice contexts among consumers and further reshape market competition. In this paper, we explore whether surprise discounts can generate strategic value for the platform when facing salient-thinking consumers, and if so, through what mechanism. We develop a game theoretical model in which a commission-based platform hosts two vertically differentiated sellers. Absent platform intervention, competition for consumer attention inevitably leads to a price-salience equilibrium, resulting in lower commission for the platform. Fixed-value surprise discounts reshape this dynamic. By lowering both prices by a constant amount, discounts amplifies the price ratio, making price more salient for recipients than for non-recipients. It enables a hybrid-salience equilibrium, in which recipients focus more on price and predominantly purchase from the low-quality seller, while non-recipients are more attentive to quality and are willing to pay a premium for the high-quality product. As a result, both sellers charge higher prices, increasing platform profit despite the cost of offering discounts. We find that offering surprise discounts is profitable only when salience intensity is above certain threshold. We also characterize how the platform calibrates discount depth and reach, showing that deeper discounts offered to fewer consumers become optimal when salience is more severe, commission rates are higher, or quality differentiation is greater.

Keywords: context-dependent preferences, platform, salience, vertical differentiation, surprise discounts

1 Introduction

> *It seems as if the human brain is as much an efficient "ignoring machine" as it is a noticing one. What the brain does is less about deciding where to allocate our attention, but instead is more about working out what won't be ignored. Gigerenzer (2007)*

When consumers compare products, they do not weigh all attributes equally. A growing body of evidence in behavioral economics and marketing demonstrates that consumers allocate disproportionate attention to whichever product attribute, price or quality, stands out most in a given choice context, while under-weighting the less prominent attribute (Bordalo et al. 2012). This phenomenon, known as salient thinking, implies that consumer preferences are context-dependent rather than fixed. For vertically differentiated products, such context dependence can meaningfully shape competitive interactions. For example, when the price ratio for two products are more pronounced than their quality ratio, consumers may focus more on product price than quality, potentially intensifying price competition and putting downward pressure on seller margins.

Retail platforms are well positioned to influence these evaluation environments. A prominent and increasingly common instrument is the surprise discount: a voucher randomly allocated to a subset of consumers upon visiting a website or opening an app. For instance, users opening the Agoda app may receive a $10 voucher applicable to any hotel booking; similar practices are widespread on platforms such as Booking, Temu, and Amazon during major shopping events. These discounts are typically viewed as demand-stimulating tools. However, given that consumers have limited attention, such interventions may also influence how consumers evaluate product price and quality.

An important but underexplored feature of surprise discounts is that they create heterogeneous choice contexts across consumers. Recipients evaluate products at discounted effective prices, while non-recipients face the full posted prices. Since salient thinking depends on relative comparisons within a consumer's own choice set, even a modest discount can shift the salient attribute between price and quality. This suggests that surprise discounts may have implications beyond demand stimulation, by affecting how consumers evaluate products and how competition unfolds among sellers on the platform.

In this paper, we explore whether surprise discounts can generate strategic value for the platform facing salient-thinking consumers beyond their conventional demand-stimulation role,

and if so, through what mechanism. In particular, we examine three research questions. First, how does salient thinking distort competitive dynamics in a vertically differentiated market absent platform intervention? Second, can a platform strategically deploy surprise discounts to reshape consumer attention and alter the competitive regime? Third, under what conditions might such attention shifts translate into improved platform performance?

To address these questions, we develop a game-theoretic model of a commission-based platform hosting two vertically differentiated sellers. The platform collects a fraction of each seller's transaction revenue as commission and can issue fixed-value discount to a randomly selected subset of consumers. After observing the platform's discount policy, the two sellers simultaneously choose their posted prices. Consumers are heterogeneous in their willingness to pay for quality and exhibit salient thinking. In evaluating products, they focus on one attribute between price and quality that displays greater proportional variation within the choice set, while discounting the less salient attribute. This context-dependent evaluation implies that perceived utilities depend not only on product characteristics but also on how those characteristics compare across available options.

A key feature of surprise discount is that randomly allocating fixed-value discount can create heterogeneous evaluation contexts. Discount recipients face lower effective prices, whereas non-recipients face the original posted prices. As a result, the two groups may attend to different attributes even when facing the same posted prices, which might impact sellers' pricing decisions.

Our analysis proceeds in three steps. We begin by characterizing market outcomes in the absence of platform intervention. We show that competition for salient consumers leads to a unique price-salience equilibrium. The intuition is driven by an asymmetry between the two sellers: under quality salience, the low-quality product is strictly dominated and thus attracts no demand. To stay alive, the low-quality seller has a strong incentive to cut prices and shift consumers' attention toward price. This survival instinct always drives the market toward *the price-salience equilibrium*. In this case, consumers are more sensitive to price so that both sellers engage in intense price competition, resulting in lower prices, compressed margins, and reduced platform commission revenue relative to a benchmark with rational consumers.

We then show that surprise discounts can break this competition trap. A key property of fixed-value discounts is that they reduce both products' prices by the same absolute amount, which compresses the effective price ratio for recipients more than for non-recipients. As a result,

the platform can induce different salience perceptions across consumers from a single set of posted prices. Under a suitably chosen discount scheme, the market transitions to *the hybrid-salience equilibrium*. Consumers who receive the discount focus more on price and predominantly purchase from the low-quality seller, while those who do not receive the discount remain more attentive to quality and are willing to pay a premium for the high-quality product. This heterogeneous attention pattern reduces direct price competition between the two sellers. The low-quality seller relies on discount recipients and therefore has less incentive to aggressively undercut its price, which in turn allows the high-quality seller to maintain a higher price. As a result, both sellers sustain higher margins, and the platform benefits from higher commission revenue that can outweigh the cost of offering discounts.

Finally, we characterize the platform's optimal discount policy and show that it follows a threshold structure. When salience distortions are sufficiently intense, the platform finds it profitable to offer discounts that induce the hybrid-salience outcome. When distortions are mild, the platform optimally refrains from intervention and the market remains in the price-salience equilibrium. This generates a non-monotonic (U-shaped) relationship between the degree of salience distortion and platform profit. Within the hybrid-salience region, the platform balances the size and rate of discounts: it prefers to preserve a large pool of non-recipients who focus on quality and are more likely to pay higher prices, while offering enough discounts to keep the low-quality seller from deviating. Numerical analysis reveals that deeper discounts allocated to fewer consumers are optimal under stronger salience distortion, higher commission rates, or greater quality differentiation.

To the best of our knowledge, this paper is the first to explore how a platform can strategically leverage context-dependent consumer preferences to manage seller competition in a vertically differentiated market. Our analysis makes two contributions. First, we uncover an attention-shaping mechanism: fixed-value surprise discounts alter effective price comparisons, inducing heterogeneous salience perceptions across consumers. Second, we show that by shaping consumer attention, the platform better disciplines seller competition: surprise discount relaxes the competitive pressure that drives markets into a price-salience equilibrium, allowing both sellers to sustain higher prices and profits. Importantly, the platform achieves this without directly controlling prices or quantities, but by shaping the contexts in which consumers evaluate product. The outcome is mutually beneficial: the platform earns higher commissions, the high-quality seller

restores its quality premium, and the low-quality seller survives without engaging in destructive price competition.

The remainder of the paper is organized as follows. Section 2 review the relevant literature. Section 3 presents the model. Section 4 develops the analysis, beginning with the no-discount benchmark and proceeding to the surprise discount equilibrium and optimal policy. Section 5 presents extensions covering asymmetric costs, heterogeneous salience susceptibility, and targeted discounts. Section 6 discusses managerial implications and concludes.

2 Literature Review

Our paper contributes to three related streams of research: promotional discounts and coupon strategy, context-dependent consumer preferences in retailing, and platform strategies for managing seller competition. We review each in turn and highlight our contributions.

2.1 Promotional Discounts and Coupon Strategy

A large literature studies how firms use discounts as strategic instruments to stimulate demand and segment consumers. The foundational framework of Narasimhan (1984) models coupons as a price discrimination device that separates price-sensitive consumers willing to incur redemption costs from less elastic consumers who pay posted prices. Extending this logic, Shaffer and Zhang (1995) show that in competitive settings, coupon can generate prisoner's dilemma outcomes in which firms are worse off relative to no-discount benchmarks. Empirical work further demonstrates that the effectiveness of promotions depends on their design: Bawa and Shoemaker (1987, 1989) show that coupon face value has nonlinear effects on redemption and profitability, while DelVecchio et al. (2007) document that the format of discounts affects consumer perceptions and purchase behavior even when the economic magnitude is identical.

A common theme in this literature is that discounts operate as demand-stimulating tools to encourage purchase and discriminate consumers. Our paper departs from this perspective by identifying a distinct channel through which discounts can create value. In our model, randomly allocated fixed-value discounts reshape the evaluation context in which consumers compare products. This attention-based mechanism is absent from classical coupon theory, where

preferences are assumed to be stable and context-independent. More broadly, we show that the value of discounts matters not only through changing consumer purchases, but also through their ability to influence seller competition.

2.2 Context-Dependent Preferences in Retailing

Our behavioral foundation builds on the salience theory of choice developed by Bordalo et al. (2013), which formalizes how consumers overweight attributes that stand out relative to a reference point while underweight others. This framework generates context-dependent preferences and provides a unifying explanation for a wide range of behavioral phenomena, including decoy effects and framing effects. Related models of attention and context dependence include Kőszegi and Szeidl (2013), who emphasize variation-based attention weights, and Bushong et al. (2021), who study relative thinking in which evaluations depend on comparison ranges.

Recent research discuss business strategies when consumers exhibit context-dependent preferences. Bordalo et al. (2016) show that salience can fundamentally alter market outcomes, with competition endogenously leading to either price-salience or quality-salience equilibria. Subsequent work examines how firms strategically respond to or exploit such preferences. For example, Zheng et al. (2019) show that probabilistic selling can reshape attention toward quality, Inderst and Obradovits (2020) demonstrate how price promotions can intensify price competition under salience, and Zhu and Dukes (2017) analyze how firms choose which attributes to make prominent when consumers have limited attention. Empirical evidence further confirms that salience manipulation can have first-order effects on consumer choice: Blake et al. (2021) show that making fees salient shifts consumers toward lower-quality products, affecting both demand and product mix.

A consistent finding in this literature is that attention-driven competition tends to intensify price rivalry and erode margins in vertically differentiated markets. Our paper contributes by identifying a mechanism through which such competitive pressure can be mitigated. In particular, surprise discounts can generate heterogeneous attention patterns across consumers, allowing price and quality salience to coexist in equilibrium. This endogenous segmentation relaxes the forces that typically drive markets toward price salience, offering a new perspective on how context-dependent preferences interact with competition.

2.3 Platform Strategy and Seller Competition

A growing literature examines how retail platforms design mechanisms to influence competition among sellers. Beyond pricing and fees, platforms use a variety of non-price instruments, such as recommendation systems (Shi and Raghu 2020), product certification (Guan et al. 2026), shipping service (Sun et al. 2025), financial support (Huang et al. 2025), privacy control (Zhu et al. 2025), and infrastructure investment (Puyang et al. 2026), to shape seller competition on the market.

