Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

The Launch Premium: Cohort Comparison and the Pricing of Perishable Goods

Chris Doyle

chris.doyle@apexeconomics.com

London School of Economics and Political Science

Research Article

Keywords: Dynamic pricing, Launch premium, Reference dependence, Inequity aversion, Algorithmic pricing, Revenue management

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10987329/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

The Launch Premium:
Cohort Comparison and the Pricing of Perishable Goods

September 2026

Abstract

A buyer who pays more than an earlier buyer for the same flight, hotel stay or event may perceive the difference as unfair, but the benchmark is a price the firm itself chose. This paper studies dynamic pricing when the comparison benchmark is endogenous to the seller's earlier pricing decision. Consumers evaluate later prices against an expectation reference determined in personal equilibrium and a comparison reference created by the firm's initial price. Expectation-based loss aversion smooths prices, while comparison-based inequity aversion compresses price differences across cohorts. Because the initial cohort creates the benchmark without itself facing an earlier-cohort comparison, the firm optimally raises the launch price above the static monopoly benchmark to relax future behavioural constraints. Stronger interpersonal comparisons therefore generate both higher launch prices and smaller subsequent price differences: under empirically plausible parameters the launch price rises by 5 to 10 per cent while cross-cohort dispersion falls by 80 to 95 per cent. The premium is bounded rather than escalating, rising with comparison intensity only up to a threshold beyond which the firm prices at the benchmark in every state. Extending the model to an arbitrary number of repricing epochs with finite capacity, the launch premium is strictly positive at every repricing frequency, and numerically it increases as the selling window is divided more finely and is largest at intermediate scarcity. Finer algorithmic repricing therefore magnifies rather than eliminates the distortion. Cohort comparison is a distinct intertemporal force in dynamic pricing, with implications for launch pricing and algorithmic pricing systems.

JEL Codes: D91; L11; M31; D42; L51

Keywords: Dynamic pricing; Launch premium; Reference dependence; Inequity aversion; Algorithmic pricing; Revenue management.

1. Introduction

Firms selling fixed, perishable capacity face a familiar dynamic pricing problem: how should prices evolve as the consumption date approaches, demand is revealed and remaining capacity changes? Airlines sell seats before departure, hotels sell rooms for a particular night, and organisers sell tickets before an event. Standard models characterise optimal pricing in terms of economic state variables such as remaining capacity, demand conditions and time to expiry (Gallego and Van Ryzin, 1994; Talluri and Van Ryzin, 2004; Phillips, 2005). This paper identifies an additional intertemporal incentive. When later consumers compare the price they are offered with prices paid by earlier buyers, the seller's initial price does more than generate current revenue and allocate capacity: it creates a benchmark against which future prices will be judged.

The resulting incentive runs opposite to the usual behavioural distortion. Where reference points are formed from a consumer's own price history, behavioural concerns push the equilibrium price *below* the non-behavioural benchmark, since building a higher reference is costly to the consumer who supplies it. Here the opposite holds. A forward-looking seller may optimally set the initial, or launch, price *above* the static monopoly price, sacrificing current profit to create a more favourable comparison benchmark for subsequent pricing. At the static monopoly price a marginal increase has no first-order effect on current profit, but it raises the price against which later consumers evaluate their treatment. Whenever disadvantageous comparison affects future demand with positive probability, this continuation-value effect is positive, so the optimal launch price exceeds the static benchmark. We call this distortion the *launch premium*.

The mechanism depends on a sequential cohort structure rather than on physical perishability. Consumers purchasing a particular flight, hotel night or event typically buy that product only once, so those whose purchases establish the benchmark differ from those who later compare their prices with it. The same logic arises for finite editions and other non-replenishable inventories sold sequentially, provided earlier and later transactions are comparable and previous prices observable; physical expiry is a sufficient rather than necessary source of the finite horizon.

Popescu and Wu (2007) also show that a forward-looking firm has an incentive to price above its myopic optimum because current prices influence consumers' future reference prices. But here the cohort structure changes the economics of that incentive. In adaptive-reference models the same consumer observes prices over time and forms a reference from her own price history, so raising the current price builds a higher future reference but simultaneously creates a current perceived loss. Reference-building therefore carries a contemporaneous behavioural cost, and in Popescu and Wu the equilibrium price remains below the price that would prevail without reference effects. Here the consumers who establish the benchmark and those who later evaluate their prices against it belong to different cohorts. The firm obtains the future benefit of a higher benchmark without paying an interpersonal behavioural cost when that benchmark is created, and this separation reverses the sign of the distortion relative to the non-behavioural benchmark.

The model distinguishes two reference objects. Consumers evaluate their realised price relative to the price they expected to pay, an expectation reference determined as a rational-expectations personal equilibrium in the sense of Kőszegi and Rabin (2006), generating *loss*

aversion when realised prices exceed expectations. They may also compare their price with that paid by an earlier cohort, a comparison reference motivated by *inequity aversion* and by evidence that consumers judge price fairness partly by comparing their treatment with that of other buyers (Kahneman, Knetsch and Thaler, 1986; Fehr and Schmidt, 1999; Xia, Monroe and Cox, 2004; Richards, Liaukonyte and Streletskaya, 2016). Unlike the expectation reference, which is an equilibrium object determined jointly by beliefs and optimal pricing, the interpersonal benchmark is directly inherited from the firm's earlier decision.

The two channels differ in their implications. Expectation-based loss aversion attenuates the response of prices to demand shocks and can generate complete smoothing over a non-degenerate range of demand states. Interpersonal comparison also compresses price differences across cohorts, but it additionally creates the forward-looking incentive to elevate the initial price. Dynamic pricing is therefore distorted at both ends of the intertemporal problem, and stronger behavioural concerns produce the distinctive combination of higher launch prices and lower subsequent dispersion.

How strongly the comparison channel operates depends not only on consumers' preferences but on the market in which they transact, and two features govern it. The first is *identity salience*. Comparison bites harder where consumers can readily observe what others paid and where they regard those others as people like themselves, so that a price difference registers as differential treatment of a peer rather than as anonymous variation. Both conditions are needed: comparison must be available, and it must be felt to matter. The second is *procedural legitimacy*. A price difference arising from a rule consumers accept (advance booking, a published fare structure, a negotiation the buyer took part in) provokes less resentment than the same difference arising from an opaque procedure. Legitimacy discounts the behavioural penalty attaching to a price difference without altering the difference itself. Neither feature constitutes a further behavioural mechanism: both scale the intensity with which the comparison already described operates, in opposite directions, and the analysis accordingly works with a composite comparison intensity that embeds them. This is what allows a single framework to accommodate the observation that intertemporal price variation is broadly acceptable in airline markets yet contentious in ticketing and ride-sharing.

The paper formalises these mechanisms in a two-period model of a monopolist selling a perishable good under demand uncertainty. The first-period price is observed before the second-period demand state is realised and becomes the comparison reference for the second cohort. The optimal launch price lies strictly above the static monopoly price whenever disadvantageous comparison occurs with positive probability. The premium rises with the prevalence and intensity of comparison concerns up to a threshold; beyond it the firm prices at the benchmark in every state, so the distortion is bounded rather than escalating in the strength of consumers' fairness concerns. Greater procedural legitimacy works in the opposite direction to identity salience, weakening the behavioural penalty attaching to a given price difference and so reducing the incentive to elevate the benchmark.

The relationship between the two references matters. The expectation reference is not an exogenous anchor the firm can manipulate independently but must be consistent with the price consumers rationally expect in personal equilibrium. The paper shows that the strategically elevated launch price lies above it. This endogenous ordering determines which behavioural regions are relevant in equilibrium and distinguishes the mechanism from models in which a single adaptive reference summarises consumers' price history.

The two-period structure is a simplification in a specific direction that fixes how often the seller may reprice. Modern algorithmic systems, however, allow firms selling perishable capacity to update prices at intervals that were previously impractical. We therefore extend the model to a selling window divided into an arbitrary number of repricing epochs with finite capacity, holding total demand over the window fixed so that dividing it more finely means repricing more often rather than selling for longer. The benchmark is the lowest price previously paid, so the initial price establishes a floor inherited by all subsequent cohorts. Finer repricing strengthens the incentive to elevate that floor: the cost of a higher launch price is borne by an increasingly small initial cohort, while its benefit extends to increasingly many future cohorts. The launch premium is strictly positive at every repricing frequency, and in the numerical analysis it increases as the window is divided more finely, remaining bounded away from zero at every frequency considered.

Capacity scarcity enters non-monotonically. Where capacity is abundant, conventional scarcity considerations play little role and the value of manipulating the benchmark is limited by the future demand over which it operates. Where scarcity is severe, the seller already prices high for rationing reasons, leaving the comparison constraint slack. The premium is therefore largest at intermediate scarcity, which distinguishes the mechanism from fairness models in which scarcity itself generates the distortion. The multi-period model also generates a declining price path of the kind usually attributed to skimming, but for a different reason. The value of maintaining a high comparison floor is greatest when most cohorts are still to arrive, so as the selling window closes the firm progressively concedes it. A price path that falls from a high launch price therefore need not indicate strong early demand, an exogenously high reference, or sequential sorting of high-valuation buyers: it can arise purely from the behavioural benchmark losing value as the horizon shortens.

These results yield empirically distinguishable implications, because the model predicts a joint pattern rather than merely reduced price flexibility. Conditional on conventional determinants of pricing, markets in which interpersonal comparisons are stronger should exhibit higher launch prices and less subsequent dispersion; the mechanism should also be stronger where identity salience is high and procedural legitimacy low. Mechanisms based solely on expectation reference dependence predict smoothing but not the launch premium, so the coexistence of an elevated initial price and compressed subsequent variation distinguishes the two. A quantitative calibration, drawing parameter values from the behavioural literature and setting the behavioural share and comparison intensity at or below the range implied by experimental estimates, indicates that the optimal launch price exceeds the static benchmark by roughly 5 to 10 per cent while cross-cohort dispersion falls by 80 to 95 per cent. The asymmetry is itself informative: the comparison channel operates far more powerfully on the dispersion of prices than on their level, so a mechanism easily missed in launch prices alone is unmistakable in the subsequent path.

The paper contributes to two literatures. The first is dynamic pricing and revenue management, in which standard models analyse how current pricing affects future profits through sales, remaining capacity, demand uncertainty and, in some settings, strategic consumer behaviour (Stokey, 1979; Gallego and Van Ryzin, 1994; Talluri and Van Ryzin, 2004; Phillips, 2005). Here an earlier price has an additional role as an endogenous behavioural state variable governing the demand response of later cohorts, so pricing changes the behavioural environment in which subsequent capacity is sold.

Peer-induced fairness provides a second connection. Li and Jain (2016) also obtain higher first-period prices when consumers dislike paying more than others for the same product. Their mechanism, however, operates through poaching competition and initial-price sensitivity in a behaviour-based pricing environment. Here, the earlier price is itself the later cohort's fairness benchmark, so raising it directly relaxes a future comparison constraint. The launch premium therefore arises even for a monopolist without customer recognition, poaching or repeat purchasing. Selove (2019) is closest in combining fairness with dynamic pricing and capacity constraints, but the benchmark there is contemporaneous and embedded in an announced pricing policy: a realised price paid by one cohort does not become the inherited benchmark for a later one. The present mechanism instead operates sequentially through realised transaction prices and therefore explains the strategic level of the initial price as well as subsequent compression. Related work examines fairness towards firms' markups (Eyster, Madarász and Michaillat, 2021), expectation-induced price stickiness (Heidhues and Kőszegi, 2008), dynamic pricing under exogenous fairness constraints (Cohen, Miao and Wang, 2025), and distributional fairness between buyers and sellers (Guo, 2015); the comparison here is instead interpersonal, between differently timed consumers purchasing equivalent units.

The distinction matters increasingly as repricing becomes cheaper. Algorithmic pricing allows firms to update prices at ever finer intervals, which in a standard revenue-management model lets the seller respond more closely to market conditions. The analysis identifies a countervailing implication: more frequent repricing increases the number of future cohorts evaluated against prices established earlier, strengthening the incentive to elevate the launch benchmark. Greater technological flexibility need not translate one-for-one into greater equilibrium price flexibility when consumers compare their treatment with that of previous buyers.

Section 2 develops the economic and behavioural foundations. Section 3 presents the model, characterises equilibrium pricing and establishes the launch-premium result. Section 4 extends the analysis to multiple repricing epochs and perishable capacity. Section 5 develops the quantitative calibration, empirical implications and policy consequences, and sets out an empirical strategy for testing the mechanism. Section 6 concludes. The appendices contain proofs, calibration details, computational methods, robustness analysis and replication information.

2. The Evolution of Dynamic Pricing and Behavioural Reference Formation

The model combines the standard dynamic-pricing problem of allocating perishable capacity over time with two behavioural responses to prices: consumers evaluate prices relative to what they expected to pay and relative to prices paid by earlier buyers. This section motivates these mechanisms and identifies the empirical regularities the model represents.

2.1 Dynamic Pricing and Endogenous Reference Points

Dynamic pricing originated in the deregulated US airline industry and has since spread through hotels, ride-sharing, e-commerce, event ticketing and retail (Smith, Leimkuhler and Darrow, 1992; Courty, 2003; Ferreira, Lee and Simchi-Levi, 2016). The fundamental problem is dynamic

because current pricing affects sales and hence the capacity available later (Gallego and Van Ryzin, 1994; Talluri and Van Ryzin, 2004).

Behavioural evidence adds a second channel. Consumers evaluate prices relative to expectations, previous prices and the treatment of other consumers (Kahneman, Knetsch and Thaler, 1986; Tversky and Kahneman, 1991; Xia, Monroe and Cox, 2004; Kőszegi and Rabin, 2006). A current price can therefore affect future demand not only through remaining capacity but by shaping the reference points against which later prices are judged. Existing behavioural pricing models establish that firms internalise endogenous reference formation (Popescu and Wu, 2007; Nasiry and Popescu, 2011) and that fairness concerns constrain price variation (Selove, 2019). The distinction required here is between a reference generated by consumers' own expectations, determined in personal equilibrium, and one created by an actual transaction with a different cohort. The latter makes an earlier price a behavioural state variable governing the demand of subsequent consumers.

2.2 Consumer Responses to Dynamic Prices

Consumer responses vary substantially across markets, suggesting that reactions depend not on the magnitude of price changes alone but on how those changes are interpreted relative to expectations, previous transactions and the prices paid by comparable consumers. The 2024 sale of Oasis tickets in the United Kingdom generated widespread perceptions of opportunistic pricing and investigations by the Competition and Markets Authority and the European Commission, despite reports that real-time algorithmic adjustment was not used; consumers nevertheless encountered substantial price differences during the sales process (BBC News, 2024; CMA, 2024, 2025). Wendy's announcement of dynamic pricing was interpreted as peak-time surcharging and subsequently reframed as off-peak discounting (Leasca, 2024), while Stonegate's peak-time pricing in UK hospitality attracted criticism focused on profiteering and transparency (London, 2023; Witts, 2023).

Such reactions are not universal. Consumers routinely accept substantial intertemporal variation in airline fares, where advance-purchase discounts, fare classes and inventory management have established expectations that prices will vary (Kimes, 1994; Williams, 2022). Major League Baseball clubs similarly introduced algorithmic single-game pricing while protecting season-ticket holders and emphasising discounted availability (Drayer, Shapiro and Lee, 2012; Shapiro and Drayer, 2012). Behavioural responses therefore depend partly on whether price differences are expected and on whether consumers regard the pricing procedure, and the comparisons it invites, as legitimate.

2.3 Behavioural Regularities

Four empirical regularities motivate the model.

First, consumers evaluate prices relative to reference points rather than in absolute terms (Kahneman and Tversky, 1979; Kahneman, Knetsch and Thaler, 1986; Kőszegi and Rabin, 2006). Reference-price effects are well established in purchase data: Winer (1986) incorporates reference prices into an estimated brand-choice model, Kalyanaram and Winer (1995) document them as an empirical generalisation, and Mazumdar, Raj and Sinha (2005) review how such references are formed and updated.

Secondly, responses around these references are asymmetric: prices above an expectation reference generate greater disutility than equivalent reductions generate utility, consistent with

loss aversion (Tversky and Kahneman, 1991; Kőszegi and Rabin, 2006). This is the expectation-based mechanism through which anticipated prices affect optimal price adjustment.

Thirdly, consumers make interpersonal comparisons as well as comparisons with their own expectations. Inequity-aversion models formalise preferences over relative outcomes (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000), while evidence on price fairness shows that consumers evaluate transactions partly by comparing their treatment with that of others (Xia, Monroe and Cox, 2004). Anderson and Simester (2010) find in the field that customers who discover they paid more than a subsequently posted price reduce later purchasing. Richards, Liaukonytė and Streletskaya (2016) provide the most direct estimate of the channel: in an incentive-compatible experiment in which subjects observe the distribution of prices paid by others for an identical good, the disutility of paying more than others is economically substantial, and raising it sufficiently drives purchase probabilities to levels at which the pricing regime ceases to be viable. Their comparison is contemporaneous and anonymous, whereas the mechanism developed below is cross-cohort and temporal, the benchmark being a price a previous cohort is known to have paid; their estimates therefore speak to the strength of the comparison channel rather than to the intertemporal structure that generates the launch premium. A related literature examines the firm's alternative response of obscuring comparison altogether (Allender, Liaukonytė, Nasser and Richards, 2021), a margin held fixed here.

Fourthly, the strength of comparison depends on the comparability of transactions and on the identity of the comparator. Li, Hardesty and Craig (2018) show that bundling a dynamically priced product makes transactions appear less comparable, reducing consumers' propensity to compare their price with that paid by others and restoring perceived fairness towards levels observed under fixed pricing. Comparison intensity also rises with identification with the comparison group (Tajfel and Turner, 1979; Akerlof and Kranton, 2000; Chen and Li, 2009) and falls where the procedure generating different prices is regarded as legitimate (Xia, Monroe and Cox, 2004).