Our paper offers a complementary perspective by focusing on the platform's role in shaping consumer attention and, through it, competitive interactions among sellers. We show that surprise discounts, commonly viewed as short-term promotional tools, can serve as instruments of attention management. This perspective highlights a new channel through which platforms can influence market outcomes.

3 Model Setup

We study a commission-based e-commerce platform hosting two sellers who offer vertically differentiated products to a unit mass of consumers with heterogeneous quality valuations. This section introduces the market primitives, formalizes salient thinking, and specifies the surprise discount mechanism and timing.

3.1 Market Primitives

The platform charges a fixed commission rate $\gamma \in (0, 1)$ on each seller's sales revenue. We treat γ as exogenous, reflecting industry practice: commission rates on major e-commerce platforms are typically set by long-run policy and remain stable relative to the higher-frequency pricing decisions of individual sellers.

Two sellers H and L offer products of distinct qualities q_H and q_L, respectively. For expositional convenience, we normalize qualities to $(q_H, q_L) = (1, k)$, where $k \in (0, 1)$ captures the degree of vertical differentiation. We refer to the two sellers as the high-quality (H) and low-quality (L) vendors. Each seller simultaneously chooses a posted price, p_H or p_L, to maximize profit net of the platform's commission.

Both sellers face the same constant marginal cost c. Equal marginal costs allow us to isolate the demand-side consequences of salience and platform intervention, rather than conflating them with cost asymmetries. This assumption is also consistent with many service and experience-good markets such as hospitality, car rentals, and air travel, where marginal operating costs are broadly similar across quality tiers, with differentiation primarily reflecting fixed investments in amenities, branding, or design. As we discuss later, the main results extend to settings with asymmetric costs, so long as providing high quality is more efficient (i.e., $q_H/c_H > q_L/c_L$).

A unit mass of consumers participates in the market. Each consumer has unit demand and is characterized by a type $\theta \sim U[0,1]$, representing marginal willingness to pay for quality. Under fully rational evaluation, the utility from purchasing product $j \in \{H, L\}$ at price p_j is $u^R(q_j, p_j) = \theta q_j - p_j$, with the outside option normalized to zero. Consumers observe their own type θ, while sellers observe only the distribution and therefore set prices to segment demand in equilibrium. This structure follows the standard vertical differentiation framework and serves as a benchmark against which we introduce salience-driven distortions and the platform's surprise discount mechanism.

3.2 Salient Thinking

With context-dependent preferences, consumers depart from the rational benchmark u^R by allocating disproportionate attention to whichever attribute, price or quality, appears most prominent. An attribute is salient when it exhibits greater relative variation across the two products $\{(q_H, p_H), (q_L, p_L)\}$ than the other attribute.

Following Bordalo et al. (2013, 2016), we assume salience is determined by the relative variation of each attribute across products. Specifically, quality is salient when the quality ratio exceeds the price ratio, $q_H/q_L > p_H/p_L$, meaning that quality varies proportionally more across products than price does. Conversely, price is salient when $p_H/p_L > q_H/q_L$, meaning that price differences are proportionally larger. In the knife-edge case $p_H/p_L = q_H/q_L$, both attributes are equally salient.

The salience condition admits an equivalent product-level interpretation: Consumers are drawn to the option that appears to offer a better deal - with the higher quality-to-price ratio q_j/p_j. Then, the attribute in which that product is relatively strong becomes salient: quality is salient when the high-quality product offers a higher quality-to-price ratio, $q_H/p_H > q_L/p_L$ and

the price is salient when the opposite holds.

Once a salient attribute is identified, consumers overweight it while discounting the non-salient attribute. We capture this distortion with a parameter $\delta \in (0,1)$ applied to the non-salient attribute. When quality is salient, consumers discount price; when price is salient, they discount quality. The resulting perceived utility from product j at price p_j is

$$u^S(q_j, p_j) = \begin{cases} \theta q_j - \delta p_j, & \text{if } q_H/q_L > p_H/p_L, \text{ where quality is salient,} \\ \theta\delta q_j - p_j, & \text{if } q_H/q_L < p_H/p_L, \text{ where price is salient,} \\ \theta q_j - p_j, & \text{if } q_H/q_L = p_H/p_L, \text{ where both attribute are equally salient.} \end{cases}$$

A smaller δ corresponds to a stronger salience distortion: consumers devote more attention to the salient attribute and increasingly neglect the other. The limit $\delta \to 1$ recovers rational evaluation. Consistent with Bordalo et al. (2016), Zheng et al. (2019), we assume a common salience parameter across consumers. We discuss in the extensions that allowing for heterogeneity in δ does not qualitatively affect the main results but significantly complicates the analysis.

3.3 Discount Mechanism and Timing

E-commerce platforms frequently employ surprise discounts to stimulate demand and capture consumer attention. We model this practice as follows. Upon visiting the platform, each consumer independently receives a fixed-value voucher worth $w > 0$ with probability $\beta \in (0,1)$, redeemable toward either seller's product. The mechanism captures key features of platform-issued promotions: discounts are not product-specific and are randomly allocated across consumers realized upon arrival. These characteristics reflect common practices such as platform-wide shopping festival vouchers on major e-commerce platforms. The notations are summarized at Table 1.

In practice, platforms deploy fixed-value discounts (e.g., "$10 off") or percentage-based discounts (e.g., "20% off"). In our main analysis, we demonstrate the strategic value of fixed-value discount by reshaping effective price ratios of two products. The mechanism does not apply to the percentage-based discounts as it maintains the price ratio after discount. We formally verify this distinction and the absence of strategic value for percentage-based discounts in the appendix.

Platforms also employ targeted discounts, tailoring offers to consumers' profiles or purchase histories. We abstract from targeting and instead adopt a randomized discount scheme in order

Table 1: Summary of Notations

k	The degree of vertical differentiation
c	Marginal cost of the two products
w	The value/level of surprise discount
β	The probability/rate of surprise discount
δ	Salient factor of customers
γ	Platform's commission rate
D_i	Market demand for seller $i \in \{H, L\}$

to isolate the strategic role of platform discounts while preserving analytical tractability. As we discuss later, incorporating targeting would grant the platform additional flexibility to optimize discount allocation based on consumer characteristics, likely amplifying rather than weakening the benefits we identify. Our results therefore provide a conservative estimate of the strategic value of surprise discounts.

We model the interaction as a three-stage game. In Stage 1, the platform commits to a discount policy (w, β), publicly announcing that a fraction β of consumers will receive a discount of value w applicable to either product. This commitment is observed by both sellers and consumers prior to any pricing or purchase decisions, reflecting common practices such as pre-announced promotions during major shopping events. In Stage 2, taking the discount policy (w, β) and the commission rate γ as given, the two sellers simultaneously choose prices p_H and p_L to maximize their respective profits. In Stage 3, consumers arrive and realize whether they have received a discount voucher. Recipients face effective prices $p_j - w$ whereas non-recipients face the posted prices p_j, $j \in \{H, L\}$. Each consumer then chooses to buy from seller H, seller L, or to opt out, based on perceived utility given her realized choice context.

This timing highlights the dual channel through which surprise discounts operate. On the demand side, discounts modify consumers' effective choice sets and, through the salience mechanism, their relative attention to price versus quality. On the supply side, because sellers observe the discount policy before setting prices, equilibrium pricing already internalizes the anticipated effects of discounts on consumer segmentation. The platform's instrument thus works simultaneously through direct context manipulation and indirect reshaping of sellers' pricing

incentives. For expositional simplicity, we focus on the parameter region $c \leq \bar{c} \equiv \delta k(1-\gamma)/2$, which guarantees that both sellers attract positive demand and remain active in equilibrium.

4 Analysis

We begin by characterizing market equilibria in the absence of the platform's surprise discount policy. To isolate the role of consumer limited attention, we distinguish between a benchmark with rational consumers and an setting in which consumers exhibit salient thinking. This comparison reveals how competition for consumer attention intensifies price rivalry and erodes seller margins, motivating the platform's intervention analyzed in Section 4.2.

We solve the model by backward induction. We first derive consumers' purchase decisions given prices. We then characterize the sellers' pricing equilibrium, where each seller $i \in H, L$ chooses its price to maximize profit:

$$\pi_i = \big((1-\gamma)p_i - c\big)D_i.$$

Finally, we analyze the platform's optimal discount policy, if any.

4.1 Benchmark without Surprise Discount

Rational consumers ($\delta = 1$). When consumers are fully rational, utility depends only on a product's own quality and price and is unaffected by the composition of the choice set. Given the choice set $\{(q_H, p_H), (q_L, p_L)\}$, the market is segmented into three groups based on quality valuations. High-valuation consumers with $\theta \in [\frac{p_H - p_L}{q_H - q_L}, 1]$ purchase product H, intermediate-valuation consumers with $\theta \in [\frac{p_L}{q_L}, \frac{p_H - p_L}{q_H - q_L})$ choose product L, and low-valuation consumers opt out. Hence, the two sellers' demands are given by $D_H = 1 - \frac{p_H - p_L}{q_H - q_L}$ and $D_L = \frac{p_H - p_L}{q_H - q_L} - \frac{p_L}{q_L}$, respectively. Solving the sellers' pricing problems simultaneously yields the unique Nash equilibrium:

$$p_H^R = \frac{2(1-k)(1-\gamma) + 3c}{(4-k)(1-\gamma)}, \tag{1}$$

$$p_L^R = \frac{k(1-k)(1-\gamma) + (2+k)c}{(4-k)(1-\gamma)}. \tag{2}$$

Salient consumers ($\delta < 1$). When consumers exhibit salient thinking, product valuations depend not only on absolute levels of quality and price but also on their relative standing within the

choice set. Sellers can therefore strategically manipulate the salient attribute through their pricing decisions, steering consumer attention toward their respective competitive advantages. However, the two sellers take fundamentally asymmetric positions in this competition for attention.

Consider first what happens under quality salience, i.e., $q_H/q_L > p_H/p_L$. Consumers focus on quality differences and discount price, becoming less price-sensitive. Although this would seem to benefit sellers by softening price competition, the low-quality product becomes strictly dominated: every consumer who derives positive utility from the low-quality product finds the high-quality product strictly more attractive. Formally, the demand interval for the low-quality seller collapses:

$$\frac{p_H - p_L}{q_H - q_L} - \frac{p_L}{q_L} = \frac{q_L/p_L - q_H/p_H}{(q_H - q_L)q_H} < 0,$$

where the inequality follows directly from the quality-salience condition $q_H/p_H > q_L/p_L$. The low-quality seller is therefore excluded from the market entirely under quality salience.

Now consider the converse case with price salience, i.e., $p_H/p_L > q_H/q_L$. Here consumers focus on price and discount quality, becoming more price-elastic. Both sellers retain positive demand but face intensified price competition due to heightened price sensitivity. The high-quality seller is disproportionately harmed as its differentiation advantage is diminished because consumers discount the quality gap between products.

The high-quality seller faces a genuine choice: cut prices aggressively to induce quality salience, or maintain higher prices and accept reduced market share under price salience. The low-quality seller has no such flexibility. Under quality salience its product attracts zero demand, so to avoid market exclusion it must cut prices to restore price salience whenever possible. As a result, any attempt by the high-quality seller to induce quality salience triggers aggressive price responses from the low-quality rival, pulling attention back to price. The competitive equilibrium therefore invariably produces price salience, and the two sellers' demands are given by $D_H = 1 - \frac{p_H - p_L}{\delta(q_H - q_L)}$ and $D_L = \frac{p_H - p_L}{\delta(q_H - q_L)} - \frac{p_L}{\delta q_L}$, respectively..

Proposition 1 *For salient consumers, the sellers' optimal pricing decisions always induce a price-salience equilibrium characterized by*

$$p_H^N = \frac{2\delta(1-k)(1-\gamma) + 3c}{(4-k)(1-\gamma)}, \tag{3}$$

$$p_L^N = \frac{k\delta(1-k)(1-\gamma) + (2+k)c}{(4-k)(1-\gamma)}, \tag{4}$$

where both equilibrium prices decrease as salience distortion intensifies, with the high-quality seller disproportionately affected, i.e., $\partial p_H^N/\partial\delta > \partial p_L^N/\partial\delta > 0$.

Eqs. 3 and 4 mirror their rational-consumer counterparts in Eqs. 1 and 2, with δ discounting the perceived quality differential. The comparative statics confirm the mechanism outlined above: as salience distortion intensifies (lower δ), consumers attend more to price. In this case, both sellers respond by cut prices and the high-quality seller's price falls more steeply ($\partial p_H^N/\partial\delta > \partial p_L^N/\partial\delta > 0$).

Corollary 1 (Lower Price, Sales and Profit under Salient Thinking) *When consumers exhibit salient thinking, both sellers charge lower prices while simultaneously serving fewer consumers relative to the rational benchmark, yielding lower profits for both sellers and reduced commission revenue for the platform.*

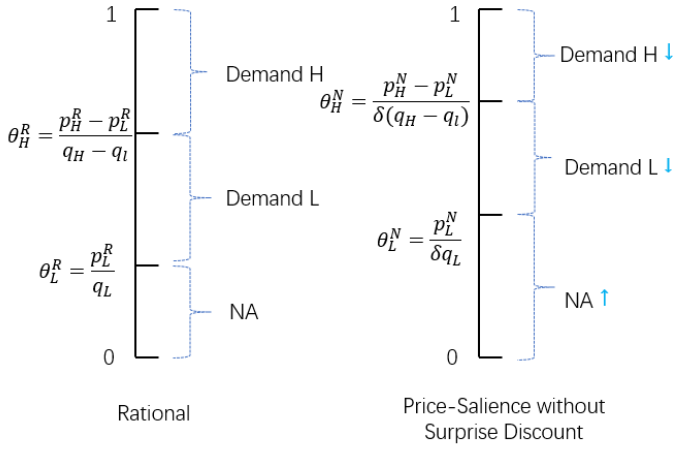

Figure 1: Market Segmentation without Platform Surprise Discounts

Comparing the equilibrium under salient thinking to the rational benchmark reveals that salient thinking leaves sellers and the platform uniformly worse off. Both sellers charge lower prices yet, contrary to the standard price-demand trade-off, also serve fewer consumers. The mechanism differs from ordinary price competition: the low-quality seller's survival incentive locks the market into price salience, where consumers discount quality and become more price-elastic. Lower prices therefore fail to expand demand; instead, the salience-induced quality discounting dominates, compressing margins and shrinking the customer base simultaneously. The platform, earning commissions on seller revenue, suffers from both lower prices and reduced transaction volume.

This outcome reflects a central paradox of competition for consumer attention. Both sellers have incentives to improve their quality-to-price ratios: the high-quality seller seeks to highlight quality differences, while the low-quality seller aims to emphasize price savings. Yet this mutual effort to capture attention intensifies price competition, eroding margins market-wide. Given that consumers' limited attention forces sellers into intensified price rivalry, a natural question arises: can the platform intervene to reshape these competitive dynamics and restore profitability? The following section demonstrates that surprise discounts can serve precisely this function.

4.2 Platform's Surprise Discounts

Having established that salient thinking traps sellers in a price-salience equilibrium with depressed margins, we now examine whether the platform can deploy its surprise discount instrument to reshape competitive dynamics.

We begin by establishing that surprise discounts offer no benefit to the platform when consumers are rational. This result serves as a baseline: any value the platform derives from discounts under salient thinking must stem from their ability to manage consumer attention, rather than from their direct effect on demand.

Lemma 1 *Facing rational consumers, the platform does not profit from offering surprise discounts.*

The intuition is straightforward. Rational consumers evaluate products independently of the choice context, so discounts create no segmentation in how consumers perceive price relative to quality. The discount simply shifts all consumers' effective prices downward by a uniform amount. Sellers anticipate this and raise posted prices to partially absorb the subsidy. Because no change in attention patterns occurs, the platform bears the cost of the discount without generating any strategic benefit—the intervention is equivalent to a lump-sum transfer that sellers partially extract through standard pricing adjustments.

The role of surprise discounts changes fundamentally when consumers exhibit salient thinking. We first establish a key property of fixed-value discounts: by lowering both prices by the same absolute amount, a discount increases the relative price dispersion across products. Formally, $(p_H - w)/(p_L - w) > p_H/p_L$ for any $w \in (0, p_L)$ and $p_H > p_L$. This property reveals how discounts influence salient consumers' perception of the salient attribute between price and

quality, established in the following lemma.

Lemma 2 *For given posted prices, salient consumers who receive a discount are more inclined to perceive price (rather than quality) as the salient attribute. As a result, competing sellers may reach one of two possible equilibrium configurations:*

(i) the price-salience equilibrium with $p_H/p_L > q_H/q_L$, where both discount recipients and non-recipients perceive price as the salient attribute

(ii) the hybrid-salience equilibrium with $p_H/p_L < q_H/q_L < (p_H - w)/(p_L - w)$, where non-recipients perceive quality as salient while recipients perceive price as salient.

To see the intuition, recall that price becomes salient when the price ratio exceeds the quality ratio, $p_H/p_L > q_H/q_L$. When a consumer receives a discount w, products are evaluated at effective prices $(p_H - w, p_L - w)$, and price salience instead requires $(p_H - w)/(p_L - w) > q_H/q_L$. Because the effective price ratio increases with w, discounts magnify the perceived disparity in prices relative to quality differences. Surprise discounts therefore do more than lower consumers' effective costs: by reshaping relative price comparisons, they actively redirect consumer attention toward price. This attention-shifting property is central to understanding how platform-sponsored surprise discounts influence sellers' pricing incentives and competitive dynamics.

A natural concern arises: if price salience intensifies competition and erodes profits, why would the platform introduce a mechanism that makes price even more salient for recipients? The answer lies in recognizing that surprise discounts do not uniformly push the entire market toward price salience. Rather, they create heterogeneous evaluation contexts from a single set of posted prices. A fraction β of consumers (discount recipients) evaluate products at effective prices $(p_H - w, p_L - w)$, while the remaining $1 - \beta$ (non-recipients) face posted prices (p_H, p_L). Critically, salience is determined within each group by comparing the price ratio to the quality ratio in its own choice context. Surprise discounts therefore generate heterogeneity in how consumers allocate attention, rather than simply intensifying price focus across the board.

Price-salience regime. When the discount policy induces price salience for both segments, the equilibrium structure resembles the no-discount benchmark, with discounts playing a limited role. In this regime, the two sellers' demands are given by $D_H = 1 - \frac{p_H - p_L}{\delta(q_H - q_L)}$ and $D_L = \beta\left(\frac{p_H - p_L}{\delta(q_H - q_L)} - \frac{p_L - w}{\delta q_L}\right) + (1 - \beta)\left(\frac{p_H - p_L}{\delta(q_H - q_L)} - \frac{p_L}{\delta q_L}\right)$, respectively.

Lemma 3 (Price-Salience Equilibrium) *If the discount policy (w, β) induces a price-salience equilibrium, then the equilibrium prices are*

$$p_H^* = \frac{(2\delta + \beta w)(1-k)(1-\gamma) + 3c}{(4-k)(1-\gamma)},$$
$$p_L^* = \frac{(k\delta + 2\beta w)(1-k)(1-\gamma) + (2+k)c}{(4-k)(1-\gamma)},$$

under which the platform finds it optimal to set $w = 0$.

The intuition parallels Lemma 1. When both recipients and non-recipients perceive price as salient, the discount does not alter how consumers weigh price relative to quality. It simply reduces effective prices at the platform's expense. Sellers raise posted prices in anticipation of the discount program, limiting the increase in transaction volume. Absent a change in consumers' attention patterns, the platform lacks leverage to reshape market outcomes, and the discount strictly reduces platform profit.

Hybrid-salience regime. We now turn to the hybrid-salience equilibrium, where discount recipients and non-recipients attend to different product attributes: recipients focus on price while non-recipients focus on quality. In this regime, the two sellers' demands are given by $D_H = \beta\left(1 - \frac{p_H - p_L}{\delta(q_H - q_L)}\right) + (1-\beta)\left(1 - \frac{\delta p_H}{q_H}\right)$ and $D_L = \beta\left(\frac{p_H - p_L}{\delta(q_H - q_L)} - \frac{p_L - w}{\delta q_L}\right) + (1-\beta)\left(\frac{p_H - p_L}{\delta(q_H - q_L)} - \frac{p_L}{\delta q_L}\right)$, respectively.

This coexistence of heterogeneous salience perceptions within a single market is the economically novel configuration enabled by surprise discounts, and it is where the platform finds scope for profitable intervention.

Lemma 4 (Hybrid-Salience Equilibrium) *If the discount policy (w, β) induces a hybrid-salience equilibrium, then the equilibrium prices are*

$$p_H^* = \frac{2\delta(1-k)(1+c\delta-\gamma) + \beta(c(3-2\delta^2(1-k)) + (1-k)(1-\gamma)w)}{(\beta(4-4\delta^2(1-k)-k) + 4\delta^2(1-k))(1-\gamma)},$$
$$p_L^* = \frac{\beta(1-\delta^2(1-k))(c(2+k) + 2(1-k)(1-\gamma)w) + \delta(1+k)(c\delta(2+k) + (1-\gamma)(2\delta w + k(1-2\delta w)))}{(\beta(4-4\delta^2(1-k)-k) + 4\delta^2(1-k))(1-\gamma))},$$

under which the platform finds it optimal to set $w(\beta) = \frac{2\beta c(2-\delta^2(2-k)) + 2\delta(c\delta(2-k) - k + k\gamma)}{\sqrt{\beta}(\beta(4-4\delta^2(1-k)-k) + 4\delta^2(1-k) - \sqrt{\beta}k)(1-\gamma))}$, a decreasing function in β.