These regularities motivate the model's distinction between expectation-based and interpersonal reference dependence. The former governs how consumers respond to departures from anticipated prices; the latter makes the realised price paid by one cohort relevant to the demand of subsequent cohorts. The firm consequently chooses prices recognising that an early transaction shapes the behavioural environment in which its remaining capacity is sold.

3. The Model

This section develops the theoretical framework. We specify the economic environment and behavioural reference objects, formulate the firm's dynamic optimisation problem, and derive the equilibrium pricing strategy.

3.1 Economic Environment

The baseline model has two periods and sequential consumer cohorts. This is the minimal structure required to isolate the mechanism. The good is perishable and purchased once (a concert seat, flight or hotel night) so consumers do not accumulate their own transaction history for the particular good. Instead, a realised price paid by one cohort can become an interpersonal comparison benchmark for a later cohort. The first cohort has no predecessor

and therefore faces no such comparison penalty, while its price establishes the benchmark inherited by the second cohort. Section 4 extends the analysis to repricing epochs and shows that dividing the selling window more finely strengthens rather than eliminates the resulting distortion.

We consider a forward-looking monopolist selling a perishable product over two periods. In each period the firm chooses its price to maximise expected discounted profits. The timing is as follows: the firm chooses the first-period price; second-period consumers observe that price; the demand shock is realised and observed by the firm; and the firm then chooses the second-period price.

Capacity is fixed and unsold units have zero salvage value after period 2. Consumer valuations are independently distributed according to cumulative distribution , with strictly positive density on a connected support. Baseline demand is

$$D(P) = 1 - F(P),$$

and is assumed to be twice continuously differentiable with $D'(P) < 0$.

Demand uncertainty enters through the period-2 demand intercept,

$$a_2 = a_1 + \varepsilon,$$

where ε is a mean-zero shock with a continuous distribution G whose density is strictly positive on a compact support $[-d, d]$. The shock is therefore atomless and has full support on this interval. Its realisation is observed by the firm before the period-2 pricing decision.

Assumption 1 (Regularity)

The benchmark profit function

$$\Pi(P) = P\,D(P)$$

is strictly concave over the relevant range,

$$2D'(P) + PD''(P) < 0,$$

and the hazard rate

$$\frac{f(v)}{1 - F(v)}$$

is non-decreasing. Marginal cost is normalised to zero, so P is the margin.

Assumption 1 guarantees a unique interior optimum in the benchmark model and is satisfied by standard demand specifications including linear and exponential demand. The analytical

foundations are given in Appendix A.1. Where closed-form expressions provide additional insight, we specialise to linear demand,

$$D(P) = a - bP, \qquad a, b > 0,$$

while the general results continue to hold under Assumption 1. The linear specification is a reduced form used for closed-form illustration rather than a special case of $1 - F$.

Behavioural considerations do not alter the firm's objective directly. They affect future demand by changing the reference objects against which subsequent prices are evaluated. Section 3.2 introduces these reference objects.

3.2 Consumer Behaviour

Consumers evaluate period-2 prices against two distinct reference objects: an expectation reference, determined by the price they expected to pay, and an interpersonal comparison reference, given by the price paid by the earlier cohort. The former generates expectation-based loss aversion; the latter captures inequity aversion from paying more than an earlier buyer. Both affect demand and hence the firm's dynamic pricing problem.

3.2.1 Consumer Behaviour and Expectation Reference

A consumer with valuation v purchasing in period t at price P_t obtains utility

$$U_t = v - P_t - \delta_t(P_t;\cdot),$$

where $\delta_t(\cdot) \geq 0$ denotes behavioural disutility. The consumer purchases whenever

$$v \geq P_t + \delta_t(P_t;\cdot).$$

so effective demand becomes

$$\widetilde{D}_t(P_t;\cdot) = 1 - F\left(P_t + \delta_t(P_t;\cdot)\right).$$

A fresh unit mass of consumers arrives in each period. A proportion $\pi \in (0,1]$ is behaviourally sensitive, while the remainder is behaviourally neutral; the firm cannot distinguish between them and posts a common price. In period 1 there is no earlier cohort and, absent demand uncertainty, the expectation reference coincides with the chosen price. The behavioural adjustment therefore operates only in period 2.

Expectation reference. Second-period consumers evaluate the realised price against the price they expected to pay. Their forward-looking reference price is

$$R_2 = E[P_2 \mid I_1],$$

Where I_1 denotes the information available after observing the firm's first-period pricing decision. Consumers rationally anticipate the firm's pricing policy, so this reference must be consistent with the equilibrium price subsequently chosen. This fixed-point condition defines a personal equilibrium; Lemma A.3 in Appendix A.1 establishes existence, uniqueness and stability.

The reference is therefore not a free parameter. Because R_2 equals the average price the firm actually charges given that consumers expect R_2, consumers are disappointed in some demand states and pleasantly surprised in others, but they are not systematically wrong. The disutility below is genuine ex post; the expectation that generates it is correct ex ante.

Loss aversion around the reference. Let $\lambda \geq 1$ denote the loss-aversion coefficient. Behavioural disutility from paying more than the expectation reference is

$$\delta_2 = (\lambda - 1)\max\{0,\, P_2 - R_2\}.$$

so that a behaviourally sensitive consumer acts as though facing the perceived price

$$\hat{P}_2 = P_2 + (\lambda - 1)\max\{0,\ P_2 - R_2\} = \begin{cases} \lambda P_2 - (\lambda - 1)R_2, & P_2 \geq R_2, \\ P_2, & P_2 < R_2, \end{cases}$$

with

$$\frac{\partial \hat{P}_2}{\partial P_2} = \begin{cases} \lambda, & P_2 > R_2, \\ 1, & P_2 < R_2. \end{cases}$$

Three features of this specification matter for what follows.

First, the perceived price is a decision weight rather than a belief about money. The consumer knows what she is paying; $\hat{P}_2$ records what that payment feels like once the disappointment of paying more than expected is included. A consumer who expected to pay 100, is quoted 110, and has $\lambda = 2$ experiences the additional 10 twice over, once as expenditure and once as the shortfall against expectation, and so behaves as though the price were 120. A consumer quoted 90 in the same circumstances simply faces 90: gains relative to the reference carry no additional weight in this specification, which is the standard asymmetry documented by Tversky and Kahneman (1991).

Secondly, the wedge is a level effect that vanishes at the reference. A consumer who pays exactly what she expected suffers no disutility, so R_2 marks the point at which the penalty begins rather than imposing a uniform behavioural cost.

Thirdly, and most importantly for the firm's problem, λ is a slope. Above the reference each additional unit of price raises the perceived price by λ rather than by one, so demand is λ times more responsive to price increases above the reference than below it. This kink in the perceived price gives rise to an attenuated pass-through (shown in Corollary 1) and of price pooling regions (characterised in Proposition 1): the firm faces a demand curve that is markedly steeper once it prices above what consumers anticipated, and over a range of demand states it prefers to price exactly at the reference rather than cross the kink.

Given a behavioural share π, aggregate second-period demand under the expectation channel is

$$\widetilde{D}_2(P_2; R_2) = \pi D(\hat{P}_2) + (1-\pi)D(P_2).$$

Both consumer types pay the posted price P_2; the behavioural adjustment affects the purchase decision, not the transfer. Revenue is therefore $P_2\widetilde{D}_2(P_2; R_2)$ throughout.

Comparison reference. Second-period consumers may also compare their realised price with the observed price paid by the first-period cohort. Let $\alpha \geq 0$ denote the intensity of inequity aversion. The corresponding perceived price is

$$\tilde{P}_2 = P_2 + \alpha\max\{0,\ P_2 - P_1\} = \begin{cases}(1+\alpha)P_2 - \alpha P_1, & P_2 \geq P_1,\\ P_2, & P_2 < P_1,\end{cases}$$

so that

$$\frac{\partial \tilde{P}_2}{\partial P_2} = \begin{cases}1+\alpha, & P_2 > P_1,\\ 1, & P_2 < P_1.\end{cases}$$

As with the expectation channel, the wedge is one-sided and vanishes at the benchmark: a consumer who pays no more than the earlier cohort suffers no comparison disutility. And as before, α is a slope rather than a level, steepening perceived demand above P_1 without imposing a standing penalty below it.

The comparison reference differs fundamentally from the expectation reference. The latter is determined through personal equilibrium; the former is the realised price directly chosen by the firm in period 1. This distinction makes P_1 a strategic comparison benchmark for subsequent consumers.

Unified perceived price. The two mechanisms may operate simultaneously. A consumer can regard a realised price as high relative to her own expectation, relative to the price paid by the earlier cohort, relative to both, or relative to neither. Table 1 summarises the four cases.

Table 1. Behavioural responses to the two reference objects.

Realised price relative to reference objects	$\boldsymbol{(P_2 \leq P_1)}$	$\boldsymbol{(P_2 > P_1)}$
$\boldsymbol{(P_2 \leq R_2)}$	No behavioural penalty	Inequity aversion only
$\boldsymbol{(P_2 > R_2)}$	Loss aversion only	Loss aversion and inequity aversion

Note: The four cells describe the complete set of behavioural evaluations of a realised price. A behavioural mechanism is activated only when the realised price exceeds its associated reference object.

Both mechanisms operate only when the realised price exceeds the corresponding reference. The unified perceived price is therefore

$$\check{P}_2 \;=\; P_2 \;+\; (\lambda - 1)\cdot \max\{0, P_2 \;-\; R_2\} \;+\; \alpha \cdot \max\{0, P_2 \;-\; P_1\},$$

and aggregate second-period demand becomes

$$\widetilde{D}_2(P_2; R_2, P_1) = \pi D(\check{P}_2) + (1-\pi)D(P_2),$$

The expectation-only specification is obtained by setting α to zero. The perceived-price function is convex and kinked at R_2 and at P_1. Lemma A.1 establishes strict concavity of the firm's objective within each interior regime.

The ordering of the two reference objects is not imposed. Proposition 4 below shows that in equilibrium the expectation reference lies strictly below the realised first-period price, $R_2 < P_1$. The two kinks in the perceived price are therefore encountered in a determinate order as the period-2 price rises: loss aversion activates first, at R_2, and inequity aversion second, at P_1, so that the marginal perceived price takes the values

$$\frac{\partial \check{P}_2}{\partial P_2} = \begin{cases} 1, & P_2 < R_2, \\ \lambda, & R_2 < P_2 < P_1, \\ \lambda + \alpha, & P_2 > P_1. \end{cases}$$

This ordering generates a five-regime pricing rule shown in Proposition 1 below and two distinct price pooling regions (illustrated in Figure 1). It is an equilibrium property rather than an assumption, and it distinguishes the mechanism from models in which a single adaptive reference summarises consumers' price history.

The one-sided form of the comparison wedge is empirically motivated. Richards, Liaukonytė and Streletskaya (2016) test the symmetric formulation of Fehr and Schmidt (1999), in which utility falls whether the consumer pays more or less than others, against a one-sided alternative in which only disadvantageous comparison reduces utility. They reject symmetry: the disadvantageous coefficient is positive and significant across every specification, the advantageous coefficient is negative and significant, and the two differ by roughly an order of magnitude in their preferred estimates. The wedge above the comparison benchmark, with no corresponding term below it, is therefore the specification their evidence supports.

The same evidence indicates that the advantageous-comparison coefficient is not merely small but favourable, so setting it to zero understates period-2 demand below the benchmark. Section 3.6 shows that this omission is conservative with respect to the paper's results.

3.2.2 Behavioural Parameters

The primitive parameters from Section 3.2.1 combine into composite indices that govern the two channels. Define

$$\kappa \equiv \pi(\lambda - 1),$$

as the effective intensity of expectation-based reference dependence. It combines the share of behaviourally sensitive consumers with the strength of loss aversion among them, and it is the only expectation-side parameter that enters the firm's problem: λ and π affect pricing exclusively through this product. Similarly, define

$$\gamma \equiv \pi\alpha,$$

as the effective intensity of comparison-based inequity aversion, combining the same behavioural share with the intensity of the interpersonal comparison. Where either channel is inactive, $\pi = 0$, or $\lambda = 1$, or $\alpha = 0$, the corresponding index is zero and the associated wedge vanishes.

The comparison channel, unlike the expectation channel, depends on features of the market as well as on preferences. As described in the Introduction, two such features modify it in opposite directions; this section gives their formal representation.

The first is identity salience. Let $\sigma \geq 0$ denote the prevalence of socially salient comparison opportunities in the market and $\varphi \geq 0$ the incremental weight consumers attach to comparisons within their own group. These enter as a product because both are required: comparison must be both available and felt to matter. Where either is absent the amplification factor $(1 + \sigma\varphi)$ equals one, and because the product is bounded below by one, identity can only strengthen the comparison channel and never weaken it.

The second is procedural legitimacy. Let $J \in [0,1]$ denote the degree of legitimacy. It enters as $(1 - J)$ because legitimacy discounts the behavioural penalty attaching to a given price difference without altering the difference itself: at $J = 1$ the comparison is fully excused and the wedge disappears, while at $J = 0$ it operates at full force.

Combining these, define

$$\Gamma \equiv \gamma(1 + \sigma\varphi)(1 - J)$$

as the effective comparison intensity. Neither identity nor legitimacy constitutes a further behavioural mechanism: the multiplicative form records that both scale a wedge already established rather than creating one. Where both are inactive, Γ collapses to γ and every result below is unchanged. The empirical motivation for all three parameters is set out in Section 2.3, and their consequences are characterised in Section 3.7; the two mechanisms of Sections 3.4 to 3.6 are logically prior and do not depend on them.

3.2.3 Summary of notation

Table 2 introduces the primitive behavioural parameters and Table 3 summarises the composite indices, reference objects and benchmark prices used throughout the model.

Table 2. Primitive behavioural parameters.

Symbol	Interpretation
λ	Loss-aversion coefficient
α	Inequity-aversion coefficient
π	Share of behaviourally sensitive consumers
β	Firm's discount factor
J	Procedural legitimacy parameter
σ	Prevalence (or intensity) of socially salient comparison opportunities
φ	Incremental strength of ingroup identity in fairness comparisons

Table 3. Behavioural Indices, Reference Objects and Benchmark Prices.

Symbol	Interpretation
Composite behavioural indices	
$\kappa \equiv \pi(\lambda - 1)$	Effective expectation-reference intensity
$\gamma \equiv \pi\alpha$	Effective comparison intensity
$\Gamma \equiv \gamma(1 + \sigma\varphi)(1 - J) = \pi\alpha(1 + \sigma\varphi)(1 - J).$	Effective comparison intensity incorporating social identity and procedural legitimacy
Reference objects	
R_2	Expectation-based reference price (personal-equilibrium benchmark)
P_1	First-period comparison reference
$\hat{P}_2$	Perceived period-2 price relative to expectations
$\tilde{P}_2$	Perceived period-2 price relative to the comparison reference
$\breve{P}_2$	Unified behavioural perceived price
Benchmark prices	
P_t^B	Static monopoly benchmark price
P_t^0	Benchmark dynamic optimum without behavioural effects

Note: R_2 and P_1 denote the expectation and comparison reference objects entering the behavioural demand system. A superscript * denotes the corresponding equilibrium value. Throughout the theoretical

development the model is first presented in terms of the general reference objects; equilibrium values are introduced subsequently when the firm's optimisation problem is solved.

The two reference objects are treated asymmetrically. The expectation reference is a fully rational object: consumers anticipate the distribution of second-period prices correctly, and R_2 is pinned down by the personal-equilibrium condition $R_2 = E[P_2]$, so consumers are not deceived about what they will be asked to pay. The comparison reference is treated differently. Consumers respond to the earlier cohort's price as a benchmark for evaluating their own treatment, without inferring that the firm may have set it strategically with precisely that evaluation in mind. It is worth being explicit that the launch premium depends on this asymmetry, and on why it is the appropriate modelling choice rather than a convenient one.

The justification is that the two references answer different questions. The expectation reference asks what the consumer will pay, and a consumer who is systematically wrong about that is being modelled as making a forecasting error; a strong assumption, and one the personal-equilibrium condition avoids. The comparison reference asks how the consumer's treatment compares with that of another buyer, and the answer does not depend on why the earlier price was what it was. Disadvantageous inequity in the sense of Fehr and Schmidt (1999) is defined over realised payoffs: a consumer who pays more than an earlier buyer for an equivalent unit has in fact been treated less favourably, and that remains true whether the earlier price was set myopically or strategically. Knowing that the firm elevated the launch price does not restore the consumer's position relative to the earlier cohort; if anything it identifies the firm as responsible for a difference the consumer might otherwise have attributed to circumstance.

This is not the only possible position. An alternative tradition, following Rabin (1993), makes inferred intentions central to fairness judgements, and under that view a consumer who understood the launch price to be strategically elevated might respond more adversely rather than less. The framework here accommodates that possibility through the procedural-legitimacy parameter. Endogenising legitimacy would be a natural extension: it would introduce a countervailing force, since a higher launch premium would reduce the legitimacy discount and hence the benefit of elevating the benchmark, and would plausibly bound the launch premium.

3.3 Firm Optimisation

The firm is fully rational and forward-looking. Taking consumer behaviour as given, it chooses prices to maximise expected discounted profits, recognising that P_1 affects continuation value through the behavioural reference objects entering second-period demand.