The hybrid-salience equilibrium transforms the competitive landscape by splitting the market into two groups with distinct attention patterns. Discount recipients focus on price and predominantly purchase from the low-quality seller; non-recipients focus on quality and are willing to

pay a premium for the high-quality product. This coexistence of attention regimes reshapes both sellers' pricing incentives.

The key mechanism is the relaxation of competitive pressure. Under price salience (Proposition 1), the low-quality seller must undercut aggressively to prevent the high-quality seller from inducing quality salience, dragging both sellers' prices down through strategic complementarity. Under hybrid salience, the low-quality seller acquires a captive base of price-focused recipients who are more inclined to favor its product. This captive base raises the low-quality seller's opportunity cost of cutting prices: an aggressive price reduction to attract non-recipients would sacrifice margins on the recipient group where it already enjoys a strong position. With low-quality seller's pricing pressure relaxed, the high-quality seller also faces softer competition and can sustain higher margins to monetize from the quality-focused non-recipients. The result is a stable equilibrium configuration in which both sellers can coexist profitably.

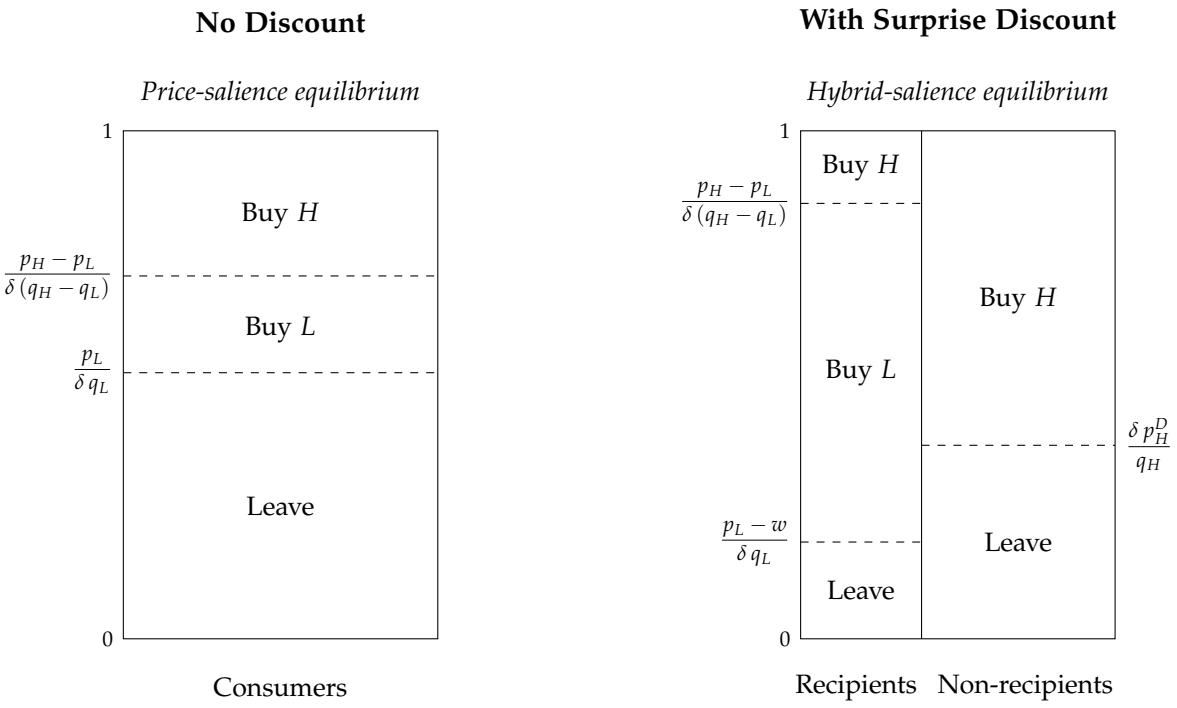

Figure 2: Market Segmentation Before and After Surprise Discount

The comparative static $\partial w(\beta)/\partial\beta < 0$ reveals a fundamental trade-off between discount reach and depth, rooted in the incentive compatibility of the low-quality seller. The hybrid-salience equilibrium is sustained only when the low-quality seller finds it more profitable to serve price-focused recipients than to undercut prices and compete for non-recipients, which would collapse

the market back to the price-salience equilibrium. When the platform exposes a larger fraction β of consumers to discounts, the recipient group that constitutes the low-quality seller's captive base grows, strengthening its incentive to remain in the hybrid regime. A smaller per-consumer subsidy then suffices to deter deviation.

This trade-off grants the platform strategic flexibility across two dimensions. The discount reach β controls the size of the price-focused group, while the discount depth w controls the intensity of attention redirection within that group. The platform exploits this flexibility to reconcile competing objectives: preserving high margins on quality-focused transactions while maintaining a viable low-quality seller. Recipients gravitate toward the low-quality product without requiring aggressive market-wide price cuts; non-recipients remain quality-focused, allowing the high-quality seller to sustain differentiation rents. The platform tunes (w, β) to balance the composition of the two groups and the competitive incentives of both sellers.

Having established that the price-salience regime offers no scope for profitable intervention (Lemma 3) while the hybrid-salience regime can relax competitive pressure and create value through attention management (Lemma 4), we now characterize the platform's optimal policy.

Proposition 2 (Optimal Discount Intervention) *There exists a threshold δ_0 such that*

(i) for $\delta < \delta_0$, the platform's optimal discount policy is characterized by unique β^* and corresponding $w^* = w(\beta^*)$, inducing a hybrid-salience equilibrium, where the platform's profit decreases with δ;

(ii) for $\delta \geq \delta_0$, the platform's optimal policy is to offer no discount, inducing a price-salience equilibrium, where the platform's profit increases with δ.

This result reveals a non-monotonic relationship between consumer salience and platform profit. When salience distortion is severe ($\delta < \delta_0$), the platform strictly benefits from offering surprise discounts. Without intervention, intense salience-driven competition compresses margins for all parties. Properly calibrated discounts prevent this outcome by inducing hybrid salience: the low-quality seller secures a viable base among price-focused recipients without engaging in ruinous undercutting, while the high-quality seller extracts rents from quality-focused non-recipients. Both sellers sustain higher posted prices, and the platform captures higher commission revenue despite subsidizing discounted transactions.

The platform's calibration problem reflects two competing objectives. On the one hand, the

platform seeks to maximize the size of the quality-focused non-recipient group, from which it earns higher commission revenue on premium-priced transactions. On the other hand, it must offer sufficient discounts to keep the low-quality seller incentive-compatible. That is, to ensure the low-quality seller finds it more profitable to serve price-focused recipients only than to undercut prices and compete for non-recipients, which would collapse the market back to the price-salience equilibrium. The optimal policy therefore minimizes discount expenditure subject to maintaining the hybrid-salience equilibrium, preserving as large a non-recipient base as possible.

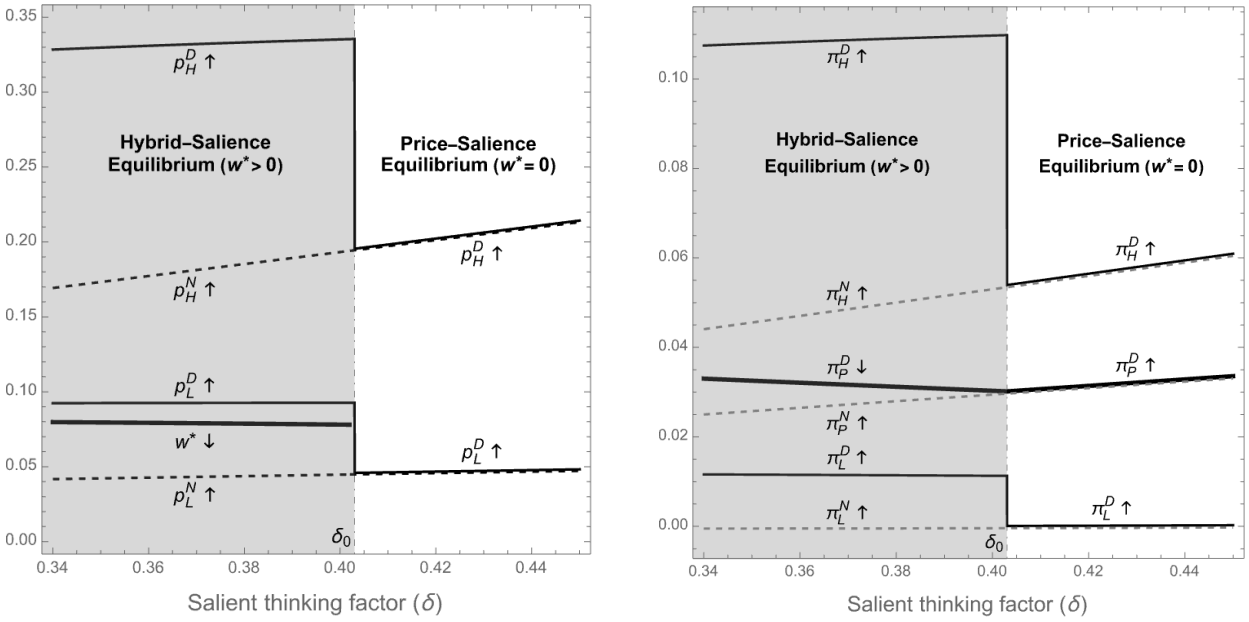

(a) Equilibrium prices (w, p_H, p_L)

(b) Equilibrium profit (π_P, π_H, π_L)

Figure 3: Equilibrium decisions (w, p_H, p_L) and profits (π_P, π_H, π_L). Solid lines correspond to the optimal policy with discounts. Dashed lines correspond to the no-discount benchmark. The platform's profit is U-shaped. When $\delta < \delta_0$ (the shaded region), offering discounts increases platform profit and benefits both sellers, a "win–win–win" outcome. When $\delta > \delta_0$, the platform does not profit from surprise discounts.

Within the hybrid-salience regime, stronger salience distortion (lower δ) further benefits the platform. Non-recipients, who perceive quality as salient, increasingly underweight price in their evaluations, allowing the high-quality seller to sustain higher margins. Simultaneously, stronger salience makes deviating toward price salience less attractive for the low-quality seller, enabling the platform to sustain segmentation with smaller discounts. Together, these forces imply that platform profit decreases as δ increases toward δ_0.

Conversely, when salience distortion is mild ($\delta \geq \delta_0$), intervention becomes unprofitable. Consumers' evaluations are sufficiently close to rational that sellers sustain reasonable margins even under price salience. Inducing hybrid salience would require discounts deep enough to compensate the low-quality seller for relinquishing the non-recipient segment, rendering the intervention unprofitable. The optimal policy is therefore no discount, and platform profit increases as salience distortions weaken further.

Taken together, these results imply a U-shaped relationship between platform profit and consumer salience severity, as illustrated in Figure 1. For low δ, stronger distortions make the discount instrument more valuable, so platform profit falls as δ increases. Beyond δ_0, the platform abandons discounts, and weaker distortions improve outcomes under price salience. This non-monotonicity underscores that consumer salience presents both a challenge and an opportunity: while severe distortions are harmful absent intervention, they become profit-enhancing when the platform deploys instruments that strategically reshape attention. This creates incentives for the platform not only to respond to consumer salience but potentially to influence it through interface design, framing, and promotional policies.

Sensitivity Analysis

We further examine how the platform calibrates the reach and depth of surprise discounts in response to market conditions. Because closed-form comparative statics on the joint optimization over (w, β) are analytically intractable, we characterize these patterns through numerical analysis across a wide range of parameter values. The results, illustrated in Figures 3 and 4, are robust across all parameterizations we examined and reveal a consistent economic logic.

This logic generates the platform's consistent trade-off between maximizing the size of the quality-focused non-recipient group to earns higher commission revenue on premium-priced transactions and offering sufficient discounts to prevent the low-quality seller from cutting prices and compete for non-recipients, which we summarize as follows.

Salience severity. When salience distortions are weaker (higher δ), consumers' attention is less responsive to context, so the low-quality seller can sustain relatively high margins even under the price-salience equilibrium. This strengthens its incentive to deviate from the hybrid-salience equilibrium and compete aggressively for non-recipients. To prevent such deviation, the platform must make the recipient group sufficiently attractive to retain the low-quality seller. At the same time, weaker salience distortions also mean that quality-focused non-recipients discount price less

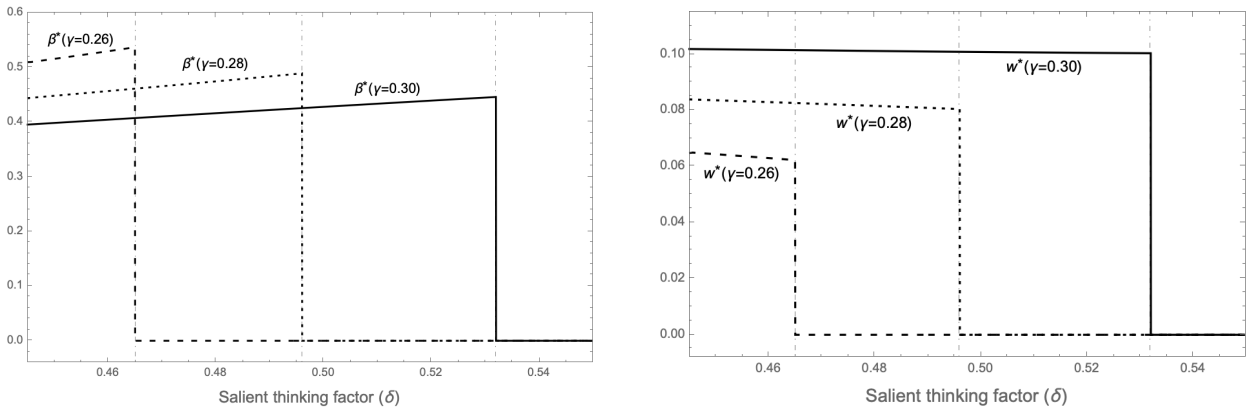

(a) Optimal discount reach β (b) Optimal discount depth w

Figure 4: Optimal discount policy changes with salience factor and commission rate

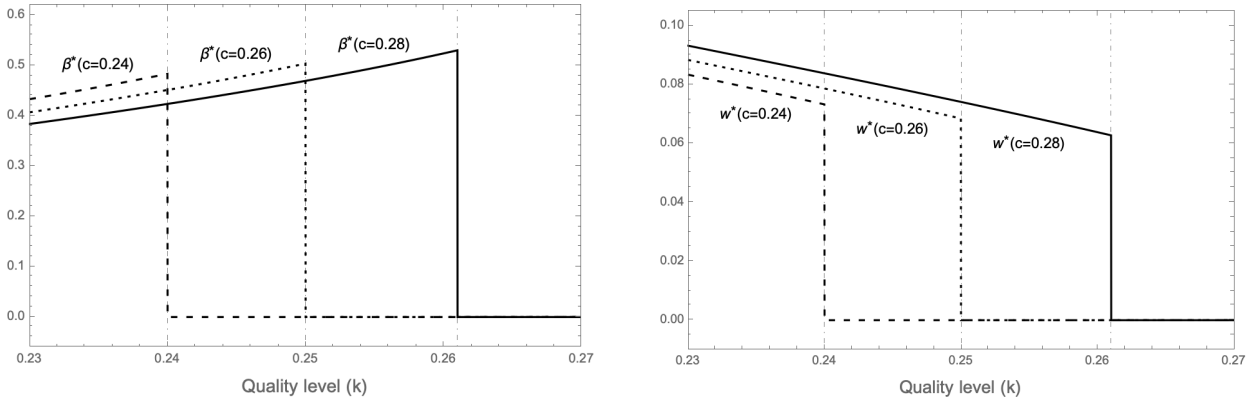

(a) Optimal discount reach β (b) Optimal discount depth w

Figure 5: Optimal discount policy changes with quality and cost

heavily in their evaluations, reducing the premium the high-quality seller can extract and thus the commission revenue the platform earns from this group. The diminished value of non-recipients, combined with the need to maintain a viable low-quality seller, leads the platform to expand discount reach and reduce discount depth. As salience distortions continue to weaken, these adjustments become increasingly costly until intervention is no longer profitable, at which point the platform abandons discounts entirely.

Commission rate. When the commission rate is high (larger γ), the platform captures a larger share of the high-quality seller's margin on quality-focused non-recipients, making this group particularly valuable. The platform therefore places greater priority on preserving a large non-recipient base and relies on deeper discounts targeted to a smaller group of recipients to keep the

low-quality seller incentive-compatible. The narrower reach conserves the lucrative non-recipient base, while deeper discounts ensure the recipient group remains sufficiently attractive to deter the low-quality seller from deviating. The net effect is higher platform profit.

Vertical differentiation. When products are closer in quality (higher k) or the production cost is lower (lower c), the low-quality seller becomes more competitive and more inclined to undercut prices to attract non-recipients. The platform responds by enlarging the recipient group to strengthen the low-quality seller's incentive to remain in the hybrid regime while reducing depth, since its smaller quality disadvantage means a shallower discount suffices to keep it profitably focused on recipients. Despite the efficient use of discount depth, the broader reach reduces the non-recipient base, lowering platform profit.

5 Extensions

5.1 Asymmetric Marginal Costs

Our main analysis assumes equal marginal costs across sellers to isolate the demand-side consequences of salience and simplify the analysis. A natural concern is whether the results are robust when the high-quality seller faces higher production costs, as is common in many markets. We examine this by allowing the two sellers to have distinct marginal costs c_H and c_L, with $c_H > c_L$, while maintaining the assumption that providing high quality is more efficient $q_H/c_H > q_L/c_L$ as in Bordalo et al. (2016). Here, we normalize marginal costs to $(c_H, c_L) = (c, tc)$, where $k < t < 1$.

The analysis proceeds as in the main model. In the benchmark without discounts, rational consumers segment between two sellers based on quality valuations, with each seller's equilibrium price now reflecting its own marginal cost. Under salient thinking, the key competitive forces are unchanged: the low-quality seller is strictly dominated under quality salience, and the price-salience equilibrium continues to arise as in main model.

When the platform introduces surprise discounts, the fixed-value discount still amplifies the effective price ratio and can induce hybrid salience. The equilibrium prices under each regime now depend on sellers' marginal costs, but the core mechanism is unchanged: the discount creates heterogeneous evaluation contexts, the low-quality seller acquires a captive recipient base that relaxes its pricing pressure, and both sellers maintain higher prices.

Lemma 5 (Equilibria under Asymmetric Marginal Costs) *(i) If the discount policy (w, β) induces a price-salience equilibrium, then the equilibrium prices are*

$$p_H^* = \frac{(2\delta + \beta w)(1-k)(1-\gamma) + (2+t)c}{(4-k)(1-\gamma)},$$
$$p_L^* = \frac{(k\delta + 2\beta w)(1-k)(1-\gamma) + (2t+k)c}{(4-k)(1-\gamma)},$$

under which the platform finds it optimal to set $w = 0$.

(ii) If the discount policy (w, β) induces a hybrid-salience equilibrium, then the equilibrium prices are

$$p_H^* = \frac{\beta(1-k)(1-\gamma)w + \beta c(2 - 2\delta^2(1-k) + t) + 2\delta(1-k)(1 + c\delta - \gamma)}{(1-\gamma)(\beta(4 - 4\delta^2(1-k) - k) + 4\delta^2(1-k))},$$
$$p_L^* = \frac{\beta(1-\delta^2(1-k))(c(k+2t) + 2(1-k)(1-\gamma)w) + \delta(1-k)(k + c\delta k - k\gamma + 2c\delta t + 2\delta(1-k)(1-\gamma)w)}{(1-\gamma)(\beta(4 - 4\delta^2(1-k) - k) + 4\delta^2(1-k))},$$

and the platform finds it optimal to set $w(\beta) = \frac{2\delta(1-k)(k+c\delta k - k\gamma - 2c\delta t) + 2\beta c(\delta^2 k^2 - 2(1-\delta^2)t + k(1+t-\delta^2(1+2t)))}{\sqrt{\beta}(1-k)(1-\gamma)(\beta(4-k) + 4(1-\beta)\delta^2(1-k) - \sqrt{\beta}k)}$, a decreasing function in β.

We solve the asymmetric-cost model numerically across a range of cost ratio t. The results confirm that the main findings are qualitatively robust. The threshold structure of the optimal discount policy persists: for sufficiently severe salience distortions, the platform profitably induces hybrid salience, and for mild distortions, it abstains. The reach-depth trade-off in the platform's optimal policy exhibits the same comparative static patterns described in the main analysis.

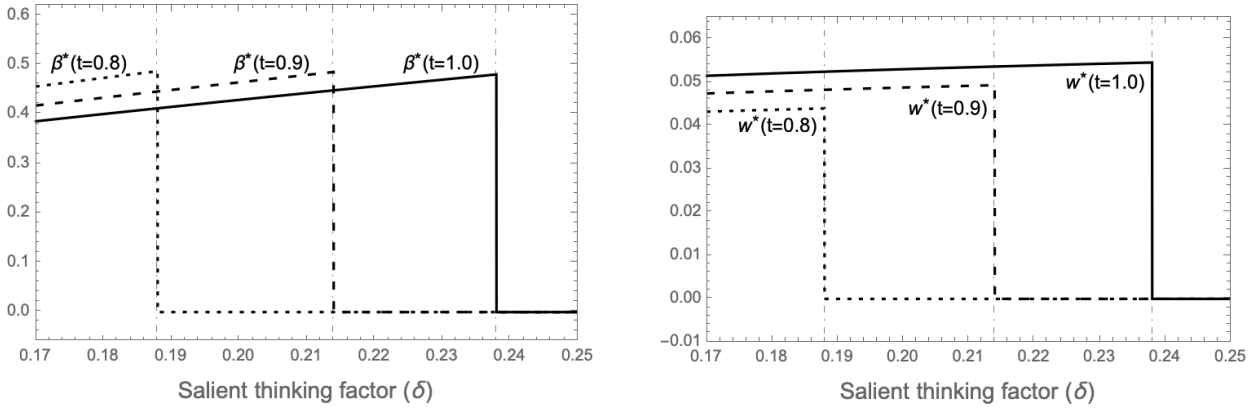

(a) Optimal discount reach β (b) Optimal discount depth w

Figure 6: Optimal discount policy changes with salience factor and cost ratios

When the production cost ratio is smaller (lower t), the low-quality seller becomes more competitive and more inclined to undercut prices to attract non-recipients. The platform responds

by enlarging the recipient group while reducing depth to strengthen the low-quality seller's incentive to remain in the hybrid regime, since its smaller quality disadvantage means a shallower discount suffices to keep it profitably focused on recipients. Despite the efficient use of discount depth, the broader reach reduces the non-recipient base, lowering platform profit.