3.3.1 Dynamic Optimisation

First-period profit is standard, while second-period demand incorporates the behavioural adjustments developed in Section 3.2:

$$\Pi_1 = P_1 D(P_1),$$

$$\Pi_2 = P_2 \widetilde{D}_2(P_2; R_2, P_1).$$

The firm's dynamic optimisation problem is

$$V_1 = \max_{P_1}\{P_1 D(P_1) + \beta \mathbb{E}_\varepsilon[V_2(P_1, \varepsilon)]\},$$

$$V_2(P_1, \varepsilon) = \max_{P_2} P_2 \widetilde{D}_2(P_2; R_2, P_1, \varepsilon).$$

Although first-period demand is behaviourally unadjusted, P_1 is not a static choice: it becomes the interpersonal comparison benchmark for the second cohort and therefore enters continuation value. The expectation reference R_2, by contrast, is determined through personal equilibrium. This asymmetry between the two reference objects generates the forward-looking pricing incentives characterised below.

3.3.2 Unified Period-2 Problem

Aggregating over the behavioural share π and behaviourally neutral consumers, second-period demand is $D_2(P_2) = (1-\pi)D(P_2) + \pi D\big(\check{P}_2\big)$, and the firm chooses P_2 to maximise $P_2 D_2(P_2)$. Two objects organise the solution and recur throughout the remainder of the paper.

Definition 1 (Dispersion region). At shock ε the firm is in the *dispersion region* if $P_2^*(\varepsilon) > P_1$, equivalently if the cross-cohort price difference $\Delta^*(\varepsilon) \equiv P_2^*(\varepsilon) - P_1$, is strictly positive (arguments in Δ are suppressed except for the parameter under consideration),

$$(1-\pi)[D(P_1) + P_1 D'(P_1)] + \pi\big[D\big(\hat{P}_1\big) + (\lambda+\alpha)P_1 D'\big(\hat{P}_1\big)\big] > 0, \qquad \hat{P}_1 \equiv \lambda P_1 - (\lambda-1)R_2,$$

which under linear demand reduces to $a_2 > \bar{A} + b\gamma P_1$, where $\bar{A} \equiv 2b(1+\kappa)P_1 - b\kappa R_2$. In the comparison-only case $\kappa = 0$ the condition simplifies to $D(P_1) + (1+\gamma)P_1 D'(P_1) > 0$, or $a_2 > bP_1(2+\gamma)$ under linear demand.

Write $\mathbf{1}_{disp}$ for the indicator of the dispersion region, equal to one at those shocks ε for which $P_2^*(\varepsilon) > P_1$ and zero otherwise, and $\mathbf{1}_{pool}$ for the indicator of the pooling region, equal to one at those ε for which $P_2^*(\varepsilon) = P_1$ and zero otherwise. Lemma 1 shows that the pooling region is an interval and measures its width. The two indicators are mutually exclusive, and their complement is the set of shocks at which the firm prices strictly below the comparison reference. Also both indicators depend on P_1 as well as on ε; arguments are suppressed where no ambiguity arises. Note that these indicators are functions of the demand shock, recording which region the period-2 optimum falls into at a given first-period price. They are distinct from indicators such as $\mathbf{1}\{P > R_2\}$ used in the proofs, which are functions of a candidate price and record which behavioural terms are active at it.

Definition 1 is strictly stronger than the condition that the non-behavioural benchmark exceeds the comparison reference. The benchmark condition is $D(P_1) + P_1 D'(P_1) > 0$, or $a_2 > 2bP_1$; between the two lies a set of shocks on which the firm prices exactly at P_1. Lemma 1 shows that this set is an interval and measures it.

Lemma 1 (Kink decoupling). Under Assumption 1 and linear demand, marginal profit is continuous in P_2 except at $P_2 = R_2$ and $P_2 = P_1$, where it falls by $b\kappa R_2$ and $b\gamma P_1$ respectively.

Neither jump depends on the position of the other reference, nor on which of the two is the lower. Consequently:

(i) the period-2 objective is globally strictly concave and the optimum is unique;

(ii) the set of demand intercepts a_2 at which the firm prices exactly at reference $r \in \{R_2, P_1\}$ is an interval of width $b\, w_r\, r$, where $w_{R_2} = \kappa$ and $w_{P_1} = \gamma$.

Proof. See Appendix A.2.1.

The width of each pooling interval therefore equals the size of the corresponding downward jump in marginal profit, and the measure of states over which the firm fully absorbs a demand shock is proportional to the strength of the behavioural motive absorbing it. Because the jumps are independent of the ordering of the two references, the pricing rule below is not conditioned on that ordering; Proposition 4 shows which ordering arises in equilibrium.

Proposition 1 (Unified Period-2 Pricing). Under Assumption 1 the period-2 objective is continuous and globally strictly concave on $\mathbb{R}_+$, hence admits a unique optimum $P_2^*(\varepsilon)$. Under linear demand $D(P) = a_2 - bP$ with $a_2 \equiv a_1 + \varepsilon$, and writing $\kappa \equiv \pi(\lambda - 1)$, $\gamma \equiv \pi\alpha$ and $\bar{A} \equiv 2b(1+\kappa)P_1 - b\kappa R_2$, the optimum is piecewise linear in a_2 with five regimes shown in Table 4. In the terminal regime, where both references bind,

$$P_2^*(\varepsilon) = \frac{a_2 + b(\kappa R_2 + \gamma P_1)}{2b(1+\kappa+\gamma)}.$$

Relative to the non-behavioural benchmark $P_2^0 = a_2/(2b)$:

(i) pass-through is attenuated by the factor $1/(1+\kappa+\gamma)$ in the terminal regime and by $1/(1+\kappa)$ in the expectation-loss regime, and is zero on the two pooling intervals (with κ and γ exchanged when $P_1 < R_2$);

(ii) dispersion $\Delta^* \equiv P_2^* - P_1$ is strictly decreasing in γ throughout the terminal regime, with

$$\frac{\partial \Delta^*}{\partial \gamma} = \frac{bP_1(1+\kappa) - a_2 - b\kappa R_2}{2b(1+\kappa+\gamma)^2} < 0; \qquad (1)$$

(iii) complete compression, $\Delta^* = 0$, obtains for every $\gamma \geq \gamma^*(\kappa)$, where

$$\gamma^*(\kappa) \equiv \max\left\{0,\ \frac{a_2 + b\kappa R_2 - 2bP_1(1+\kappa)}{bP_1}\right\},$$

and equivalently, for fixed γ, over an interval of demand shocks of width $b\gamma P_1$ in a_2. The corresponding plateau at R_2 has width $b\kappa R_2$.

The taxonomy does not depend on the ordering of the two references: for $P_1 < R_2$ the five regimes are obtained by exchanging (R_2, κ) with (P_1, γ) throughout, the jump sizes and interval widths being unchanged by Lemma 1. The expectation-only and comparison-only cases obtain by setting $\gamma = 0$ and $\kappa = 0$ respectively.

Proof. See Appendix A.2.2.

Table 4. The five period-2 pricing regimes, case $R_2 < P_1$.

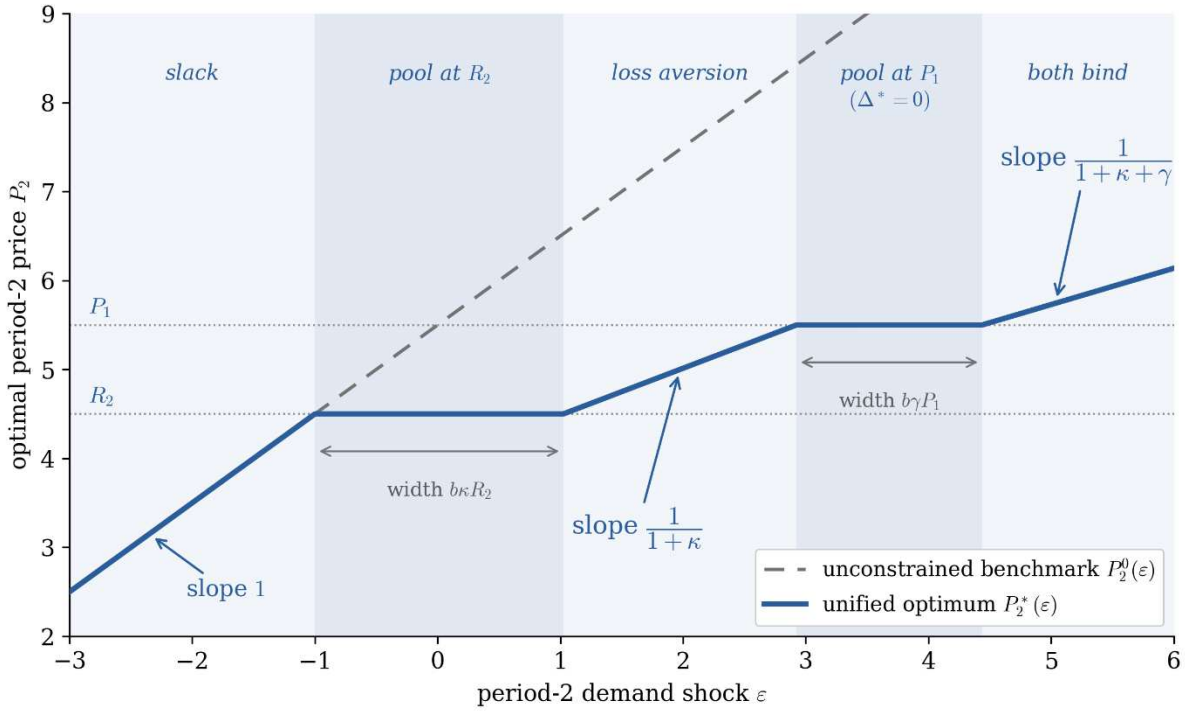
Regime	Condition on $a_2 = a_1 + \varepsilon$	P_2^*
Slack	$a_2 \le 2bR_2$	$\frac{a_2}{2b}$
Pool at R_2	$2bR_2 < a_2 \le (2+\kappa)bR_2$	R_2
Expectation loss only	$(2+\kappa)bR_2 < a_2 \le \bar{A}$	$\frac{a_2 + b\kappa R_2}{2b(1+\kappa)}$
Pool at P_1	$\bar{A} < a_2 \le \bar{A} + b\gamma P_1$	P_1
Both bind	$a_2 > \bar{A} + b\gamma P_1$	$\frac{a_2 + b(\kappa R_2 + \gamma P_1)}{2b(1+\kappa+\gamma)}$

Note: $\bar{A} \equiv 2b(1+\kappa)P_1 - b\kappa R_2$. The pooling intervals have widths $b\kappa R_2$ and $b\gamma P_1$, equal to the downward jumps in marginal profit at the two kinks. For $P_1 < R_2$, exchange (R_2, κ) with (P_1, γ).

Proposition 1 produces five pricing regimes. When neither reference binds, the firm follows the conventional dynamic-pricing benchmark. It then pools at R_2 over an interval of shocks of width $b\kappa R_2$; above that interval, expectation-based loss aversion attenuates price adjustment by the factor $1/(1+\kappa)$; at P_1 interpersonal comparison generates a second pooling interval, of width $b\gamma P_1$; and above it both mechanisms bind, attenuating adjustment by $1/(1+\kappa+\gamma)$.

Because both behavioural penalties operate only when price exceeds the corresponding reference, adjustment is asymmetric: favourable demand shocks are attenuated once the relevant reference binds, while equivalent downward movements incur no corresponding behavioural penalty. Figure 1 illustrates the resulting pricing rule.

Figure 1. Unified period-2 pricing rule under expectation- and comparison-based reference formation.



Note: The figure illustrates the five regimes of Proposition 1 and the two pooling widths of Lemma 1. It is drawn for $R_2 < P_1$, an ordering shown in Proposition 4 to arise endogenously in equilibrium, and normalised so that $2b = 1$, so that the non-behavioural benchmark has unit slope and the annotated segment slopes are $1/(1+\kappa)$ and $1/(1+\kappa+\gamma)$ exactly.

3.4 Expectation-Based Reference Dependence

We first consider the behavioural implications of expectation-based reference dependence. Consumers compare the realised second-period price with the price they expected to pay, so prices above the expectation benchmark generate an additional behavioural cost for behaviourally sensitive consumers. Consequently, firms become less willing to translate temporary demand shocks into immediate price increases. Expectation-based reference dependence therefore moderates the responsiveness of optimal prices to changing market conditions. Corollary 1 formalises this result.

Corollary 1 (Reference-Dependent Price Smoothing). Under Assumption 1, let R_2^* denote the personal-equilibrium reference price and consider a positive demand shock ε in period 2. Let $P_2^*(\varepsilon)$ denote the optimal period-2 price under the expectation-augmented demand $\widetilde{D}_2$, and let $P_2^0(\varepsilon)$ denote the corresponding optimum under the non-behavioural benchmark. Then:

(a) $P_2^*(\varepsilon) \leq P_2^0(\varepsilon)$, with strict inequality whenever $P_2^0(\varepsilon) > R_2^*$.

(b) On the interior expectation-loss regime,

$$\frac{\partial P_2^*}{\partial \varepsilon} = \frac{1}{1+\kappa} \cdot \frac{\partial P_2^0}{\partial \varepsilon} = \frac{1}{2b(1+\kappa)},$$

so the response to demand shocks is attenuated by the factor $1/(1+\kappa) = 1/[1+\pi(\lambda-1)]$ relative to the non-behavioural benchmark, this factor lying strictly between zero and one whenever $\pi > 0$ and $\lambda > 1$. On the pooling interval at R_2^*, which by Lemma 1 has width $b\kappa R_2^*$ in a_2, the response is zero.

(c) $\mathrm{Var}_\varepsilon(P_2^*) < \mathrm{Var}_\varepsilon(P_2^0)$ whenever the demand-shock distribution assigns positive probability to outcomes for which the benchmark optimum lies in the behavioural loss region.

Proof. See Appendix A.2.3.

Expectation-based reference dependence therefore acts as a behavioural smoothing mechanism, attenuating firms' responses to temporary demand shocks while preserving the underlying optimisation problem. Consumers evaluate realised prices relative to the prices they expected to pay, so firms become less willing to translate temporary demand shocks into immediate price increases. Behavioural considerations therefore attenuate, rather than eliminate, firms' responses to changing market conditions. The strength of this effect depends on the composite parameter κ, which combines the prevalence of behaviourally sensitive consumers with the degree of loss aversion. The next section considers a distinct behavioural mechanism arising from consumers' comparisons with prices paid by earlier purchasers.

3.5 Comparison-Based Inequity Aversion

Section 3.4 established how expectation-based reference dependence moderates firms' responses to demand shocks. Consumers, however, frequently evaluate prices against a second behavioural reference object: the prices paid by other consumers purchasing the same product under comparable circumstances. Unlike expectation-based reference dependence, which concerns whether consumers pay more than they expected, comparison-based inequity aversion concerns whether they pay more than others. The relevant comparison benchmark in the present model is the realised first-period price observed by the second-period cohort.

This distinction is economically important. Expectation-based reference dependence influences the responsiveness of prices to changing market conditions. Comparison-based inequity aversion instead influences the firm's willingness to create price differences between successive consumer cohorts. The two behavioural mechanisms therefore affect dynamic pricing through different channels. Corollary 2 establishes the implications of the comparison benchmark.

Given the behavioural specification introduced in Section 3.3, behaviourally sensitive consumers evaluate the realised second-period price relative to the observed first-period price. The strength of this comparison is governed by the inequity-aversion parameter α, with the effective behavioural intensity summarised by the composite parameter γ. As comparison-based inequity becomes stronger, firms face progressively greater incentives to avoid large differences between the prices charged to successive consumer cohorts.

Corollary 2 (Comparison-Based Compression). Under Assumption 1, suppose a share $\pi \in (0,1)$ of consumers are inequity-averse with intensity $\alpha > 0$, so that period-2 behaviourally sensitive demand is evaluated at the comparison-adjusted perceived price $\tilde{P}_2 = (1+\alpha)P_2 - \alpha P_1$, where P_1 is the observed first-period comparison price. Hold P_1 fixed and suppose the firm is in the dispersion region of Definition 1 with $\kappa = 0$, that is

$$D(P_1) + (1+\gamma)P_1 D'(P_1) > 0, \qquad \text{under linear demand } a_2 > bP_1(2+\gamma).$$

Then the optimal upward adjustment $\Delta^*(\alpha,\pi) \equiv P_2^*(\alpha,\pi) - P_1$ satisfies:

(a) $0 < \Delta^*(\alpha,\pi) < P_2^0 - P_1$: inequity aversion strictly compresses dispersion relative to the no-inequity benchmark;

(b) $\partial\Delta^*/\partial\alpha < 0$ and $\partial\Delta^*/\partial\pi < 0$, so Δ^* is strictly decreasing in the composite comparison intensity $\gamma = \pi\alpha$.

Under linear demand $D(P) = a_2 - bP$,

$$P_2^*(\gamma) = \frac{a_2 + \gamma b P_1}{2b(1+\gamma)}, \qquad \Delta^*(\gamma) = \frac{a_2 - bP_1(2+\gamma)}{2b(1+\gamma)}, \qquad \frac{\partial\Delta^*}{\partial\gamma} = \frac{bP_1 - a_2}{2b(1+\gamma)^2} < 0.$$

Proof. See Appendix A.2.4.

Remark 1. The dispersion condition is strictly stronger than $P_2^0 > P_1$. On the intermediate set $2bP_1 < a_2 \le bP_1(2+\gamma)$ the comparison kink binds, the firm prices exactly at P_1, and the derivatives in part (b) are zero. That case is the subject of Corollary 3 below.

Comparison-based inequity aversion instead operates through a different channel, reducing the profitability of price dispersion across successive consumer cohorts. Whereas reference dependence moderates firms' responses to demand shocks, comparison-based inequity aversion reduces the profitability of charging different prices across successive consumer cohorts. The behavioural effect therefore operates on equilibrium price dispersion rather than on the responsiveness of prices to changing market conditions.

3.6 Complete Compression

Corollary 2 established that comparison-based inequity aversion progressively compresses price differences between successive consumer cohorts. An important question is whether this behavioural effect remains incremental or whether sufficiently strong inequity concerns can eliminate such differences altogether. The next result identifies the critical behavioural threshold beyond which further price differentiation ceases to be optimal.