5.2 Heterogeneous Salient Thinking

In the main model, all consumers share a common salience parameter. In practice, consumers likely vary in their degree in salient thinking — some exhibit stronger limited attention and thus are more context dependent than others. We examine robustness by considering a market in which a fraction $\alpha \in (0,1)$ of consumers are salient thinkers (with parameter $\delta < 1$) and the remaining fraction $1-\alpha$ are fully rational ($\delta = 1$).

Both types of consumers observe the same posted prices and, if applicable, receive the same discount. However, rational consumers are unaffected by the salience mechanism: their product evaluations depend only on absolute quality and price, regardless of the choice context. Salient consumers continue to overweight the more prominent attribute as in the main model. This heterogeneity adds a new dimension: the platform's discount instrument affects how salient consumers allocate attention but has no attention-shifting effect on rational consumers, for whom it functions purely as a price reduction.

When the platform introduces surprise discounts, the hybrid-salience mechanism operates on the salient consumer subpopulation. Among salient consumers, recipients perceive price as salient while non-recipients may perceive quality as salient, as in the main model. Rational consumers, whether or not they receive a discount, continue to evaluate products without distortion. The platform's instrument therefore generates attention heterogeneity within the salient subpopulation while leaving rational consumers' evaluations unchanged.

Lemma 6 (Hybrid-Salience Equilibria under Heterogeneous Salient Thinking) *(i) If the discount policy (w,β) induces a price-salience equilibrium, then the equilibrium prices are*

$$p_H^* = \frac{-\alpha\left(c(2+t)+\beta(k-1)(\gamma-1)w\right)+\delta\left(c(\alpha-1)(2+t)+(k-1)(\gamma-1)\left(-2+\beta(\alpha-1)w\right)\right)}{(k-4)\left(\delta(\alpha-1)-\alpha\right)(\gamma-1)},$$

$$p_L^* = \frac{-\alpha\left(c(k+2t)+2\beta(k-1)(\gamma-1)w\right)+\delta\left(-k^2(\gamma-1)+2(\alpha-1)\left(ct-\beta(\gamma-1)w\right)+k\left(c(\alpha-1)+(\gamma-1)\left(1+2\beta(\alpha-1)w\right)\right)\right)}{(k-4)\left(\delta(\alpha-1)-\alpha\right)(\gamma-1)},$$

under which the platform finds it optimal to set $w=0$.

(ii) If the discount policy (w, β) induces a hybrid-salience equilibrium, then the equilibrium prices are

$$p_H^* = \frac{-2(-1+\beta)c\delta^2(k-1)\alpha - \beta\alpha[3c+(k-1)(\gamma-1)w] + \delta[3c(\alpha-1)+(k-1)(\gamma-1)(-2+\beta(\alpha-1)w)]}{\left[\delta(-4+k)(\alpha-1) - \beta(-4+k)\alpha + 4(-1+\beta)\delta^2(k-1)\alpha\right](\gamma-1)},$$

$$p_L^* = \frac{c(2+k)[\delta(1-\alpha)+\beta\alpha][\delta+\beta\alpha+\delta(1-(1-\beta)\delta(1-k))\alpha] + (1-k)(1-\gamma)[\delta k(\delta(1-\alpha)+\beta\alpha) + 2\beta(\delta(1-\alpha)+\alpha)(\delta+\beta\alpha+\delta(1-(1-\beta)\delta(1-k))\alpha)w]}{[\delta(1-\alpha)+\beta\alpha][\delta(4-k)(1-\alpha)+\beta(4-k)\alpha+4(1-\beta)\delta^2(1-k)\alpha](1-\gamma)},$$

and the optimal discount is $w(\beta) = \frac{2[c(-2\delta(\alpha-1)+2\beta\alpha+(-1+\beta)\delta^2(-2+k)\alpha)+\delta k(\gamma-1)]}{\beta[\delta k(\alpha-1)-k\alpha+\sqrt{\frac{\delta+\alpha-\delta\alpha}{\delta+\beta\alpha-\delta\alpha}}(\delta(-4+k)(\alpha-1)-\beta(-4+k)\alpha+4(-1+\beta)\delta^2(k-1)\alpha)](\gamma-1)}$*, a decreasing function in* β*.*

Parallel to Proposition 2, we find a threshold policy in which the platform offers surprise discounts when salient consumers are sufficiently prevalent.

Proposition 3 *There exists a threshold α_0 such that*

(i) for $\alpha \geq \alpha_0$, the platform's optimal discount policy is characterized by unique β^* and corresponding $w^* = w(\beta^*)$, inducing a hybrid-salience equilibrium among salient consumers, where the platform's profit increases with α;

(ii) for $\alpha < \alpha_0$, the platform's optimal policy is to offer no discount, inducing a price-salience equilibrium among salient consumers, where the platform's profit decreases with α.

This result mirrors the threshold structure in Proposition 3, with the prevalence of salient consumers α playing a role analogous to the intensity of salience distortion δ. When salient consumers constitute a sufficiently large share of the market ($\alpha \geq \alpha_0$), the platform's attention-management mechanism operates on a meaningful portion of demand, justifying the cost of intervention.

Conversely, when salient consumers are relatively scarce ($\alpha < \alpha_0$), the rational majority anchors pricing close to the standard vertical differentiation outcome, limiting the distortions caused by salience. The platform's discount instrument then affects too small a share of the market to be profitable, and the platform optimally refrains from intervention. In this case, an increase in α intensifies attention competition without triggering intervention, thereby reducing platform profit.

The resulting non-monotonic relationship between α and platform profit reinforces the central insight of the paper: consumer salience is both a challenge and an opportunity. When salient consumers are few, they primarily erode margins; when they are sufficiently prevalent, the platform can profitably exploit their attention patterns through surprise discounts. These findings

demonstrate that our main results are robust to consumer heterogeneity and do not rely on the assumption of homogeneous salience.

5.3 Targeted Surprise Discounts

In the main model, we abstract away the platform's capability to offer targeted discount to consumers to isolate the strategic value of surprise discount in managing consumer attention. In practice, platforms sometime offer targeted discounts to specific consumer groups based on observable characteristics such as browsing history or past purchases. We examine targeted discounts to understand whether the platform can improve upon random allocation and to clarify the distinct role of attention management in our mechanism.

Suppose the platform can condition discount allocation on consumers' quality valuations θ. In line with the literature on targeted discount (Villas-Boas 2004, Acquisti and Varian 2005), we consider a policy that offers a fixed-value discount w to β fraction of consumers with $\theta < \theta_0$ for some threshold $\theta_0 \in (0, 1)$, while consumers with $\theta > \theta_0$ pay the posted prices. This targets discounts to low-valuation consumers, who might not make any purchases in the absence of discounts.

Targeted discounts with rational consumers. We first establish that, unlike surprise discounts (Lemma 1), targeted discounts benefit the platform even when consumers evaluate products rationally.

Lemma 7 *The platform always benefits from offering discounts that target low-valuation consumers with $\beta^* = 1$, even when consumers are fully rational.*

This is in contrast with Lemma 1. Surprise discounts subsidize consumers indiscriminately, including high-valuation consumers who would have purchased at the posted price regardless, and sellers partially offset the subsidy by raising posted prices. The net effect is a costly transfer with no strategic benefit. Targeted discounts, by contrast, concentrate the subsidy on the extensive margin: low-valuation consumers whose willingness to pay falls below the posted price of the low-quality product. By drawing these consumers into the market, the platform generates incremental transactions and earns commissions on sales that would not otherwise occur. This demand-expansion mechanism operates independently of salience and establishes that the value

of targeting stems from efficient subsidy allocation, not from attention management.

Targeted surprise discounts with salient consumers. Under salient thinking, targeted surprise discounts combine two effects: the demand-expansion effect described above and the attention-management mechanism from the main analysis. With targeting, however, the platform avoids subsidizing high-valuation consumers who would not patronize the low-quality seller and thus do not contribute to sustaining the hybrid-salience equilibrium. This cost saving enhances the platform's ability to profitably induce hybrid salience.

Proposition 4 *Facing salient consumers, the platform profits from offering surprise discounts to low-valuation consumers with* $\beta^* < 1$.

The intuition is that targeting eliminates a source of waste inherent in random allocation. Under surprise discounts, the platform subsidizes some high-valuation consumers who, being predisposed toward the high-quality product, do not help sustain the low-quality seller's captive recipient base. These discounts are costly without contributing to the incentive compatibility condition that prevents the low-quality seller from deviating. Targeting removes this inefficiency by directing the subsidy toward low-valuation consumers who are more likely to purchase from the low-quality seller and whose participation strengthens the recipient base that anchors the hybrid-salience equilibrium.

Numerical analysis confirms that, relative to random allocation, the platform under targeting offers deeper discounts (larger w) to a smaller group of low-valuation consumers (smaller β). The deeper discount reflects the platform's ability to redirect the savings from eliminating wasteful subsidies toward more intensive attention manipulation within the targeted group. The narrower allocation preserves a larger pool of non-recipients, including the high-valuation consumers no longer receiving unnecessary discounts, from whom the platform earns higher commission on sales of high-quality products.

6 Conclusion

This paper examines the strategic value of platform-issued surprise discounts in markets where consumers exhibit salient thinking. We develop an analytical model of a commission-based platform hosting two vertically differentiated sellers, where consumers allocate disproportionate

attention to whichever product attribute, price or quality, is more prominent in their evaluation context. Our analysis reveals that salient thinking traps the market in a price-focused equilibrium that depresses margins for both sellers and reduces the platform's commission income. We show that fixed-value surprise discounts offer the platform a powerful instrument to reshape this competitive landscape.

The central mechanism operates through the creation of heterogeneous evaluation contexts. A fixed-value voucher, by reducing both products' prices by the same absolute amount, amplifies the effective price ratio for discount recipients while leaving non-recipients' evaluation context unchanged. When calibrated appropriately, this generates a hybrid-salience equilibrium in which recipients attend to price differences and non-recipients attend to quality differences. The resulting coexistence of distinct attention patterns within a single market relaxes the competitive pressure that drives the price-salience trap. The low-quality seller, anchored by a base of price-focused recipients, reduces its incentive to undercut prices aggressively. The high-quality seller, facing quality-focused non-recipients, can sustain premium pricing. Both sellers charge higher posted prices than in the no-discount benchmark, and the platform's commission gains more than compensate for the cost of subsidizing discounted transactions.

The platform's optimal policy balances two competing objectives: maximizing the pool of quality-focused non-recipients who generate high commission revenue, and providing sufficient discounts to deter the low-quality seller from deviating back to aggressive price competition. This balance gives rise to a threshold structure that the platform intervenes only when salience distortion is sufficiently severe and produces a U-shaped relationship between salience severity and platform profit. Within the intervention region, the platform calibrates discount depth and allocation probability in response to market primitives: stronger salience distortion, higher commission rates, and greater quality differentiation all shift the optimal policy toward deeper discounts allocated to fewer consumers.

Our findings carry several managerial implications. For platform operators, surprise discounts should be understood not merely as demand-stimulation devices but as instruments of attention governance. The format of the discount matters: fixed-value vouchers possess strategic properties that percentage-based discounts lack, because only the former reshapes relative price comparisons across products. The allocation design also matters: random distribution, even without consumer targeting, generates the evaluation-context heterogeneity necessary to sustain profitable market

configurations.

For sellers, our analysis suggests that platform-initiated surprise discount programs can be beneficial as they raise posted prices. Both sellers earn higher profits under the hybrid-salience equilibrium than under the no-discount benchmark, with the high-quality seller benefiting disproportionately. Sellers should therefore view surprise discounts as a form of platform-mediated competitive relief rather than as a threat to pricing autonomy.

More broadly, our results reveal that consumer salience is not merely a source of market distortion. Instead, it is a strategic variable that market intermediaries can actively manage. Severe attention distortions, which are detrimental in the absence of intervention, become a source of profit when platforms deploy appropriate instruments. This creates incentives for platforms not only to respond to existing patterns of consumer attention but potentially to shape them through interface design, information architecture, and promotional framing.

References

Acquisti, A. and Varian, H. R. (2005). Conditioning prices on purchase history. *Marketing Science*, 24(3):367–381.

Bawa, K. and Shoemaker, R. W. (1987). The effects of a direct mail coupon on brand choice behavior. *Journal of marketing Research*, 24(4):370–376.

Bawa, K. and Shoemaker, R. W. (1989). Analyzing incremental sales from a direct mail coupon promotion. *Journal of marketing*, 53(3):66–78.

Blake, T., Moshary, S., Sweeney, K., and Tadelis, S. (2021). Price salience and product choice. *Marketing Science*, 40(4):619–636.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2012). Salience theory of choice under risk. *The Quarterly Journal of Economics*, 127(3):1243–1285.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, 121(5):803–843.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2016). Competition for attention. *The Review of Economic Studies*, 83(2):481–513.

Bushong, B., Rabin, M., and Schwartzstein, J. (2021). A model of relative thinking. *The Review of Economic Studies*, 88(1):162–191.

DelVecchio, D., Krishnan, H. S., and Smith, D. C. (2007). Cents or percent? the effects of promotion framing on price expectations and choice. *Journal of marketing*, 71(3):158–170.

Gigerenzer, G. (2007). *Gut eelings: The Intelligence of the Unconscious*. Viking, New York.

Guan, X., Jiang, Y., and Wang, Y. (2026). Platform certification and seller disclosure in online selling. *Manufacturing & Service Operations Management*, 28(2):381–399.

Huang, R., Lai, G., Wang, X., and Xiao, W. (2025). Platform financing vs. trade credit for lending to third-party sellers. *Management Science*, 71(7):5589–5604.

Inderst, R. and Obradovits, M. (2020). Loss leading with salient thinkers. *The RAND Journal of Economics*, 51(1):260–278.

Kőszegi, B. and Szeidl, A. (2013). A model of focusing in economic choice. *The Quarterly journal of economics*, 128(1):53–104.

Narasimhan, C. (1984). A price discrimination theory of coupons. *Marketing Science*, 3(2):128–147.

Puyang, X., Gu, Z., Chen, R. R., and Li, J. (2026). Vertical competition on a common platform. *Marketing Science*, 45(1):175–188.

Shaffer, G. and Zhang, Z. J. (1995). Competitive coupon targeting. *Marketing Science*, 14(4):395–416.

Shi, Z. and Raghu, T. (2020). An economic analysis of product recommendation in the presence of quality and taste-match heterogeneity. *Information Systems Research*, 31(2):399–411.

Sun, G., Cavusoglu, H., and Raghunathan, S. (2025). Unveiling the strategic impacts of extending membership-based free shipping programs beyond the online marketplaces. *Information Systems Research*.

Villas-Boas, J. M. (2004). Price cycles in markets with customer recognition. *RAND Journal of Economics*, pages 486–501.

Zheng, Q., Pan, X. A., and Carrillo, J. E. (2019). Probabilistic selling for vertically differentiated products with salient thinkers. *Marketing Science*, 38(3):442–460.

Zhu, R., Moorthy, S., and Shi, X. (2025). Advertising platforms and privacy. *Marketing Science*, 44(6):1381–1406.

Zhu, Y. and Dukes, A. (2017). Prominent attributes under limited attention. *Marketing Science*, 36(5):683–698.

Appendix of Platform's Surprise Discounts with Salient Consumers

A.1 Percentage-Based Discounts

Our main analysis establishes the strategic value of surprise discounts as a tool for the platform to steer consumer attention and shape seller competition. A natural question is whether this value is specific to fixed-value discounts, or whether percentage-based discounts that are also prevalent in practice, can serve the same role. We show formally that they cannot.

Suppose the platform offers a percentage-based discount $\alpha \in (0, 1)$ instead of a fixed-value discount. A recipient purchasing product j at posted price p_j pays the effective price $(1-\alpha)p_j$. The key observation is that a percentage discount scales both prices by the same factor, preserving the price ratio: $(1-\alpha)p_H/((1-\alpha)p_L) = p_H/p_L$. Because the salience attribute depends on whether the price ratio exceeds the quality ratio, a recipient's salience perception is identical to that of a non-recipient. The discount therefore cannot shift consumer evaluation contexts; all consumers, regardless of whether they receive a discount, perceive the same attribute as salient.

Without the ability to create distinct attention patterns across consumers, a percentage-based discount loses its strategic value. It effectively reduces the platform's commission rate, since the platform earns commissions from both sellers' revenue. As in the rational-consumer case (Lemma 1) and the price-salience case (Lemma 3): sellers anticipate the discount and adjust posted prices upward, partially absorbing the subsidy. The platform bears the cost of the discount without inducing any shift in the competitive regime.

This result sharpens the contribution of the main analysis. The strategic value of surprise discounts does not arise from price reductions per se. Instead, it stems from the different effective price ratio that generates heterogeneous evaluation contexts across consumers. Fixed-value discounts are uniquely capable of this because they compress price ratios by a magnitude that varies with the posted price level, a property that percentage-based discounts inherently lack.

A.2 Proof of Main Results

Proof of Proposition 1

We derive the equilibrium pricing strategies of the two sellers. We begin by ruling out the existence of a quality-salience equilibrium with $p_L > c$.

Suppose, for contradiction, that there exists a quality-salience equilibrium in which $p_L > c$. Under quality salience, the low-quality seller can not attracts any demand thus earns zero profit. Instead, by slightly lowering its price, it can deviate to the price-salience regime and attract positive demand. This deviation contradicts equilibrium. Therefore, no quality-salience equilibrium with $p_L > c$ exists.

We thus focus on the price-salience equilibrium. In this regime, both products are evaluated with price as the salient attribute, and the sellers' profit functions are given by

$$\pi_H^R = ((1-\gamma)p_H - c)(1 - \frac{p_H - p_L}{\delta(q_H - q_L)}),$$
$$\pi_L^R = ((1-\gamma)p_L - c)(\frac{p_H - p_L}{\delta(q_H - q_L)} - \frac{p_L}{\delta q_L}).$$

Solving the two sellers' profit maximization problems simultaneously yields a unique Nash equilibrium. The equilibrium prices are $p_H^S = \frac{2\delta(1-k)(1-\gamma)+3c}{(4-k)(1-\gamma)}$ and $p_L^S = \frac{k\delta(1-k)(1-\gamma)+(2+k)c}{(4-k)(1-\gamma)}$. In equilibrium, both sellers must attract strictly positive demand, which imposes the condition $c \leq \bar{c} \equiv \delta k(1-\gamma)/2$.

Proof of Proposition 2

Recall that consumers who receive the discount are more likely to perceive price as the salient attribute. This gives rise to three potential equilibrium configurations: (i) a quality-salient equilibrium in which both groups perceive quality as salient, (ii) a hybrid equilibrium in which recipients exhibit price as salient while non-recipients exhibit quality as salient, and (iii) a price-salient equilibrium in which both groups perceive price as salient.

We then rule out the existence of a quality-salient equilibrium under surprise discounts. Suppose, for contradiction, that such an equilibrium exists. In this case, the low-quality seller attracts no demand and earns zero profit. However, given that a fraction β of consumers receives the discount, a price reduction by the low-quality seller would make price sufficiently salient for

these consumers and allow the seller to capture positive demand. This profitable deviation to the hybrid-salience regime contradicts the existence of a quality-salient equilibrium. Therefore, only the price-salient equilibrium and the hybrid-salience equilibrium can arise in the presence of surprise discounts.