Corollary 3 (Complete Compression). Under Assumption 1 with $\kappa = 0$ and $\gamma \equiv \pi\alpha > 0$, define

$$\gamma^*(0) \equiv \max\left\{0, \frac{a_2 - 2bP_1}{bP_1}\right\}.$$

Then:

(a) for $\gamma < \gamma^*(0)$ the firm is in the dispersion region, $\Delta^*(\gamma) = [a_2 - bP_1(2+\gamma)]/[2b(1+\gamma)] > 0$, and Δ^* is strictly decreasing in γ and hence in each of α and π;

(b) for every $\gamma \geq \gamma^*(0)$, $\Delta^*(\gamma) = 0$: the firm prices exactly at the comparison reference P_1, and complete compression once attained is not reversed by further increases in comparison intensity;

(c) equivalently, for fixed γ the set of demand shocks at which $\Delta^* = 0$ is an interval of width $b\gamma P_1$ in a_2, coinciding with the pooling region of Proposition 1(iii) and Lemma 1(ii).

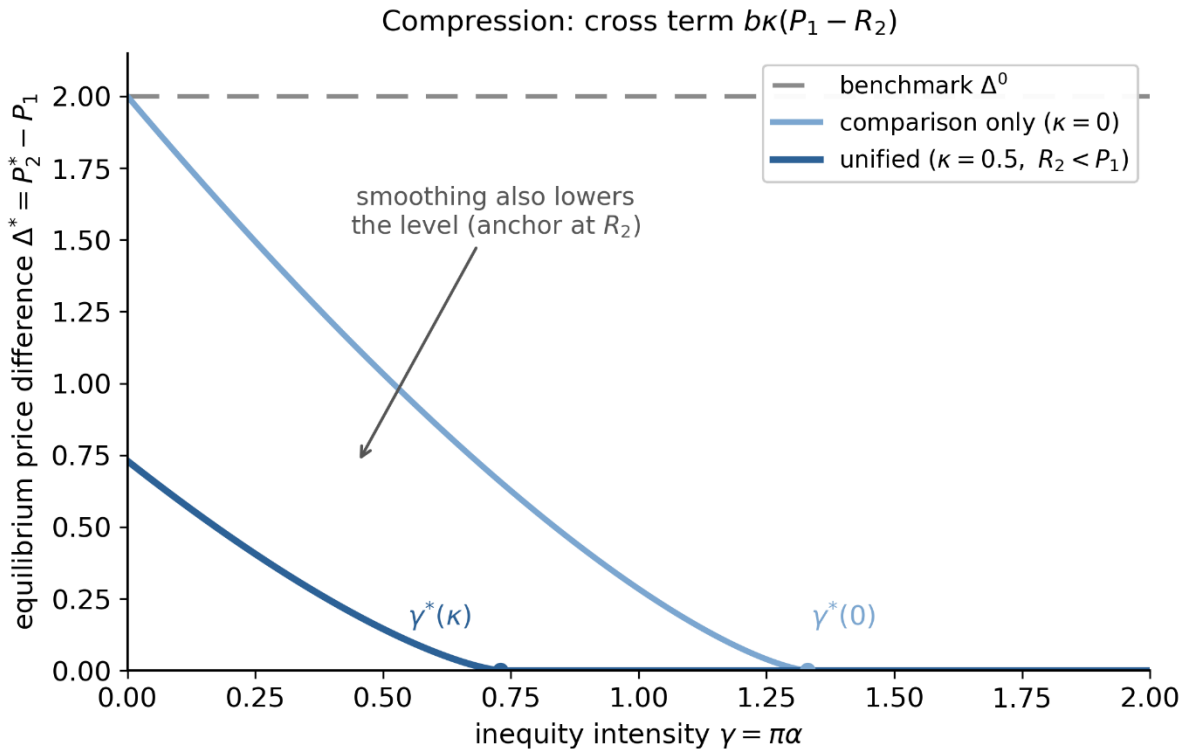
Proof. See Appendix A.2.4.

Corollary 3 identifies the limiting case of comparison-based inequity aversion. As behavioural responses strengthen, the gains from intertemporal price differentiation are progressively outweighed by the behavioural costs of charging different prices. Beyond a critical threshold, the optimal policy converges to complete compression, with both consumer cohorts facing the same price.

The specification adopted here omits a behavioural gain below the comparison benchmark, and Corollary 3 is the natural place to record why that omission is conservative. In the estimates of Richards, Liaukonytė and Streletskaya (2016) the advantageous-comparison coefficient is not merely small but favourable: consumers who learn they paid less than others become more likely to purchase, not less. Setting that term to zero therefore understates period-2 demand at prices below P_1. Restoring it would raise demand precisely in the region where complete compression occurs, increasing the profitability of pricing at the benchmark and lowering the threshold $\gamma^*(0)$ at which the dispersion region is exhausted. Complete compression would arise for a wider set of parameters, not a narrower one. The same logic strengthens the launch premium of Theorem 1, since the firm pools at the benchmark in exactly the states where a gain term would apply, so an elevated benchmark would be worth more rather than less. The cost of incorporating it is a demand system with a third kink and no closed-form solution; the analytical results below would survive in sign but not in the explicit form of Corollary 3 and Proposition 1.

Figure 2 illustrates how increasing comparison intensity progressively compresses equilibrium price dispersion. Expectation-based reference dependence reinforces this effect by shifting the entire compression schedule downward and bringing forward the point at which complete compression occurs.

Figure 2. Equilibrium Price Dispersion under Increasing Comparison Intensity



Note. The figure plots optimal period-2 dispersion $\Delta^* = P_2^* - P_1$ against comparison intensity $\gamma = \pi\alpha$, holding P_1 fixed. Dispersion falls continuously until the threshold $\gamma^*(\kappa)$ and is identically zero thereafter, as Corollary 3(b) states. The expectation channel shifts the whole schedule downward through the cross term $b\kappa(P_1 - R_2)$ and brings the threshold forward from $\gamma^*(0)$ to $\gamma^*(\kappa)$. The vertical intercept of the unified schedule is the expectation-only dispersion, not the unconstrained benchmark Δ^0; the two coincide only when $\kappa = 0$.

3.7 Identity and Procedural Legitimacy

The preceding analysis has treated comparison-based inequity aversion as homogeneous across consumers. Behavioural evidence, however, suggests that fairness judgements depend systematically upon social context. Social Identity Theory predicts that individuals attach greater weight to comparisons involving members of salient social groups (Tajfel and Turner, 1979). Building on this insight, Akerlof and Kranton (2000) incorporate identity directly into economic preferences, while Chen and Li (2009) show experimentally that social identity systematically affects economic behaviour. Consistent with this evidence, perceptions of unfair treatment become stronger when consumers compare themselves with members of their own reference group and weaker when observed price differences arise from legitimate procedures. These factors do not create additional behavioural mechanisms. Rather, they modify the strength of the comparison-based channel developed in the previous section.

We therefore consider two extensions that modify the effective intensity of comparison-based inequity aversion while leaving the underlying optimisation problem unchanged. The first incorporates social identity, allowing comparison-based responses to become stronger when consumers identify with those against whom they compare themselves. The second incorporates procedural legitimacy, recognising that transparent, predictable and well-understood pricing rules can attenuate perceived unfairness.

The legitimacy parameter has a direct experimental counterpart. Richards, Liaukonytė and Streletskaya (2016) compare a posted-price regime, in which the buyer accepts or rejects a price set by the seller, with a price-discovery regime, in which the buyer opens a negotiation and thereby participates in forming the price paid. The manipulation alters the procedure generating price differences while leaving the differences themselves unchanged, which is precisely the variation J is intended to capture. Its effect is large. In their preferred specification the disadvantageous-comparison effect is strongly positive under posted prices, but the interaction with the price-discovery treatment more than offsets it, so that the net effect among participating buyers changes sign: subjects who regarded the price as unfair relative to others were nonetheless more likely to purchase than comparable subjects facing posted prices. Procedural participation does not merely attenuate the comparison response in their data; it reverses it. This supports values of J at the upper end of the unit interval where consumers genuinely participate in price formation, and suggests that the attenuation in Proposition 3 may be conservative for procedures involving active participation rather than disclosure alone.

Every result stated above for γ holds with Γ in place of γ, since identity and legitimacy enter the period-2 problem only through the effective intensity of the comparison term. In particular $\Gamma^*(\kappa)$ denotes the complete-compression threshold of Proposition 1(iii) with Γ substituted for γ.

Proposition 2 (Identity Amplification). Under Assumption 1, with Γ as defined above, the optimal dispersion $\Delta^*(\varphi, \sigma)$ satisfies:

(a) $\partial\Delta^* / \partial\varphi \leq 0$, with strict inequality whenever $\alpha > 0, \pi > 0, \sigma > 0, J < 1$ and $\Gamma < \Gamma^*(\kappa)$;

(b) $\partial\Delta^* / \partial\sigma \leq 0$, with strict inequality whenever $\alpha > 0, \pi > 0, \varphi > 0, J < 1$ and $\Gamma < \Gamma^*(\kappa)$;

(c) for two sectors A and B with common π, α, J and parameters (φ_A, σ_A) and (φ_B, σ_B), $\Delta_A^* < \Delta_B^*$ whenever $\sigma_A\varphi_A > \sigma_B\varphi_B$ and $\Gamma_A < \Gamma^*(\kappa)$. If $\Gamma_B \geq \Gamma^*(\kappa)$ both sectors exhibit complete compression and the comparison is degenerate.

Identity therefore increases the effective comparison intensity Γ, strictly compressing dispersion until complete compression is reached and having no further effect thereafter.

Proof. See Appendix A.2.5.

Proposition 2 shows that social identity amplifies comparison-based behavioural responses without altering the structure of the equilibrium. Greater identification with earlier consumers increases the behavioural importance of price comparisons and therefore strengthens the firm's incentive to compress intertemporal price differences.

Proposition 3 (Procedural Legitimacy). Under Assumption 1, procedural legitimacy attenuates the effective comparison intensity according to $\Gamma = \gamma(1 + \sigma\varphi)(1 - J)$, so that $\partial\Gamma / \partial J = -\gamma(1 + \sigma\varphi) < 0$ for $J < 1$. Then:

(a) for values of J at which $\Gamma < \Gamma^*(\kappa)$, the optimal price difference $\Delta^*(\Gamma)$ is strictly increasing in J;

(b) for values at which $\Gamma \geq \Gamma^*(\kappa)$, $\Delta^* = 0$ and $\partial\Delta^* / \partial J = 0$;

(c) as $J \to 1$, $\Gamma \to 0$ and $\Delta^*(\Gamma)$ converges to the expectation-only dispersion

$$\Delta^{\kappa} \equiv \frac{a_2 + b\kappa R_2}{2b(1+\kappa)} - P_1,$$

which coincides with the unconstrained benchmark $\Delta^0 = a_2/(2b) - P_1$ if and only if $\kappa = 0$.

Proof. See Appendix A.2.6.

Remark 2. Complete procedural legitimacy therefore removes the comparison distortion but not the expectation distortion. The residual gap is

$$\Delta^0 - \Delta^{\kappa} = \frac{\kappa(a_2 - bR_2)}{2b(1+\kappa)} > 0 \qquad \text{whenever } \kappa > 0,$$

which is strictly positive throughout the region in which the expectation reference binds, since there $a_2 > 2bR_2 > bR_2$. It is the vertical intercept gap between the unified and comparison-only schedules in Figure 2. It should not be confused with the cross term $b\kappa(P_1 - R_2)$ appearing in the numerator of (1) in Proposition 1, which governs the *slope* of the compression schedule rather than its intercept.

Remark 3 (calibration). The numerical work of Section 4 and Appendix A.4 sets $J = 0$, so that $\Gamma = \gamma(1 + \sigma\varphi)$ throughout. The reported distortions are therefore the largest attainable along the legitimacy margin, and Proposition 3 implies that any positive J would reduce them.

Proposition 3 identifies procedural legitimacy as the natural counterweight to comparison-based inequity aversion. When consumers perceive pricing procedures to be transparent and legitimate, observed price differences carry less behavioural weight. The resulting reduction in comparison intensity allows firms greater pricing flexibility without altering the underlying optimisation problem.

3.8 Equilibrium Pricing

The preceding propositions characterise the behavioural mechanisms individually. We now combine them to derive the firm's equilibrium dynamic pricing strategy. The central distinction is that the two reference objects play different roles in the optimisation problem: the expectation reference is an equilibrium object the firm takes as given in the fixed-point sense, whereas the comparison reference is the firm's own earlier price. The second is therefore an instrument as well as a constraint, and the resulting forward-looking incentive raises rather than lowers the initial price; the opposite of the distortion arising when a single consumer's own price history supplies the reference.

3.8.1 Behavioural Reference Formation and Dynamic Optimisation

Although both behavioural reference objects influence future demand, they differ in one respect. The expectation-based reference price R_2 is determined through consumers' rational expectations and therefore satisfies the personal-equilibrium condition. By contrast, the

comparison benchmark is the realised first-period price, which is directly observed by later consumers and cannot subsequently be revised. The firm therefore faces two distinct behavioural reference state variables. The expectation reference cannot be manipulated directly because it must satisfy personal equilibrium. By contrast, the comparison reference is determined by the firm's own first-period pricing decision and therefore becomes a strategic choice variable. This asymmetry is central to the equilibrium characterised below.

Theorem 1 shows that this asymmetry gives rise to a forward-looking pricing incentive: firms optimally elevate initial prices because current prices become future comparison benchmarks.

Theorem 1 (Forward-Looking Behavioural Pricing). Under Assumption 1, let P_1^B denote the static monopoly benchmark and define the saturation threshold

$$\bar{\gamma} \equiv \max\left\{0, \frac{a_1 + d - b(2+\kappa)P_1^B}{bP_1^B}\right\},$$

above which the dispersion region is empty at every $P_1 \geq P_1^B$. Then:

(i) whenever $\beta > 0$ and $0 < \gamma < \bar{\gamma}$, every period-1 optimum satisfies $P_1^* > P_1^B$; under linear demand

$$P_1^* = P_1^B \; + \; (\beta/2)\mathbb{E}_\varepsilon\left[\gamma P_2^*(\varepsilon)\mathbf{1}_{disp}\right] \; + \; \left(\frac{\beta}{2b}\right)\mathbb{E}\left[(a_2 - 2bP_1)\mathbf{1}_{pool}\right].$$

The two terms are the two channels through which the launch price relaxes the period-2 behavioural constraint. The second is strictly positive throughout the pooling region and is what sustains the premium in the saturation region of part (iii), where the first term vanishes;

(ii) the wedge $P_1^* - P_1^B$ is non-decreasing in the discount factor β for every $\gamma > 0$, and strictly increasing at any β at which the period-1 optimum is unique, reflecting the greater value of relaxing future behavioural constraints;

(iii) the wedge is non-decreasing in γ and bounded: it equals zero at $\gamma = 0$, is strictly positive for every $\gamma > 0$, and is constant for every $\gamma \geq \hat{\gamma}$, where $\hat{\gamma}$ is the fixed point at which the comparison intensity exactly exhausts the dispersion region at the equilibrium launch price, $\hat{\gamma} = \bar{\gamma}\left(P_1^*(\hat{\gamma})\right)$. Raising γ increases the marginal value of an elevated comparison reference but simultaneously compresses P_2^* and shrinks the dispersion region; once the latter is exhausted the comparison term is never active in equilibrium and further intensification leaves the firm's problem unchanged. The premium does not revert: beyond $\hat{\gamma}$ the firm prices at the comparison reference in every state where it would otherwise price above it, and an elevated reference continues to raise continuation value through the price charged when pooling. The increase on $(0, \hat{\gamma})$ is strictly throughout the parameter range shown in Section 4, though the offsetting effects of γ are not signed analytically;

(iv) through Γ, the wedge inherits (i) to (iii) with Γ in place of γ, so it is non-decreasing in identity salience $\sigma\varphi$ and in the complement of procedural legitimacy $1 - J$, and constant once the composite intensity exceeds the saturation threshold.

Proof. See Appendix A.2.7.

Remark 4 (period-1 concavity). $V_2(P_1, \varepsilon)$ is a maximum of piecewise-quadratic functions of P_1 and is not established to be concave: raising P_1 raises P_2^* but lowers the probability of the dispersion region, so the sign of the second derivative of the continuation term is ambiguous. Theorem 1(i) establishes that P_1^B is not optimal and that every optimum lies strictly above it, but not that the optimum is unique. Uniqueness is verified numerically over the parameter range reported in Appendix A.4 and is not claimed outside it.

Remark 5 (saturation). The plateau in part (iii) is the launch-price counterpart of the complete compression of Corollary 3. Once comparison intensity is high enough that the firm never prices above the cohort benchmark, the benchmark binds in every relevant state, and the only remaining channel through which P_1 affects continuation value is the level at which the firm pools. That channel is independent of γ, so the premium becomes invariant to further increases in fairness sensitivity. The economically distinctive prediction is therefore a threshold followed by a plateau, not an interior maximum: markets differing in fairness intensity above $\hat{\gamma}$ should exhibit the same launch premium, while those below it should differ systematically.

Proposition 4 (Reference Ordering). Under Assumption 1, the equilibrium first-period price strictly exceeds the personal-equilibrium expectation reference, $P_1^* > R_2^*$, whenever $\beta > 0, \gamma = \pi\alpha > 0$ and the period-2 optimisation reaches the dispersion region with positive probability. Hence at the equilibrium references the comparison kink lies strictly above the expectation kink, so the ordering $R_2 < P_1$ illustrated in Figure 1 arises endogenously in equilibrium.

Proof. See Appendix A.2.8.