Next, consider the two sellers' pricing equilibrium. The high-quality seller's profit function is given by

$$\pi_H^D = \begin{cases} ((1-\gamma)p_H - c)\Big(1 - \frac{p_H - p_L}{\delta(q_H - q_L)}\Big), & \text{if } p_H > \frac{p_L}{k} < p_H, \\ ((1-\gamma)p_H - c)\Big(\beta\Big(1 - \frac{p_H - p_L}{\delta(q_H - q_L)}\Big) + (1-\beta)(1-\delta p_H)\Big), & \text{if } \frac{p_L - (1-k)w}{k} < p_H < \frac{p_L}{k}. \end{cases}$$

Similarly, the low-quality seller's decision problem is

$$\pi_L^D = \begin{cases} ((1-\gamma)p_L - c)\Big(\beta\Big(\frac{p_H - p_L}{\delta(q_H - q_L)} - \frac{p_L - w}{\delta q_L}\Big) + (1-\beta)\Big(\frac{p_H - p_L}{\delta(q_H - q_L)} - \frac{p_L}{\delta q_L}\Big)\Big), & \text{if } p_L < kp_H, \\ ((1-\gamma)p_L - c)\beta\Big(\frac{p_H - p_L}{\delta(q_H - q_L)} - \frac{p_L - w}{\delta q_L}\Big), & \text{if } kp_H < p_L < kp_H + (1-k)w. \end{cases}$$

Solving the two sellers' profit maximization problems simultaneously yields two possible Nash equilibria. If $w \leq \frac{2\delta k(1-\gamma)-4c}{\sqrt{\beta}(4-k\beta-k\sqrt{\beta})(1-\gamma)}$, there exists a price-salience equilibrium, in which both consumer groups perceive price as salient. The equilibrium prices are $p_H^N = \frac{(2\delta+\beta w)(1-k)(1-\gamma)+3c}{(4-k)(1-\gamma)}$ and $p_L^N = \frac{(k\delta+2\beta w)(1-k)(1-\gamma)+(2+k)c}{(4-k)(1-\gamma)}$. If $w \geq \frac{2\beta c(2-\delta^2(2-k))+2\delta(c\delta(2-k)-k+k\gamma)}{\sqrt{\beta}(\beta(4-4\delta^2(1-k)-k)+4\delta^2(1-k)-\sqrt{\beta}k)(1-\gamma))}$, there exists a hybrid-salience equilibrium, in which discount recipients perceive price as salient while non-recipients perceive quality as salient. Prices are $p_H^D = \frac{2\delta(1-k)(1+c\delta-\gamma)+\beta(c(3-2\delta^2(1-k))+(1-k)(1-\gamma)w)}{(\beta(4-4\delta^2(1-k)-k)+4\delta^2(1-k))(1-\gamma)}$ and $p_L^D = \frac{\beta(1-\delta^2(1-k))(c(2+k)+2(1-k)(1-\gamma)w)+\delta(1+k)(c\delta(2+k)+(1-\gamma)(2\delta w+k(1-2\delta w)))}{(\beta(4-4\delta^2(1-k)-k)-4\delta^2(1-k))(1-\gamma))}$.

Proof of Lemma 3

Suppose that $w \leq \frac{2\delta k(1-\gamma)-4c}{\sqrt{\beta}(4-k\beta-k\sqrt{\beta})(1-\gamma)}$. In the resulting price-salience equilibrium, the platform's profit function is given by $\pi_P^D(w;\beta) = \Big[c(1-\gamma)(\delta k(4+5k)\gamma + \beta(8-k^2(1-\gamma)+k(2+8\gamma))w) - c^2(4+5k)\gamma - (1-\gamma)^2(\delta^2 k(4-3k-k^2)\gamma + \beta\delta k(12-8\gamma-k(3-8\gamma))w + \beta((4-k)^2 - \beta(1-k)(4(2+\gamma)-k(2-\gamma)))w^2)\Big] \Big/ \Big[\delta(4-k)^2 k(1-\gamma)^2\Big]$. We can verify that $\partial\pi_P^D/\partial w = \beta\Big[c(8-k^2(1-\gamma)+k(2+8\gamma)) - (1-\gamma)(\delta k(12-8\gamma-k(3-8\gamma)) + 2((4-k)^2 - \beta(1-k)(8-2k+4\gamma+k\gamma))w)\Big] \Big/ \Big[\delta(4-k)^2 k(1-\gamma)\Big] < 0$ for $w \leq \frac{2\delta k(1-\gamma)-4c}{\sqrt{\beta}(4-k\beta-k\sqrt{\beta})(1-\gamma)}$, implying that the platform's profit is decreasing in w within this region. Therefore, the platform optimally sets $w = 0$.

Proof of Lemma 4

Suppose that $w \geq \frac{2\beta c(2-\delta^2(2-k))+2\delta(c\delta(2-k)-k+k\gamma)}{\sqrt{\beta}(\beta(4-4\delta^2(1-k)-k)+4\delta^2(1-k)-\sqrt{\beta}k)(1-\gamma))}$. In the resulting hybrid-salience equilibrium, the platform's profit function is given by $\pi_P^D(w;\beta) = \Big[(1-\beta)\delta(1-k)k\gamma(2\delta(1-k)(1+c\delta-\gamma)+\beta(c(3-2\delta^2(1-k))+(1-k)(1-\gamma)w))(2\delta^2(1-k)(1-c\delta-\gamma)-\beta(c\delta(3-2\delta^2(1-k))-(1-\gamma)(4-4\delta^2(1-k)-k-\delta(1-k)w)))-\beta(1-k)\gamma(\beta(c(2-\delta^2(2-k))-2(1-\delta^2(1-k))(1-\gamma)w)+\delta(c\delta(2-k)-(1-\gamma)(k+2\delta w-2\delta kw)))(\beta(1-\delta^2(1-k))(c(2+k)+2(1-k)(1-\gamma)w)+\delta(1-k)(c\delta(2+k)+(1-\gamma)(2\delta w+k(1-2\delta w))))+\beta k\gamma(2\delta(1-k)(1+c\delta-\gamma)+\beta(c(3-2\delta^2(1-k))+(1-k)(1-\gamma)w))(\delta(\beta(4-4\delta^2(1-k)-k)+4\delta^2(1-k))(1-k)(1-\gamma)-2\delta(1-k)(1+c\delta-\gamma)-\beta(c(3-2\delta^2(1-k))+(1-k)(1-\gamma)w)+\beta(1-\delta^2(1-k))(c(2+k)+2(1-k)(1-\gamma)w)+\delta(1-k)(c\delta(2+k)+(1-\gamma)(2\delta w+k(1-2\delta w))))-\beta(\beta(4-4\delta^2(1-k)-k)+4\delta^2(1-k))(1-k)(1-\gamma)w(-\beta(c(1-\delta^2(1-k))(2+k)-(1-\gamma)(\delta(4-k)k+4\delta^3(1-k)k-(2+k)w+2\delta^2(1-k^2)w))-\delta(1-k)(c\delta(2+k)-(1-\gamma)(2\delta w-k(1-4\delta^2-2\delta w))))\Big]\Big/\Big[\delta(\beta(4-4\delta^2(1-k)-k)+4\delta^2(1-k))^2(1-k)k(1-\gamma)^2\Big]$. We can verify that $\partial\pi_P^D/\partial w < 0$ for $w \geq \frac{2\beta c(2-\delta^2(2-k))+2\delta(c\delta(2-k)-k+k\gamma)}{\sqrt{\beta}(\beta(4-4\delta^2(1-k)-k)+4\delta^2(1-k)-\sqrt{\beta}k)(1-\gamma))}$, implying that the platform's profit is decreasing in w within this region. Therefore, the platform optimally sets $w = \frac{2\beta c(2-\delta^2(2-k))+2\delta(c\delta(2-k)-k+k\gamma)}{\sqrt{\beta}(\beta(4-4\delta^2(1-k)-k)+4\delta^2(1-k)-\sqrt{\beta}k)(1-\gamma))}$.

Proof of Proposition 2

For the price-salience equilibrium, the platform's optimal profit is $\pi_P^D = \Big[(c\delta k(4+5k)(1-\gamma)-c^2(4+5k)+\delta^2k(4-3k-k^2)(1-\gamma)^2)\gamma\Big]\Big/\Big[\delta(4-k)^2k(1-\gamma)^2\Big]$. It is easy to verify that $\partial\pi_P^D/\partial\delta = \frac{((4-3k-k^2)\delta^2k(1-\gamma)^2+c^2(4+5k))\gamma}{(4-k)^2\delta^2k(1-\gamma)^2} > 0$, implying that the platform's profit is increasing in δ in the price-salience equilibrium.

For the price-salience equilibrium, the platform's optimal profit is $\pi_P^D(\beta) = \Big[\beta^3c^2(-4-5k+4\delta^6(k-1)^2k+\delta^2(8+9k-8k^2)-\delta^4(4+8k-15k^2+3k^3))\gamma-2\sqrt{\beta}\delta^2(k-1)(-k^2(-1+4\delta^2-\gamma)(\gamma-1)^2+c^2\delta^2(-4-2k\gamma+k^2(1+2\gamma))+c\delta k(\gamma-1)(4\delta^2(k-2)+4\gamma-k(2+3\gamma)))+\delta^2(2c\delta(k-2)k(\gamma-1)(2-\gamma+k(2+\gamma))+k(\gamma-1)^2(4\delta^2\gamma+k(-2+\gamma-8\delta^2\gamma)+k^2(-2+(-1+4\delta^2)\gamma))-c^2\delta^2(8-4\gamma+4(2+\delta^2)k\gamma+k^3(2+\gamma+4\delta^2\gamma)-k^2(6+(5+8\delta^2)\gamma)))-2\beta^{5/2}c(\delta k(8-2k-\delta^2(16-14k+k^2)+4\delta^4(2-3k+k^2))(\gamma-1)+c(1+\delta^2(k-1))(4+k(2+3\gamma)+\delta^2(-4-2k\gamma+k^2(1+2\gamma))))+2\beta^{3/2}\delta(-\delta(4+4\delta^2(k-1)-k)k^2(\gamma-1)^2-ck(\gamma-1)(-8\delta^4(2-3k+k^2)-4\gamma+k(3+4\gamma)+\delta^2(k-4)(-4+k-\gamma+k\gamma))+c^2\delta(-8+k(2-5\gamma)+k^2(3+5\gamma)+2\delta^2(k-1)(-4-2k\gamma+k^2(1+2\gamma))))+\beta\delta(-\delta k(\gamma-1)^2(4(-1+\delta^2)\gamma+(3-8\delta^2)k\gamma+k^2(-2+\gamma+4\delta^2\gamma))+ck(\gamma-1)(4(-2+\gamma)-3k(2+$

$\gamma) + 2\delta^2(4 + k^2(-3 + \gamma) - 4\gamma + k(4 + 3\gamma))) + c^2\delta(8(-2 + \gamma) + 12\delta^4(k-1)^2k\gamma + 2k^2(3 + \gamma) - k(4 + 9\gamma) + \delta^2(16 - k^3(-4 + \gamma) - 12\gamma + 8k\gamma + k^2(-12 + 5\gamma)))) - \beta^2c(\delta k(\gamma - 1)(-4(-1 + \delta^2)\gamma + 2\delta^2k^2(-1 + 2\gamma) + k(-4 + 4\delta^2 + 5\gamma)) + c(8 - 2\delta^2(8 + 2k + 3k^2(\gamma - 1) - 8\gamma) - 4\gamma + k(4 + \gamma) + 12\delta^6(k-1)^2k\gamma + \delta^4(8 + k^3(2 - 5\gamma) - 12\gamma - 8k\gamma + k^2(-6 + 25\gamma))))\Big] \Big/ \Big[\delta k(4\delta^2(1-k) - \sqrt{\beta}k + \beta(4 - 4\delta^2(1-k) - k))^2(1-\gamma)^2\Big]$. We can verify that $\partial\pi_P^D/\partial\beta < 0$, implying that there exist a unique β that maximizes the platform's profit.

Following the envelope theorem, we can find $\partial\pi_P^D/\partial\delta < 0$, implying that the platform's profit is decreasing in δ in the hybrid-salience equilibrium.

Note that the platform never finds it optimal to offer surprise discounts when facing rational consumer ($\delta = 1$). Taking together, there exists a unique threshold δ_0 such that the platform offers a surprise discount (w^*, β^*) if $\delta < \delta_0$, and offers no discount otherwise.