Theorem 1 and Proposition 4 are complementary, and the dependency runs in one direction only. Theorem 1 is proved for either ordering of the two references, since the envelope argument uses only that the comparison reference enters the period-2 perceived price through the term $\alpha\max\{0, P_2 - P_1\}$. Proposition 4 then invokes Theorem 1 to sign $\Psi(P_1^*) - P_1^*$ and thereby selects the ordering $R_2^* < P_1^*$. Because the expectation reference must satisfy personal equilibrium whereas the comparison reference is determined by the firm's own initial pricing decision, only the latter can be manipulated strategically, and forward-looking firms exploit this asymmetry by elevating the initial price in order to relax future behavioural constraints.

Theorem 1(iii) sharpens the empirical content of the mechanism in a way that distinguishes it from a simple monotone behavioural effect. The launch premium is not proportional to fairness sensitivity. It rises steeply where comparison concerns are weak, and becomes insensitive to them once they are strong enough that the firm never prices above the cohort benchmark. Sectors above that threshold should exhibit indistinguishable launch premia despite differing markedly in measured fairness sensitivity, while sectors below it should differ systematically. Combined with the inverted-U in capacity established in Section 4.4, the joint prediction is that the premium is greatest in markets with moderate expected load factors and fairness sensitivity

at or above the saturation threshold, and that variation in the premium across high-sensitivity markets should be driven by capacity rather than by fairness.

This is a more demanding test than a monotone prediction, because it implies a specific pattern of *non*-variation. A finding that launch premia vary continuously with measured fairness sensitivity across the whole range would be evidence against the mechanism, not for it.

4. Extension: T Selling Periods, Perishable Capacity and the Continuous-Repricing Limit

The baseline model has two selling periods. This section shows that the two-period structure is not a restriction that generates the result but the smallest case of a general one, and that the launch-price distortion is larger, not smaller, when the selling window is divided into many repricing epochs. It also addresses the perishable capacity constraint directly rather than as a robustness check.

4.1 The T-Period Model

A perishable good is sold over a window divided into T epochs, after which unsold units have zero salvage value. A fresh cohort of mass $1/T$ arrives in each epoch, so total potential demand is held fixed as T varies and the limit as T grows is the continuous-repricing limit rather than a longer horizon. Total capacity is K. In the launch epoch demand is deterministic; in every subsequent epoch the firm observes the demand shock and then posts a price. Sales are the lesser of realised demand and remaining inventory.

In the T-period model the fairness wedge is applied to aggregate demand at the composite intensity Γ rather than to a sensitive subpopulation. Under linear demand this is without loss: the mixture of Section 3.2 aggregates to $a - bP - b\pi\alpha(P - m)^+$, which is the demand of a representative consumer facing the wedge at intensity Γ. Demand, revenue and hence the inventory path coincide under the two formulations, so nothing in what follows depends on which is adopted. Suppressing the expectation reference, the behavioural state reduces to the prevailing comparison floor, which with remaining inventory constitutes the state.

Because the good is consumed once and by a different cohort in each epoch, no consumer accumulates a memory of her own past prices, and the adaptive anchor of the reference-price tradition has no counterpart. The only benchmark available to a buyer is the price paid by earlier buyers of the same good. We take the salient benchmark to be the lowest price any earlier cohort has paid, that is, the running minimum of past prices, which is the object of resentment reported in the fairness literature and the anchor used by Nasiry and Popescu (2011). The effective price faced by a behaviourally sensitive consumer adds a wedge equal to the composite fairness intensity multiplied by the excess of the current price over that floor. The first cohort has no predecessor, so its wedge is zero: the launch price is set free of any fairness penalty, and it establishes the floor that every later cohort inherits.

Two features distinguish this state variable from an adaptive reference. First, it is a ratchet: the floor falls when the firm cuts below it and never rises, so a markdown is an irreversible concession. Secondly, it is set once, at launch, by the one cohort that pays no penalty for it. The firm's problem is a dynamic program in the pair consisting of the floor and remaining inventory.

4.2 The Launch Premium at General T

Let the continuation value denote expected profit from the second epoch onward under optimal play, as a function of the prevailing floor. Because the launch cohort faces no wedge, total value is the sum of launch revenue, which carries weight $1/T$, and the continuation value evaluated at the launch price. At the static monopoly benchmark the launch-revenue term is stationary, so the sign of the distortion is the sign of the derivative of the continuation value with respect to the floor.

Proposition 5 (Launch Premium at General T).

Suppose the composite fairness intensity is strictly positive and the unconstrained period-2 optimum exceeds the prevailing floor with positive probability. Then the equilibrium launch price strictly exceeds the static monopoly benchmark for every T of at least two and for any capacity K.

Proof. See Appendix A.2.9.

The argument requires neither linear demand nor slack capacity, and it reduces to Theorem 1 when T equals two. The proviso is not cosmetic. If demand is stable enough that the firm's unconstrained period-2 optimum never exceeds the prevailing floor, the floor is slack in every state, an elevated floor confers no benefit, and no distortion arises. Note that strong fairness concerns alone do not produce this case: they widen the band over which the firm pools at the floor without moving the unconstrained optimum, and pooling at a higher floor is itself valuable; this is the plateau of Theorem 1(iii). The launch premium is a property of markets in which the firm expects to want to raise prices later, which is precisely the setting that motivates dynamic pricing in the first place.

4.3 The Continuous-Repricing Limit

Proposition 6 (Refinement Monotonicity of the Continuation Value). Suppose the conditions of Proposition 5 hold, and let $W_T(m, k)$ denote the value of the firm's problem from the second epoch onward when the selling window is divided into T epochs, given prevailing floor m and remaining inventory k. Then:

(i) *(refinement monotonicity)* whenever T divides T', $W_{T'}(m, k) \geq W_T(m, k)$ for every (m, k): dividing the selling window more finely weakly raises the value of any given floor;

(ii) *(comparative static in the launch weight)* holding the continuation problem fixed, the optimal launch price is non-decreasing as the weight on launch revenue falls, since the launch-revenue term is strictly decreasing in price above the static benchmark and therefore the objective exhibits decreasing differences in the launch price and that weight;

(iii) *(uniform positivity)* by Proposition 5 the launch premium is strictly positive at every $T \geq 2$ and for every capacity K, so it does not vanish as the window is refined.

Proof. See Appendix A.2.10.

Remark 6. Parts (i) and (ii) push the launch price in the same direction: refining the window both raises the value of an elevated floor and reduces the cost of elevating it. They do not combine into a proof that the premium is monotone in T, because (i) signs the level of W_T but not its derivative in the floor, which is the object entering the launch first-order condition. The monotonicity of the premium in T, and its convergence as the window is refined without bound, are therefore established numerically in Table 5 rather than analytically, and are described as such in Section 4.6.

The economic content is that continuous algorithmic repricing does not compete the launch distortion away. A firm that reprices once has little to gain from an elevated benchmark, whereas a firm that reprices continuously is judged against that benchmark by every subsequent buyer, and by Proposition 6(iii) the premium is strictly positive at every repricing frequency. Table 5 reports the launch premium computed numerically under linear demand and uniformly distributed shocks with slack capacity. The premium is exactly zero at every horizon when fairness concerns are absent, confirming that the distortion is behavioural rather than an artefact of the horizon or of capacity. It rises as the window is divided more finely and then flattens, consistent with Proposition 6(i) and (ii) and with the convergence discussed in Section 4.6.

The fairness dimension exhibits the saturation established in Theorem 1(iii). The premium rises steeply at low intensity, then becomes exactly invariant above a threshold of roughly 0.6 under this calibration. Beyond that point the firm never prices above the cohort benchmark in any state, so the comparison term is inactive in equilibrium and further increases in fairness sensitivity leave the launch decision unchanged. What sustains the premium there is not the avoidance of a wedge but the level at which the firm pools: with the benchmark binding in every state, a higher benchmark is simply a higher price. The premium therefore plateaus rather than reverting, and the plateau is reached at a lower intensity the more finely the window is divided.

Table 5. Launch premium as the selling window is divided more finely and as fairness intensity rises (slack capacity, shock dispersion 0.5).

Fairness intensity Γ	$T = 2$	$T = 4$	$T = 8$	$T = 16$
0 (no fairness)	0.000	0.000	0.000	0.000
0.10	+0.014	+0.022	+0.024	+0.025
0.25	+0.028	+0.042	+0.045	+0.045
0.50	+0.040	+0.057	+0.060	+0.061
0.75	+0.043	+0.059	+0.063	+0.063
1.00	+0.043	+0.059	+0.063	+0.063
1.50	+0.043	+0.059	+0.063	+0.063
2.00	+0.043	+0.059	+0.063	+0.063
3.00	+0.043	+0.059	+0.063	+0.063
4.00	+0.043	+0.059	+0.063	+0.063
6.00	+0.043	+0.059	+0.063	+0.063

Note: Under linear demand $D(P) = 1 - P$ the static monopoly benchmark is $P^B = 0.5$; shocks independent and uniform on $[-0.5, 0.5]$; the launch epoch faces no demand shock, so the no-fairness launch price is exactly P^B; slack capacity; no within-window discounting. The launch epoch faces no demand shock, so the no-fairness launch price is exactly P^B. The premium rises in fairness intensity and then plateaus at a strictly positive level, as Theorem 1(iii) predicts. The saturation point $\hat{\gamma}$ is 0.68 at $T = 2$, 0.61 at $T = 4$, 0.60 at $T = 8$ and 0.58 at $T = 16$; at each, the corresponding $\hat{\gamma}(P_1^*)$ is $0.76, 0.68, 0.66$ and 0.66, confirming that saturation occurs where the comparison intensity exhausts the dispersion region at the equilibrium launch price. The period-1 objective was verified to be single-peaked at every grid point, as Remark 4 requires. Values computed on a price grid of width 0.001 with fifteen-node Gauss–Legendre quadrature over the shock.

4.4 Perishable Capacity

Capacity enters the T-period problem as a second state variable, and its effect on the behavioural distortion is not monotone. Where capacity is slack the constraint is irrelevant to the launch decision and the distortion takes its slack-capacity value exactly; the premium begins to rise only once the constraint binds. Where capacity is severely scarce the firm already prices high in order to ration a fixed stock, so the floor is slack, the firm rarely wishes to price above it, and the fairness wedge is seldom paid. The distortion is consequently largest at intermediate scarcity. Table 6 isolates the behavioural effect by reporting, at each capacity level, the difference between the launch price with and without fairness concerns, so that the scarcity premium common to both is differenced out.

Table 6. Fairness-attributable launch premium by capacity (T = 8, shock dispersion 0.5).

Capacity K	Launch price, no fairness	Launch price, with fairness	Fairness premium
0.90	0.500	0.563	+0.063
0.70	0.500	0.564	+0.064
0.60	0.501	0.580	+0.079
0.50	0.532	0.625	+0.093
0.45	0.566	0.656	+0.090
0.35	0.655	0.730	+0.075
0.25	0.756	0.818	+0.062
0.15	0.872	0.913	+0.041

Notes: Fairness intensity is set at $\Gamma = 1.00$, within the saturation region of Theorem 1(iii), so that the capacity comparative static is not confounded by movement along the Γ schedule. Shocks are i.i.d. uniform on [−0.5, 0.5]; the launch epoch faces no demand shock, so the no-fairness launch price equals the static benchmark $P^B = 0.500$ wherever capacity is slack. The capacity constraint begins to bind between $K = 0.70$ and $K = 0.60$, and the premium is invariant to capacity above that point: at $K = 0.90$ it coincides with the slack-capacity plateau of Table 5 at $T = 8$, confirming that the two exercises share a calibration. Solved by backward induction on the price and inventory grids described in Appendix A.4, with 81 inventory nodes.

This inverted-U separates the present mechanism sharply from Selove (2019), in which scarcity is what generates the fairness distortion, since rationing during peak demand is the object of that analysis. Here scarcity ultimately attenuates the distortion, and does so below the slack-capacity level: an event certain to sell out has little need of a behavioural floor. Relative to the slack benchmark of +0.063, moderate scarcity amplifies the launch premium by roughly half, to +0.093 at $K = 0.50$, while severe scarcity reduces it by a third, to +0.041 at $K = 0.15$. The prediction is testable, since the launch premium should be largest for events and routes with moderate, rather than extreme, expected load factors, and should be smaller for those that reliably sell out than for those that reliably do not.

4.5 The Price Path and the Joint Prediction

The equilibrium path declines. The value of maintaining a high floor is proportional to the number of cohorts still to arrive, so as the window closes the firm progressively concedes the ratchet and marks down. The model therefore generates skimming endogenously, with the highest price at launch, for a reason unavailable in the adaptive-reference tradition, where a declining path requires an exogenously high initial reference. Table 7 reports simulated paths and confirms at general T the joint prediction developed in Section 5: as fairness concerns strengthen, the launch price rises while the dispersion of prices across cohorts falls.

Table 7. Launch price and cross-cohort dispersion ($T = 8$, shock dispersion 0.5, slack capacity; 4,000 simulated shock paths).

Fairness intensity	Launch price	Mean spread (max less min)	Mean standard deviation
0.00	0.500	0.374	0.123
0.25	0.545	0.222	0.074
0.50	0.560	0.164	0.056
1.00	0.563	0.146	0.049
2.00	0.563	0.144	0.048

Notes: Dispersion statistics are computed across the T posted prices within each simulated selling window and then averaged across paths. Shocks are drawn on the quadrature nodes used in the dynamic program, so the simulation and the optimisation face an identical shock support; shocks are i.i.d. uniform on [−0.5, 0.5] and the launch epoch faces no demand shock. The launch price saturates at $\Gamma \approx 0.75$ and coincides with the slack-capacity plateau of Table 5 at $T = 8$ and with the $K = 0.90$ row of Table 6. Dispersion continues to fall marginally beyond saturation because the floor, though no longer raising the launch price, continues to compress prices in later epochs. The mean spread falls by 61 per cent between $\Gamma = 0$ and $\Gamma = 1.00$ while the launch price rises by 12.6 per cent, which is the joint prediction in its sharpest form: the two move in opposite directions and both saturate together.

4.6 Limitations

Propositions 5 and 6 hold for general demand and for any capacity, but the magnitudes reported in Tables 5 to 7 are numerical, computed under linear demand, uniformly distributed shocks and no discounting within the selling window. Both the monotonicity of the launch premium in the number of repricing epochs and its convergence as the window is divided more finely are observed rather than proved: Proposition 6 establishes refinement monotonicity of the continuation value and the direction of the launch-weight comparative static, but with two state variables, the floor and remaining inventory, and a kinked payoff, these do not combine into a closed-form monotone comparative static in T. The comparative static in capacity is likewise numerical. These results establish the direction and the robustness of the distortion rather than its precise magnitude.

The saturation threshold reported in Table 5 is calibration-dependent. It is determined by the position of the upper support point of the demand shock relative to twice the equilibrium launch price, so it will be higher in markets with more volatile demand and lower in markets where the launch price is already close to half the maximum realisable intercept. The qualitative pattern, strict increase, then exact invariance, is a property of the model rather than the calibration, but the numerical location of the threshold should not be read as a quantitative prediction.

5. Implications and Empirical Predictions

Section 3 established that forward-looking firms optimise over two behavioural reference objects alongside conventional economic state variables: an expectation reference, determined

in personal equilibrium, which smooths responses to demand shocks; and a comparison reference, created by the firm's own initial price, which compresses cross-cohort price differences and generates the forward-looking elevation of the initial price. This section examines the resulting empirical content: the distinctive joint prediction that high-comparison markets exhibit both higher initial prices and smaller subsequent price differences; the magnitude of the equilibrium distortions under plausible behavioural parameters; and the implications for pricing strategy across industries.

The remainder of this section examines the implications of this behavioural equilibrium from three complementary perspectives. First, it considers whether the behavioural mechanisms identified by the model are consistent with existing experimental and field evidence on reference dependence, inequity aversion, social identity and procedural legitimacy. Secondly, it assesses the magnitude of the equilibrium distortions generated under empirically plausible behavioural parameters. Finally, it considers the broader implications of the framework for algorithmic pricing, market transparency, competition policy and consumer welfare. Throughout, the emphasis is not simply on whether behavioural preferences influence pricing decisions, but on how firms internalise endogenous behavioural reference formation when solving the dynamic optimisation problem.

5.1 Distinctive Empirical Predictions

The framework generates an empirical prediction that separates it from the nearest existing theories. Conventional models explain intertemporal price variation through changes in demand, inventories, capacity constraints and learning (Gallego and Van Ryzin, 1994; Talluri and Van Ryzin, 2004; Phillips, 2005; Krasikov and Lamba, 2026). Behavioural models show that reference dependence and fairness influence consumers' responses to observed prices (Kahneman, Knetsch and Thaler, 1986; Xia, Monroe and Cox, 2004; Kőszegi and Rabin, 2006), and a smaller literature allows firms to internalise the references their own prices create (Popescu and Wu, 2007; Nasiry and Popescu, 2011; Selove, 2019). The present framework belongs to that third group but differs within it. Where the reference is formed from a consumer's own price history, building it carries a contemporaneous behavioural cost and the equilibrium price lies below the non-behavioural benchmark; where the reference is inherited across cohorts, that cost is absent and the distortion reverses sign.

The empirical implication is a joint pattern rather than a single one. Conditional on demand, inventories, capacity constraints and other conventional determinants, markets characterised by greater comparison salience should exhibit lower subsequent price differences throughout, and higher initial prices up to a saturation point beyond which further increases in salience leave initial prices unchanged even as dispersion remains fully compressed. Reduced price flexibility alone is consistent with several behavioural accounts; adaptive-reference mechanisms predict compression accompanied by a *lower* initial price. The coexistence of an elevated launch price and compressed subsequent variation is therefore difficult to generate without a forward-looking benchmark inherited across cohorts.

The saturation property makes the prediction sharper than a monotone comparative static, and correspondingly easier to refute. Sectors with comparison salience above the threshold should show indistinguishable initial-price elevation despite differing markedly in measured salience, while sectors below it should differ systematically. A finding that initial prices rise continuously with comparison salience across the whole observed range would count against the

mechanism rather than for it. Section 5.5 sets out two identification strategies that exploit this structure.

5.2 Behavioural Evidence and Cross-Sector Variation

The behavioural mechanisms developed in Section 3 help explain why dynamic pricing produces markedly different outcomes across industries despite increasingly similar pricing technologies. Existing empirical evidence demonstrates that consumers exhibit reference dependence, inequity aversion and fairness concerns in a wide variety of settings (Kahneman, Knetsch and Thaler, 1986; Xia, Monroe and Cox, 2004). However, the intensity of these behavioural responses varies systematically across markets. The present framework provides an explanation for this heterogeneity by distinguishing between expectation-based and comparison-based behavioural reference points and by showing how each affects firms' optimal pricing strategy.

The expectation reference primarily influences the responsiveness of equilibrium prices to changing market conditions. Markets in which consumers form strong expectations regarding future prices should therefore exhibit smoother equilibrium pricing paths because firms have incentives to avoid creating unexpected price increases that trigger loss aversion. This mechanism is consistent with the relatively gradual adjustment of prices observed in industries such as airlines and hotels, where consumers have become accustomed to predictable pricing rules and repeated interactions over time (Talluri and Van Ryzin, 2004; Phillips, 2005).

The comparison reference operates through a different channel. Where prices paid by different consumers are highly visible, firms face stronger incentives to limit profitable subsequent price differences because later prices are evaluated relative to earlier transactions. This prediction is consistent with the greater controversy surrounding dynamic pricing in live entertainment, sports ticketing and ride-sharing, where consumers can readily compare prices across purchasers and where perceived inequities frequently generate adverse behavioural responses (Xia, Monroe and Cox, 2004; Haws and Bearden, 2006).

Direct estimates of the comparison parameters are scarce, but Richards, Liaukonytė and Streletskaya (2016) provide a benchmark for π and α. Their disadvantageous-comparison coefficient is estimated as a random parameter, and the implied share of subjects with a positive coefficient (the empirical counterpart of the behaviourally sensitive share π) lies between roughly 0.57 and 0.81 across their specifications. The calibration adopted below sets $\pi = 0.5$ at the baseline and $\pi = 0.3$ in the pre-reform counterfactual of Section 5.4, both at or below the lower end of that range. Expressed relative to their estimated price coefficient, their comparison intensity likewise exceeds the range of α considered here. Their setting differs from the present one in ways that caution against transporting the magnitudes directly: the good is a single discretionary purchase against a strong outside option, the subjects are students, and comparison is contemporaneous rather than cross-cohort. The relevant point is one of direction. On the only direct evidence available, the calibration understates rather than overstates the behavioural intensities, so the launch premia reported in Section 5.3 should be read as conservative.

Propositions 2 and 3 further imply that identical pricing algorithms need not generate identical market outcomes. Social identity strengthens behavioural comparisons by increasing the weight consumers attach to prices paid by others, while procedural legitimacy moderates these effects by increasing the perceived acceptability of price differences. Industries therefore differ

not simply because firms employ different pricing technologies, but because consumers interpret identical pricing behaviour through different institutional and social contexts. The model consequently provides an explanation for why dynamic pricing has become widely accepted in some markets while remaining persistently contentious in others.

Taken together, these observations suggest that cross-industry variation in dynamic pricing reflects differences in behavioural environments as well as differences in underlying economic conditions. Conventional determinants such as demand volatility, inventories and capacity constraints remain important, but they are complemented by behavioural reference formation, social comparison and procedural legitimacy. The equilibrium diversity of pricing strategies observed across markets is therefore consistent with firms optimising jointly over economic conditions and consumer behavioural responses rather than over economic conditions alone.

The sectoral pattern this implies is not simply that comparison-prone markets are more constrained. Section 5.4 shows that the behavioural distortion is largest at intermediate capacity scarcity: where a sell-out is near-certain the firm prices high for rationing reasons and the behavioural floor is slack, while where capacity is abundant there is little future rent for an elevated floor to protect. The attenuation at high scarcity is not merely relative: at the tightest capacity considered the launch premium falls to roughly two-thirds of its slack-capacity value, so venues that reliably sell out are predicted to exhibit smaller launch premia than those operating with comfortable capacity, and not merely smaller premia than those in between. Airlines and hotels, which combine institutionalised fare structures with high and variable load factors, therefore sit at one end of this distribution; event ticketing and live entertainment, where comparison is salient, legitimacy contested and capacity moderately binding, sit at the other. The prediction is that launch premia should be largest in the latter, and this is where public controversy over dynamic pricing has in fact concentrated.

5.3 Quantitative Implications

Theorem 1 establishes the paper's central prediction: because current prices become future behavioural comparison benchmarks, forward-looking firms optimally elevate the initial price in anticipation of future consumer evaluations. Proposition 1 and the accompanying corollaries show how expectation-based reference dependence attenuates responses to demand shocks and how comparison-based inequity aversion compresses equilibrium price differences. The numerical implementation reported below and in Appendix A.4 complements these analytical results by illustrating the economic magnitude of the resulting distortions under representative behavioural parameter values. The objective is not structural estimation but an assessment of how far empirically plausible behavioural responses move equilibrium pricing from conventional dynamic-pricing benchmarks.

The calibration demonstrates that the behavioural mechanisms developed in Section 3 generate quantitatively meaningful effects. The expectation reference produces noticeable attenuation of firms' responses to demand shocks, while the comparison reference substantially compresses equilibrium price differences across consumer cohorts. Theorem 1's forward-looking pricing incentive also proves economically significant, although the resulting elevation of the optimal initial price remains modest because the expectation channel simultaneously smooths subsequent price adjustments. Stronger comparison salience systematically increases the compression of price dispersion, while the initial-price elevation reflects the interaction between the comparison and expectation channels.

The broader numerical exercises, including the disclosure-reform counterfactual reported in Appendix A.4, also illustrate the interaction between behavioural responses and procedural legitimacy. Increasing legitimacy weakens the behavioural penalty associated with observed price differences, allowing firms to sustain greater equilibrium price dispersion while reducing the incentive identified in Theorem 1 to elevate initial prices. This result complements Proposition 3 by demonstrating that institutional design influences not only consumer acceptance of dynamic pricing but also firms' optimal pricing strategies.

More broadly, the numerical analysis demonstrates that behavioural effects are not merely theoretical curiosities. Under plausible behavioural parameter values, endogenous behavioural reference formation generates economically meaningful compression of cross-cohort price differences together with the forward-looking elevation of the initial price predicted by Theorem 1. Although this elevation is typically only of the order of a few to ten per cent, it represents a qualitatively novel distortion absent from conventional dynamic-pricing models. The quantitative results therefore reinforce the central message of the paper: behavioural reference formation represents an economically important component of firms' dynamic optimisation problem rather than a minor refinement to existing pricing models.

Table 8 reports representative equilibrium distortions implied by the unified model under the central behavioural calibration (Appendix A.4). Rather than estimating behavioural parameters, the exercise illustrates the economic magnitude of the principal theoretical mechanisms developed in Section 3. Panel (i) quantifies the attenuation of demand-shock pass-through predicted by Proposition 1 and Corollary 1. Panel (ii) reports the width of the complete-smoothing region. Panel (iii) quantifies the forward-looking pricing incentive characterised in Theorem 1, showing that the optimal initial price is elevated above the static benchmark, but by a modest amount once the expectation channel is incorporated. The compression of cross-cohort price differences and the moderating role of procedural legitimacy (Proposition 3) are examined separately in the disclosure-reform counterfactual. Panel (iii) also displays a non-monotonicity across columns, which the notes explain: because the behaviourally sensitive share enters both behavioural composites, the two channels work against one another as it rises.

Table 8. Representative quantitative implications of the unified behavioural dynamic pricing model.

(i) Shock attenuation pure $1/(1+\kappa)$ → combined $1/(1+\kappa+\gamma)$ [Corollary 1 → Proposition 1]

	$\boldsymbol{\pi = 0.3}$	$\boldsymbol{\pi = 0.5}$	$\boldsymbol{\pi = 0.7}$
$\boldsymbol{\lambda = 1.5}$	0.87 → 0.69	0.80 → 0.57	0.74 → 0.49
$\boldsymbol{\lambda = 2.0}$	0.77 → 0.62	0.67 → 0.50	0.59 → 0.42
$\boldsymbol{\lambda = 2.5}$	0.69 → 0.57	0.57 → 0.44	0.49 → 0.36

(ii) Complete-smoothing region (pooling at $\boldsymbol{R_2}$) width $\kappa/2 = \pi(\lambda-1)/2$ (fraction of demand intercept) [Proposition 1]

	$\boldsymbol{\pi = 0.3}$	$\boldsymbol{\pi = 0.5}$	$\boldsymbol{\pi = 0.7}$
$\boldsymbol{\lambda = 1.5}$	0.08	0.13	0.18
$\boldsymbol{\lambda = 2.0}$	0.15	0.25	0.35
$\boldsymbol{\lambda = 2.5}$	0.23	0.38	0.53

(iii) Forward-looking period-1 price elevation full-equilibrium value, % above static benchmark ($\lambda = 2, \beta = 1$) [Theorem 1]

	$\boldsymbol{\pi = 0.3}$	$\boldsymbol{\pi = 0.5}$	$\boldsymbol{\pi = 0.7}$
$\boldsymbol{\alpha = 0.5}$	3.9	5.4	5.5
$\boldsymbol{\alpha = 1.0}$	6.6	8.0	7.7
$\boldsymbol{\alpha = 1.5}$	8.8	9.5	7.7

(iv) Sensitivity of the elevation to shock dispersion % above static benchmark (α = 1, λ = 2, β = 1) [Theorem 1(i), pooling term]

Shock support	$\boldsymbol{\pi = 0.3}$	$\boldsymbol{\pi = 0.5}$	$\boldsymbol{\pi = 0.7}$
$\boldsymbol{d = 0.9}$	6.6	8.0	7.7
$\boldsymbol{d = 0.5}$	5.0	4.0	1.5
Pr[dispersion], $\boldsymbol{d = 0.5}$	0.114	0.000	0.000

Notes. The baseline sets $\lambda = 2, \alpha = 1$ and $\pi = 0.5$, so $\kappa = \pi(\lambda - 1) = 0.5$ and $\gamma = \pi\alpha = 0.5$; demand is linear, $D(P) = a - bP$, with $\varepsilon \sim \text{Uniform}[-d, d], \beta = 1$ and static benchmark $P_1^B = \frac{a_1}{2b}$. Panels (i) to (iii) set $d = 0.9$.

Panel (i) reports the pass-through of a period-2 demand shock to the optimal price. The pure factor $1/(1 + \kappa)$ is the expectation channel alone (Corollary 1, $\gamma = 0$); the combined factor $1/(1 + \kappa + \gamma)$ is the unified value where both reference kinks bind (Proposition 1), with $\alpha = 1$ throughout. A smaller factor means stronger smoothing, so comparison-based inequity aversion reinforces expectation-based smoothing and the combined pass-through lies uniformly below the pure value (0.67 → 0.50 at the baseline).

Panel (ii) reports the width of the complete-smoothing region around the expectation reference R_2; under the benchmark calibration it equals $\kappa/2$ as a fraction of the demand intercept and is unaffected by the comparison channel. The values are rounded-up.

Panel (iii) is obtained by solving jointly for the optimal first-period price and the personal-equilibrium reference $R_2 = E(P_2)$ ($\sigma\varphi = 0$). Because π enters both the comparison composite $\gamma = \pi\alpha$ and the smoothing composite $\kappa = \pi(\lambda - 1)$, the comparative static across columns is confounded: the non-monotonicity in π reflects the offsetting influence of the two channels and does not bear on the monotonicity in γ established in Theorem 1(iii). Incorporating the expectation channel attenuates the forward-looking incentive because smoother period-2 prices reduce the behavioural gains from raising the initial price, so the equilibrium elevation remains in the single digits. The $\alpha = 1.0, \pi = 0.3$ cell coincides with the pre-reform elevation of the $\sigma\varphi = 0$ row of Table 9, and the $\alpha = 1.5, \pi = 0.3$ cell with that of the $\sigma\varphi = 0.5$ row, since the two exercises share the same effective comparison intensity and smoothing composite in those configurations.

Panel (iv) reports the sensitivity of the elevation to the width of the shock support. Panels (i) to (iii) use $d = 0.9$, a wider support than the $d = 0.5$ of Tables 5 to 7; the two sets of exercises concern different models — the two-period unified model here, the T-period comparison-only extension there; so their magnitudes are not directly comparable. Narrowing the support reduces the elevation, by roughly a quarter at $\pi = 0.3$ and by four-fifths at $\pi = 0.7$, because a narrower shock distribution leaves fewer states in which the firm would price above the cohort benchmark. The elevation nonetheless remains strictly positive at every configuration, including those in which the dispersion region is empty in equilibrium: at $\pi = 0.5$ and $\pi = 0.7$ under the narrower support the firm never prices strictly above the comparison benchmark, yet the launch price still exceeds the static benchmark by 4.0 and 1.5 per cent respectively. This is the pooling term of Theorem 1(i) in isolation, the elevation is sustained not by avoiding the comparison wedge but by the level at which the firm pools, and it is the two-period counterpart of the plateau established in Theorem 1(iii).

The loss-aversion coefficient $\lambda = 2$ is representative of estimates in the behavioural and marketing literature (DellaVigna, 2009; Hardie, Johnson and Fader, 1993; Pope and Schweitzer, 2011); the behaviourally sensitive share π is calibrated using evidence on the prevalence of Fehr–Schmidt inequity aversion (Bellemare, Kröger and van Soest, 2008), and lies at or below the range implied by the estimates of Richards, Liaukonytė and Streletskaya (2016) discussed in Section 5.2. The reported values are therefore intended to illustrate the economic significance of the theoretical mechanisms rather than to provide structural estimates for a particular industry.

The broader sensitivity analysis reported in Appendix A.4 demonstrates that these results are not driven by a particular parameter configuration. Across a wide range of behavioural parameter values, stronger comparison-based responses systematically compress equilibrium price differences, and greater procedural legitimacy moderates both distortions. The forward-looking elevation behaves differently, and instructively so. It is monotone in the comparison intensity Γ, as Theorem 1(iii) establishes, but not in the primitives that determine it: raising the

behaviourally sensitive share π strengthens the comparison channel while simultaneously strengthening the expectation channel, and beyond a threshold the second dominates. The non-monotonicity visible in panel (iii) is therefore a property of the primitives rather than of the mechanism, and the saturation it reflects is itself a prediction of the model.

The numerical analysis therefore supports the paper's central theoretical conclusion. Endogenous behavioural reference formation generates substantial compression of cross-cohort price differences together with a more modest, but qualitatively novel, elevation of the optimal initial price. These results demonstrate that behavioural reference formation is not merely an ancillary feature of dynamic pricing but an economically significant state variable shaping firms' intertemporal pricing decisions.

5.4 Implications for Pricing Strategy and Market Design

The framework carries broader implications for the design of pricing systems. Transparency requirements, disclosure obligations and algorithmic governance are conventionally treated as external constraints on firms' pricing behaviour. The present framework identifies an additional channel: anything that alters consumers' ability to observe, interpret or compare prices, whether a regulatory disclosure rule, a comparison platform or a firm's own presentation of its price history, also modifies the behavioural reference points that determine future demand. Such changes therefore alter the environment over which forward-looking firms optimise and, through the mechanism identified in Theorem 1, change firms' incentives to manage the reference points created by today's pricing decisions.

This distinction is particularly relevant in digital markets, where comparison websites, online platforms and social media have substantially increased the visibility of historical prices and inter-consumer price differences. At the same time, policymakers have proposed or introduced measures requiring greater transparency regarding pricing practices and algorithmic decision-making. Such reforms are generally justified on the grounds that they improve consumer information and facilitate more informed purchasing decisions. The present framework suggests that they also influence firms' equilibrium pricing strategies by altering the behavioural environment within which pricing decisions are made.

To illustrate this mechanism, Table 9 reports an equilibrium counterfactual in which comparison salience increases through a larger share of consumers engaging in comparison-based behavioural evaluations, while the underlying demand environment remains unchanged. The counterfactual illustrates the comparative statics implied by Proposition 1 and the forward-looking pricing incentive characterised in Theorem 1 rather than evaluating a particular regulatory intervention.

Table 9. Equilibrium counterfactual: a disclosure reform raising the behavioural share (π: $0.30 \rightarrow 0.60$; $\alpha = 1, \beta = 1, \lambda = 2$)

Representative sector (identity salience $\sigma\varphi$)	Shock pass-through $1/(1 + \kappa + \Gamma)$, pre → post	Period-1 elevation, pre / post	Subsequent Price Compression vs. Benchmark, pre / post
Airlines, hotels ($\sigma\varphi = 0$)	0.63 → 0.45	6.6% / 8.2%	−82% / −100%
Ride-sharing, sport ($\sigma\varphi = 0.5$)	0.57 → 0.40	8.8% / 8.5%	−91% / −100%
Live entertainment ($\sigma\varphi = 1.0$)	0.53 → 0.36	10.0% / 8.5%	−96% / −100%

Notes: Period-1 elevation is the full-equilibrium percentage above the static monopoly benchmark $P_1^B = a_1/2b$, solved jointly with the personal-equilibrium reference $R_2 = E[P_2]$ ($\lambda = 2, \beta\ 1$) on the calibration of Appendix A.4. The pre-reform columns use $\pi = 0.30$ and the post-reform columns $\pi = 0.60$; α and $\sigma\varphi$ are held fixed within each row, so the reform raises both the comparison composite $\Gamma = \pi\alpha(1 + \sigma\varphi)$ and the smoothing composite $\kappa = \pi(\lambda - 1)$. Because $\sigma\varphi$ enters Γ but not κ, the three post-reform cells share $\kappa = 0.60$ and differ only in Γ. Every post-reform cell lies in the saturation region of Theorem 1(iii): the equilibrium probability of the dispersion region is 0.035, 0.000 and 0.000 respectively, so the comparison term is inactive and the premium is sustained entirely by the level at which the firm pools. Elevation is therefore invariant to further increases in Γ and converges across sectors, even though the pre-reform values differ by more than three percentage points. The shock pass-through column is the interior factor $1/(1 + \kappa + \Gamma)$ of Table 8 panel (i); a smaller value means stronger smoothing. Compression is measured relative to the behaviour-free benchmark ($\Gamma = \kappa = 0$) as the ratio of expected upward cross-cohort dispersion, taken as an expectation over the shock distribution; post-reform the firm pools at the comparison benchmark P_1 in every state ($\Delta^* = 0$), so compression saturates at −100%. Setting $\lambda = 1$ reproduces the comparison-only results.

The counterfactual illustrates a broader implication of the framework. Policies that strengthen opportunities for behavioural comparison do not have proportionally larger effects in markets where comparison is already salient. Where the pre-reform environment leaves the firm pricing above the cohort benchmark in a substantial set of states, the reform raises the launch premium materially. Where comparison intensity is already high enough that the benchmark binds in nearly every state, the reform operates almost entirely on the compression margin and leaves the launch premium close to unchanged. The prediction is therefore one of convergence rather than divergence: sectors that differ appreciably in launch behaviour before a transparency reform are predicted to look alike after one, with the residual variation appearing in the pass-through of demand shocks rather than in the level of the launch price. This is a demanding test, since it requires the cross-sectional dispersion in launch premia to narrow while the dispersion in shock pass-through persists.

A comparison with the stability analysis of Richards, Liaukonytė and Streletskaya (2016) clarifies what the present mechanism does and does not predict. In their simulations, progressively widening the dispersion of prices across buyers under a posted-price regime drives purchase probability from roughly thirteen per cent to under five per cent, a level at which they treat the pricing regime as commercially unviable; under negotiated pricing the same widening leaves purchase probability broadly intact. Their threshold is one of market collapse:

dispersion is imposed, and demand eventually fails to support it. The threshold in the present model is of a different kind. Here dispersion is chosen, and the firm anticipates the behavioural response, so the constraint binds before demand collapses: as comparison intensity rises the firm compresses dispersion to zero and prices at the cohort benchmark, and the launch premium saturates rather than the market failing. The two results are complementary. Theirs identifies the dispersion a market will tolerate; this paper identifies the dispersion a forward-looking firm will choose, and shows that the second is reached first.

The model predicts that greater comparison salience does not simply constrain firms' ability to price discriminate. Instead, forward-looking firms respond endogenously by adjusting the entire equilibrium price path. Stronger behavioural comparisons increase the incentive to establish more favourable comparison reference points through higher initial prices while simultaneously compressing subsequent price differences in order to reduce the behavioural costs associated with visible price disparities. The resulting equilibrium therefore reflects firms adapting optimally to a changed behavioural environment rather than merely complying with an external regulatory constraint.

More broadly, the analysis suggests that the effectiveness of dynamic-pricing policy depends not only on the pricing outcomes regulators seek to influence but also on the behavioural environment within which those outcomes emerge. Measures that improve procedural legitimacy, clarify pricing rules and increase consumer understanding may reduce adverse behavioural responses without eliminating the efficiency benefits associated with dynamic pricing. More fundamentally, the framework suggests that policy influences firms' dynamic pricing strategies by changing the endogenous behavioural reference points over which forward-looking firms optimise, rather than merely by constraining pricing behaviour directly.

5.5 An Empirical Strategy for Testing the Mechanism

The model is estimable in principle but is not directly testable in levels, because the static monopoly benchmark against which the launch premium is defined is unobservable. What the model does deliver, and what can be taken to data, are two comparative statics, each of which admits an identification strategy using observable ticketing data. We set them out here because the framework should be judged partly by whether its distinctive predictions can be refuted.

The first design exploits plausibly exogenous variation in the salience of price comparisons. All-in pricing requirements, which require sellers to display the full mandatory price a buyer will pay rather than a base price followed by additional fees, have been introduced at different dates across jurisdictions by legislatures and regulators. There is direct experimental evidence that such presentation matters in ticket markets. In a large-scale randomized field experiment on StubHub, Blake et al. (2021) show that disclosing mandatory fees upfront changes purchasing behaviour and that price shrouding makes comparison across ticket offers more difficult. All-in pricing therefore provides a source of variation in the salience and comparability of transaction prices that is plausibly exogenous to event-level pricing decisions. To the extent that buyers observe or communicate prices offered to other buyers of the same event, greater price salience should in turn raise the effective intensity of cross-cohort comparison represented by the model.

The prediction concerns the shape of the price path rather than the displayed headline price alone. Following an increase in comparison intensity, the model predicts a higher initial all-in price together with smaller subsequent price differences across cohorts: the launch price rises

while the inter-cohort price path compresses. Prices should therefore be measured on a common all-in basis both before and after a disclosure reform, adding mandatory fees to pre-reform headline prices where necessary. This avoids confusing the predicted increase in the economic launch price with the purely mechanical incorporation of previously shrouded fees into the displayed price. The joint prediction is what gives the test discriminatory power. A change in the headline price alone may reflect presentation; a rise in the reconstructed all-in launch price accompanied by compression of subsequent cohort differentials is not generated mechanically by fee disclosure. Moreover, this joint movement is opposite to the prediction of the adaptive-reference benchmark considered above, in which the behavioural channel lowers rather than raises the optimal initial price.

A stronger implementation of this design would exploit variation in the observability of earlier prices. All-in pricing directly increases the comparability of contemporaneously displayed prices, but does not by itself reveal what earlier cohorts paid. The model therefore predicts that the response to an all-in pricing reform should be larger where earlier prices are more readily observable; for example through secondary-market platforms, price-tracking services, or other public records of ticket prices. An interaction between the disclosure reform and pre-existing cross-cohort price observability would help distinguish the mechanism proposed here from a conventional search-cost or price-salience effect.

The second design tests the capacity result of Section 5.4 and has the advantage that it discriminates between this framework and existing fairness models rather than merely corroborating it. A single artist plays venues of very different capacities on the same tour, so artist-by-tour fixed effects absorb much of the unobservable component of the benchmark price, while identification comes from variation in venue capacity and local market demand across dates of the same tour. Regressing the launch price on expected scarcity, measured by the expected load factor, and its square, the model predicts an inverted-U relationship between scarcity and the magnitude of the behavioural distortion. Equivalently, conditional on the linear effect of scarcity, the coefficient on the quadratic term should be negative. The distortion is largest at intermediate scarcity: where a sell-out is near-certain, the promoter prices high for rationing reasons and the behavioural floor is relatively unimportant; where capacity is abundant, there is little future rent for an elevated initial-price benchmark to protect. The prediction is the inverse of that implied by models in which scarcity itself generates the fairness distortion, so the curvature of the relationship provides a means of separating the two mechanisms.

The data requirements are demanding but largely observable. Primary-market prices for individual events can be collected from the major ticketing platforms, and repeated retrieval from the initial on-sale date to the event date can recover both the launch price and the subsequent price path. To make comparisons meaningful, prices would need to be recorded at the ticket-category or seat-quality level, and on an all-in basis including mandatory fees. Venue capacities are published. Expected demand, and hence the expected load factor, can be proxied using predetermined measures such as artist popularity, venue-market characteristics, and the sell-out record of previous tour dates. Secondary-market listings and transactions, where observable, provide an independent measure of scarcity and buyer valuations and can be used to validate the demand proxy. The principal practical constraint is that historical primary-market price paths are difficult to reconstruct retrospectively, so the panel is most credibly assembled prospectively.

Neither design identifies the structural behavioural parameters, and neither is attempted here. Their purpose is instead to show that the mechanism generates predictions that are sharp, jointly signed, and distinguishable from those of nearby alternative theories, and that the principal variables required to test those predictions are observable. In particular, the first design combines a policy-induced change in price salience with a prediction about the entire inter-cohort price path, while the second exploits a non-monotonic prediction relating the behavioural distortion to scarcity. Failure of either predicted pattern would constitute evidence against the mechanism rather than merely an inability to estimate its structural parameters.

6. Conclusion

Dynamic pricing has traditionally been analysed as an optimisation problem under demand uncertainty, capacity constraints and intertemporal consumer heterogeneity. This paper identifies an additional dynamic force: consumers evaluate prices not only by what they pay, but also by how those prices compare with behavioural reference points formed through expectations and previous transactions. Forward-looking firms therefore optimise not only over current revenues, but also over the future behavioural environment that current prices help to create. Reference points are no longer simply exogenous constraints on pricing; they become endogenous state variables in the firm's intertemporal optimisation problem.

The analysis develops a tractable dynamic model in which expectation-based reference dependence, comparison-based inequity aversion, social identity and procedural legitimacy jointly influence equilibrium pricing. The two-period model yields three principal results. First, expectation-based reference dependence attenuates firms' responses to positive demand shocks, generating smoother equilibrium price paths than the conventional benchmark. Secondly, comparison-based inequity aversion compresses profitable price differences across cohorts, with the effect strengthened when social identity makes cross-buyer comparisons more salient. Thirdly, and most importantly, because current prices establish reference points against which subsequent prices are evaluated, forward-looking firms optimally elevate the initial price. The launch premium is therefore not a conventional scarcity or skimming effect: it arises because the initial price has strategic value in shaping future behavioural responses.

The mechanism extends beyond the two-period setting. When the perishable good is sold over multiple repricing epochs and the comparison benchmark is the lowest price paid by an earlier cohort, the launch premium remains strictly positive for any finite horizon and capacity. Numerical analysis indicates that the premium increases as repricing becomes more frequent, although the continuous-time limit is not established analytically. The distortion is bounded as comparison intensity increases and is non-monotonic in scarcity: it is largest at intermediate scarcity, because severe scarcity already induces high prices for rationing purposes, while abundant capacity leaves little future rent for an elevated comparison benchmark to protect. The value of maintaining a high benchmark also declines as the selling horizon closes, generating an endogenous downward price path without requiring an exogenously high initial reference price.

The calibration suggests that the mechanism can be economically significant. For empirically plausible behavioural parameters, cross-cohort price differences are compressed by approximately 80–95 per cent relative to the non-behavioural benchmark, while initial prices rise by approximately 5–10 per cent. The framework also provides a possible explanation for the markedly different reception of algorithmic dynamic pricing across markets. Dynamic pricing is

well established in sectors such as airlines and hotels but remains more contentious in ride-sharing, ticketing and live entertainment. The model suggests that these differences may reflect variation in the salience of cross-buyer comparisons, the strength of identification with the relevant comparison group, and the perceived legitimacy of the process generating price differences.

These mechanisms have implications for both firms and policymakers. For firms, increasingly sophisticated pricing algorithms need to account for the fact that current revenue maximisation can alter the behavioural constraints facing future prices. A pricing rule that ignores the reference points created by its own previous decisions can therefore be dynamically suboptimal even when it is optimal conditional on current demand and capacity. For policymakers, transparency interventions are not necessarily behaviourally neutral. Measures such as all-in pricing can increase the salience and comparability of transaction prices and thereby alter the behavioural environment over which firms optimise. Their effects may consequently extend beyond better-informed consumer choice to changes in the level and intertemporal structure of prices themselves.

Three limitations of the analysis suggest directions for further work. First, although the launch premium is analytically positive at every finite repricing frequency considered by the model, its behaviour as repricing approaches continuous time is established numerically rather than in closed form. Secondly, the comparison reference is inherited passively: consumers compare their treatment with earlier prices without inferring that those prices may themselves have been chosen strategically to influence subsequent comparisons. Allowing procedural legitimacy to respond endogenously to perceived strategic manipulation would provide an important extension and would naturally constrain the firm's incentive to elevate the initial benchmark. Thirdly, the model considers a monopolist. Extending the analysis to competition, where rival prices may enter the comparison set and firms may compete not only over current demand but also over reference-point formation, is perhaps the most consequential theoretical extension.

The central implication is therefore broader than the particular behavioural mechanisms studied here. Once reference formation is endogenous to the firm's intertemporal problem, dynamic pricing cannot be characterised solely as optimisation over economic state variables such as demand, inventories and capacity. Today's price affects both current revenue and the behavioural environment in which tomorrow's price will be evaluated. Theorem 1 captures this forward-looking distortion: the current price becomes a strategic instrument for shaping future behavioural responses. The contribution is therefore not simply to add behavioural preferences to dynamic pricing, but to show that endogenous reference formation changes the structure of the dynamic pricing problem itself.

Appendix A.1 Analytical Foundations

A.1.1. Lemma A.1 (Strict Concavity and Uniqueness)

Under Assumption 1, the firm's period-2 objective function is strictly concave in every interior behavioural regime considered in the paper. Consequently, whenever an interior behavioural solution exists, the first-order condition characterises the unique optimum.

Proof. Within any interior behavioural regime, the unified perceived price of a behaviourally sensitive consumer is affine in the realised price. Specifically,

$$\breve{P}_2 = P_2 + (\lambda - 1)\max\{0, P_2 - R_2\} + \alpha\max\{0, P_2 - P_1\},$$

has constant marginal slope

$$s \equiv \frac{\partial \breve{P}_2}{\partial P_2} \in \{1, \lambda, 1 + \alpha, \lambda + \alpha\},$$

according to whether neither, one, or both behavioural reference objects are exceeded. The population share π enters only through the weighting of behavioural and standard consumers and therefore does not affect the local slope within a given regime.

The firm's period-2 objective is

$$\Pi_2(P_2) = P_2[(1 - \pi)D(P_2) + \pi D(\breve{P}_2)].$$

Differentiating twice yields

$$\frac{\partial^2 \Pi_2}{\partial P_2^2} = (1 - \pi)[P_2 D''(P_2) + 2D'(P_2)] + \pi\left[s^2 P_2 D''(\breve{P}_2) + 2sD'(\breve{P}_2)\right].$$

By Assumption 1,

$$P\,D''(P) + 2D'(P) < 0$$

for every relevant price argument. Since $s > 0$ and $s^2 > 0$, both components of the second derivative are strictly negative. The objective is therefore a convex combination of two strictly concave revenue functions with non-negative weights $1 - \pi$ and π, implying

$$\frac{\partial^2 \Pi_2}{\partial P_2^2} < 0$$

throughout every interior behavioural regime. Hence the period-2 optimisation problem is strictly concave, the first-order condition is sufficient, and the interior optimum is unique. ■

A.1.2 Lemma A.2 (Characterisation of behavioural pricing regimes)

The period-2 optimisation problem admits exactly three mutually exclusive pricing regimes:

(i) a no-loss interior optimum;

(ii) an interior behavioural optimum above the reference price; and

(iii) a corner solution at the reference price.

The third regime corresponds to the complete behavioural smoothing region characterised in Corollary 1.

Proof. The reference-dependent period-2 objective is kinked at

$$P_2 = R_2.$$

Let

$$\Pi_-(P_2)$$

and

$$\Pi_+(P_2)$$

denote the objective evaluated immediately below and above the reference price. At the kink,

$$\Pi'_+(R_2) = \Pi'_-(R_2) - \kappa R_2[-D'(R_2)].$$

Since

$$\pi > 0, \kappa > 0, R_2 > 0, D'(R_2) < 0,$$

the right derivative is strictly smaller than the left derivative.

Exactly three cases are therefore possible.

(i) No-loss interior optimum. If

$$\Pi'_-(R_2) \leq 0,$$

the objective is decreasing at the reference price from below. The optimum lies weakly below R_2 and is characterised by the standard monopoly first-order condition.

(ii) Interior behavioural optimum. If

$$\Pi'_+(R_2) \geq 0,$$

the objective is increasing immediately above the reference price. The optimum lies strictly above R_2 and, by Lemma A.1, is unique.

(iii) Complete-smoothing corner. Otherwise,

$$\Pi'_-(R_2) \geq 0 \quad \text{and} \quad \Pi'_+(R_2) \leq 0.$$

The objective increases up to the reference price and decreases immediately afterwards. Hence the unique optimum is

$$P_2 = R_2.$$

This is the complete-smoothing region referred to in Corollary 1.

Under linear demand,

$$D(P_2) = a - bP_2,$$

the complete-smoothing region has width

$$\kappa b R_2,$$

or

$$\frac{\kappa b R_2}{a}$$

as a proportion of the baseline demand intercept. ■

Remark (unified two-kink taxonomy). When the comparison reference P_1 is also present, the period-2 objective has a second upward kink at $P_2 = P_1$, where the perceived-price slope rises by α. The same single-crossing argument applies independently at each kink: the marginal perceived slope $s = \partial \breve{P}_2 / \partial P_2$ is constant within each interval, taking the values $1, \lambda, 1+\alpha$ and $\lambda + \alpha$, and rises at every kink, so marginal revenue falls discontinuously there. The optimum therefore lies in exactly one region, slack, pooling at the lower reference, an interior region between the references, pooling at the upper reference, or the interior region above both, and within each interior region is unique by strict concavity. This is the regime taxonomy invoked in the proof of Proposition 1.

A.1.3 Lemma A.3 (Existence, Uniqueness and Stability of the Personal-Equilibrium Reference Point)

Maintain Assumption 1 and the shock specification of the model section, under which ε has a continuous, mean-zero distribution G that is atomless with full support on a compact interval $[-d, d]$. Following Kőszegi and Rabin (2006), the expectation reference (behavioural reference point in their terminology) is determined as a personal equilibrium. Let the behavioural reference point be the rational expectation of the equilibrium period-2 price conditional on the information available after the first-period pricing decision,

$$R_2 = E[P_2 \mid I_1]$$

and define the personal-equilibrium operator

$$\Psi(R_2) \equiv E[P_2^*(R_2, \varepsilon) \mid I_1]$$

where $P_2^*(R_2, \varepsilon)$ is the firm's optimal period-2 price when consumers hold reference point R_2 and ε is the period-2 shock, and let $\mathcal{R} \subset \mathbb{R}$ denote the closed, bounded interval of admissible reference points. Then the operator $\Psi: \mathcal{R} \to \mathcal{R}$ admits a unique fixed point $R_2 \in \mathcal{R}$,

$$R_2^* = \Psi(R_2^*)$$

and this fixed point is globally stable: for every initial belief $R_2^{(0)} \in \mathcal{R}$, the sequence defined by $R_2^{(n+1)} = \Psi(R_2^{(n)})$ converges to R_2^*, with geometric error bound $\left|R_2^{(n)} - R_2^*\right| \leq k^n \left|R_2^{(0)} - R_2^*\right|$ for some constant $k \in [0,1)$.

Proof. The argument proceeds in four steps.

Step 1 — Ψ *is single-valued*. By Lemma A.1, Assumption 1 renders the firm's period-2 objective strictly concave in P_2 for every realisation of ε and every $R_2 \in \mathcal{R}$. The maximiser $P_2^*(R_2, \varepsilon)$ is therefore unique for each pair (R_2, ε), and $\Psi(R_2) = E[P_2^*(R_2, \varepsilon) \mid I_1]$ is a well-defined, single-valued function on $\mathcal{R}$.

Step 2 — Ψ *maps the reference set into itself.* For each $R_2 \in \mathcal{R}$ the optimal price $P_2^*(R_2, \varepsilon)$ lies in the admissible price interval, which coincides with $\mathcal{R}$. Because $\mathcal{R}$ is convex, the conditional expectation $\Psi(R_2)$ of an $\mathcal{R}$-valued random variable also lies in $\mathcal{R}$; hence $\Psi: \mathcal{R} \to \mathcal{R}$.

Step 3 — Ψ is a contraction. By Proposition 1 the period-2 rule is piecewise linear and continuous in R_2 across five regimes. Write

- $N(R_2)$ for the no-loss event, on which $P_2^* = a_2/(2b)$;
- $M_R(R_2)$ for the event that the firm pools at the expectation reference, $P_2^* = R_2$;
- $L(R_2)$ for the interior expectation-loss event, on which $P_2^* = (a_2 + b\kappa R_2)/[2b(1+\kappa)]$;
- $M_P(R_2)$ for the event that the firm pools at the comparison reference, $P_2^* = P_1$;

- $B(R_2)$ for the interior both-bind event, on which $P_2^* = [a_2 + b(\kappa R_2 + \gamma P_1)]/[2b(1+\kappa+\gamma)]$.

Differentiating each branch,

$$\frac{\partial P_2^*}{\partial R_2} = \begin{cases} 0, & \text{on } N, \\ 1, & \text{on } M_R, \\ c_L, & \text{on } L, \\ 0, & \text{on } M_P, \\ c_B, & \text{on } B, \end{cases}$$

where

$$c_L := \frac{\kappa}{2(1+\kappa)} \qquad \text{and} \qquad c_B := \frac{\kappa}{2(1+\kappa+\gamma)},$$

both constant in ε and satisfying $0 < c_B < c_L < 1$ for $\kappa > 0$ and $\gamma > 0$. The two interior slopes differ because the comparison wedge adds γ to the effective demand slope wherever it binds, further attenuating the pass-through of the expectation reference. P_2^* is continuous across all four regime boundaries, and because G is atomless the boundaries carry zero probability, so Ψ is differentiable on $\mathcal{R}$ with

$$\begin{aligned} \Psi'(R_2) &= \Pr[M_R(R_2)] + c_L \Pr[L(R_2)] + c_B \Pr[B(R_2)] \\ &= 1 - \Pr[N(R_2)] - \Pr[M_P(R_2)] - (1-c_L)\Pr[L(R_2)] - (1-c_B)\Pr[B(R_2)], \end{aligned}$$

the second line following from $\Pr[N] + \Pr[M_R] + \Pr[L] + \Pr[M_P] + \Pr[B] = 1$. Every term is non-negative, so $0 \le \Psi'(R_2) \le 1$, with equality on the right only if M_R occurs almost surely.

That case is excluded. The pooling band at the expectation reference is the interval $(2bR_2, (2+\kappa)bR_2]$ in a_2, of width $b\kappa R_2$, whereas $a_2 = a_1 + \varepsilon$ has full support on an interval of width $2d$. Hence $\Pr[M_R(R_2)] < 1$ at every $R_2 \in \mathcal{R}$ whenever

$$2d > b\kappa \max\{R_2 : R_2 \in \mathcal{R}\},$$

so that the support of a_2 is never contained in the band; equivalently, at every candidate reference the firm strictly departs from it with positive probability. Under this condition $\Pr[N] + \Pr[L] + \Pr[M_P] + \Pr[B] > 0$ and therefore $\Psi'(R_2) < 1$ pointwise.

Since G admits a strictly positive density and the regime thresholds vary continuously in R_2, each of the five regime probabilities is continuous in R_2, and hence so is Ψ'. As $\mathcal{R}$ is compact, the supremum

$$k := \sup_{R_2 \in \mathcal{R}} \Psi'(R_2) < 1$$

is attained. By the mean value theorem, for all $R_2, R_2' \in \mathcal{R}$,

$$|\Psi(R_2) - \Psi(R_2')| \le k|R_2 - R_2'|,$$

so Ψ is a contraction on $\mathcal{R}$ with modulus k. ■

Step 4 — *Existence, uniqueness and stability.* The set $\mathcal{R}$ is a closed, bounded subset of $\mathbb{R}$ and hence a complete metric space under the Euclidean metric. By the Banach contraction-mapping theorem, Ψ possesses a unique fixed point $R_2 \in \mathcal{R}$ satisfying $R_2^* = \Psi(R_2^*)$. Moreover, the Picard iterates $R_2^{(n+1)} = \Psi(R_2^{(n)})$ converge to R_2^* from any starting point $R_2^{(0)} \in \mathcal{R}$, with the geometric error bound

$$\left|R_2^{(n)} - R_2^*\right| \leq k^n \left|R_2^{(0)} - R_2^*\right|$$

This establishes existence, uniqueness and global stability of the personal-equilibrium reference point under iterative updating of consumers' expectations. ■

Throughout the paper, all propositions are evaluated at this unique personal-equilibrium behavioural reference point, R_2^*.

Remark (numerical verification). In the disclosure calibration of Section 4.4 the contraction modulus, the derivative of T evaluated at the equilibrium reference, does not exceed 0.19 across the reported sectors, with no-loss probability roughly 0.36–0.40 at each fixed point; the personal equilibrium is therefore strongly stable, well inside the contraction region.

Remark (location of the fixed point). Because the firm pools at the reference over a band of positive width $b\kappa R_2^*$ in ε (Lemma A.2), P_2^* carries a mass point at R_2^*, which is therefore both a mode of the period-2 price and, by the personal-equilibrium condition, its mean. Cancelling this common pooling mass from $R_2^* = E[P_2^*]$ yields the equivalent characterisation $R_2 * = E[P_2^* \mid \text{the firm does not pool at } R_2^*]$: the unattenuated downside of the slack region ($P_2^* < R_2^*$) exactly offsets the attenuated upside of the loss, compression and both-bind regions ($P_2^* \geq R_2^*$). The fixed point therefore sits at the lower edge of the complete-smoothing band in benchmark terms ($P_2^0 = R_2^*$), with slack states below and attenuated states above; this is the sense in which the pooling height drawn in Figure 1 is consistent with $R_2^* = E[P_2^*]$, the offsetting no-loss mass being roughly 0.36–0.40 in the Section 4.4 calibration.

A.1.4 Lemma A.4 (Monotonicity of the Value Function in the Fairness Floor)

In the T-period model of Section 4, the firm's value function is weakly increasing in the prevailing price floor, for any epoch and any level of remaining inventory. It is strictly increasing whenever the optimal price exceeds the floor with positive probability and demand is interior.

Proof. Fix two floors, and let the second be strictly higher than the first. Let a policy be optimal at the lower floor. Construct a policy at the higher floor by mimicking: whenever the original policy posts a price, the constructed policy posts instead the nominal price that reproduces exactly the same behaviourally adjusted price, that is, the price at which the nominal price plus the fairness wedge measured against the higher floor equals the nominal price plus the wedge measured against the lower floor. Because the wedge is weakly smaller at the higher floor, the mimicking price is weakly above the original price, and strictly above it whenever the original price exceeds the lower floor.

Under the mimicking policy the behaviourally adjusted price is identical in every epoch and every state, so demand is identical, sales are identical, and the inventory carried forward is

identical. The capacity constraint therefore binds at exactly the same times and the entire inventory path is unchanged. The floor transition is the minimum of the prevailing floor and the price posted, and since both arguments are weakly higher under the mimicking policy, the transitioned floor is weakly higher, so the construction iterates forward. Revenue in each epoch is the nominal price multiplied by sales, and sales are unchanged while the nominal price is weakly higher, so revenue is weakly higher in every epoch, and strictly higher in any epoch in which the original price exceeded the lower floor with positive demand.

The mimicking policy is feasible at the higher floor but need not be optimal there, so the value at the higher floor weakly exceeds the value of the mimicking policy, which weakly exceeds the value of the optimal policy at the lower floor. The inequality is strict under the stated condition. ■

The lemma removes the only channel through which a higher floor might have reduced continuation value, namely faster depletion of perishable inventory. Because the mimicking policy reproduces the inventory path exactly while collecting weakly higher revenue at every date, no such trade-off arises, and the monotonicity does not depend on the salvage value or on the shadow price of capacity.

Convention (shock consistency and forecast timing). The demand shock process is specified at the finest resolution under consideration, and the epoch shock at resolution T is the average of its constituent finer shocks. Aggregate demand over any interval of the selling window is therefore invariant to the resolution at which the window is divided. At the start of each interval of length $1/T$ the firm observes the aggregate shock for that interval; within the interval it observes constituent shocks as they arrive. The comparison across T thus holds the firm's forecast of near-term aggregate demand fixed and varies only the frequency at which it may reprice.

A.1.5 Lemma A.5 (Refinement replication)

Divide the selling window into T epochs with cohort mass $1/T$ each, and let $W_T(m, I)$ denote the continuation value defined in Proposition 6. If T divides T', then $W_{T'}(m, I) \geq W_T(m, I)$ for every floor m and inventory I.

Proof. Let $T' = \ell T$ and partition the T' epochs into T consecutive blocks of ℓ epochs each. Given any feasible policy for the T-epoch problem, construct a policy for the T'-epoch problem that posts, throughout block j, the price the T-epoch policy posts in epoch j. We show this policy is feasible and attains the same expected value.

Feasibility. The T-epoch price in epoch j is measurable with respect to the aggregate shock $\bar{\varepsilon}_j$ together with the floor and inventory entering epoch j. By the forecast convention $\bar{\varepsilon}_j$ lies in the T'-firm's information set at the start of block j, and by the induction below the floor and inventory entering block j coincide with those entering epoch j. The replicating price is therefore measurable with respect to the T'-firm's information at the start of block j and constant thereafter, so the policy is adapted.

Cohort mass. Block j comprises ℓ epochs of mass $1/T'$ each, so the mass facing the replicating price is $\ell/T' = 1/T$, exactly the mass facing it in the T-epoch problem.

Floor. The floor is the running minimum of posted prices. Since the replicating price is constant within a block, the floor is unchanged within blocks. The behavioural wedge depends on the current price and the prevailing floor alone, and is therefore identical.

Demand, sales and revenue. By shock consistency, aggregate demand over block j at the common price is the same random variable as demand in epoch j of the T-epoch problem. Sales are the lesser of realised demand and remaining inventory; since demand is non-negative and the price is constant within the block, sequential rationing across the ℓ epochs yields cumulative block sales equal to the lesser of aggregate block demand and inventory entering the block—which is precisely epoch-j sales in the T-epoch problem. Revenue over the block is therefore identical realisation by realisation, as is inventory carried forward.

Induction. The floor and inventory entering block $j+1$ thus coincide with those entering epoch $j+1$. Taking as the induction hypothesis that they coincide at the start of block j, which holds at $j=1$ by construction, the argument above establishes it at $j+1$. The two paths coincide throughout, and the replicating policy attains the same expected value as the T-epoch policy.

Hence every T-epoch policy is attainable in the T'-epoch problem at equal expected value. Since the T'-firm may in addition vary its price within blocks and condition on constituent shocks as they arrive, taking the maximum over policies gives $W_{T'}(m,I) \geq W_T(m,I)$, with strict inequality wherever some within-block price variation is strictly profitable.

Appendix A.2 Proofs

Unless otherwise stated, all proofs invoke Assumption 1 and Lemma A.1 to ensure strict concavity and uniqueness of the relevant optimisation problem.

A.2.1 Proof of Lemma 1 (Kink decoupling)

Write $w(P) \equiv \kappa \mathbf{1}\{P > R_2\} + \gamma \mathbf{1}\{P > P_1\}$. Under linear demand, aggregate second-period demand is

$$D_2(P) = a_2 - b\big[\big(1 + w(P)\big)P - \kappa R_2 \mathbf{1}\{P > R_2\} - \gamma P_1 \mathbf{1}\{P > P_1\}\big],$$

so profit $\Pi(P) = PD_2(P)$ is continuous, and on each of the three intervals (two, in the non-generic case $R_2 = P_1$) delimited by R_2 and P_1 it is quadratic with $\Pi''(P) = -2b(1 + w) < 0$. Hence Π is strictly concave on each piece.

Consider the breakpoint $P = R_2$. Immediately below it the expectation term is inactive and marginal profit is $a_2 + b\gamma P_1 \mathbf{1}\{R_2 > P_1\} - 2b(1 + w^-)R_2$, where $w^- = \gamma \mathbf{1}\{R_2 > P_1\}$; immediately above it the same expression obtains with κ added to w^- and $b\kappa R_2$ added to the intercept. The difference between the right- and left-derivatives is therefore

$$b\kappa R_2 - 2b\kappa R_2 = -b\kappa R_2,$$

which is strictly negative and independent of the indicator $\mathbf{1}\{R_2 > P_1\}$, hence of whether the comparison term is active and of the ordering of the two references. The identical computation at $P = P_1$ gives a jump of $-b\gamma P_1$. This proves the jump sizes and their independence.

For (i): a continuous function that is strictly concave on each piece and whose derivative jumps downward at every breakpoint has a strictly decreasing derivative correspondence on the whole line, and is therefore globally strictly concave. Existence and uniqueness of the optimum follow, without recourse to the regime taxonomy of Lemma A.2.

For (ii): by global strict concavity the firm prices exactly at a breakpoint r if and only if the left-derivative at r is non-negative and the right-derivative non-positive. Both derivatives are affine in a_2 with unit coefficient, so the set of a_2 satisfying both is a closed interval whose length equals the size of the derivative jump at r, namely $b\kappa R_2$ at $r = R_2$ and $b\gamma P_1$ at $r = P_1$. ■

A.2.2. Proof of Proposition 1 (Unified Period-2 Pricing)

Behaviourally sensitive consumers perceive $\breve{P}_2 = P_2 + (\lambda - 1)\max\{0, P_2 - R_2\} + \alpha \cdot \max\{0, P_2 - P_1\}$, with marginal value $s \equiv \partial \breve{P}_2 / \partial P_2 \in \{1, \lambda, 1 + \alpha, \lambda + \alpha\}$. The firm's first-order condition is $(1 - \pi)[D(P_2) + P_2 D'(P_2)] + \pi[D(\breve{P}_2) + s \cdot P_2 D'(\breve{P}_2)] = 0$.

Existence and uniqueness. Lemma A.1 establishes strict concavity of the objective on each interior regime under Assumption 1. By Lemma 1 marginal profit falls discontinuously at each breakpoint by $\pi\kappa \cdot (-D')R_2$ at $P_2 = R_2$ and by $\pi\gamma \cdot (-D')P_1$ at $P_2 = P_1$, which under linear demand are $b\kappa R_2$ and $b\gamma P_1$ respectively. A continuous function that is strictly concave on each piece and whose derivative jumps downward at every breakpoint has a strictly decreasing

derivative correspondence on the whole line, and is therefore globally strictly concave. The optimum exists and is unique, and is interior because the objective vanishes at $P_2 = 0$ and at the choke price.

Regime boundaries. Equating the one-sided derivatives to zero at each breakpoint and inverting for a_2 delivers the boundaries of Table 4, the pooling intervals having widths equal to the respective jumps. The argument nowhere uses the relative position of R_2 and P_1, so the taxonomy holds for either ordering.

The terminal regime. Under linear demand, where both references bind ($\check{P}_2 = (\lambda + \alpha)P_2 - (\lambda - 1)R_2 - \alpha P_1$ and the first-order condition solves to $P_2^* = [a_2 + b(\kappa R_2 + \gamma P_1)]/[2b(1 + \kappa + \gamma)]$. Differentiating with respect to a_2 gives the attenuation factor $1/(1 + \kappa + \gamma)$ of part (a); setting $\gamma = 0$ gives $\frac{1}{1+\gamma}$. Pass-through is zero on the pooling intervals because the price is pinned to a breakpoint independent of a_2.

Part (b). Differentiating $\Delta^* = P_2^* - P_1$ with respect to γ yields (1). Its sign follows from the regime condition: the terminal regime requires $a_2 > \bar{A} + b\gamma P_1 = 2b(1 + \kappa)P_1 - b\kappa R_2 + b\gamma P_1$, whence $a_2 + b\kappa R_2 > 2b(1 + \kappa)P_1 > b(1 + \kappa)P_1$,so the numerator of (1) is strictly negative.

Part (c). Setting $\Delta^* = 0$ and solving for γ gives the threshold $\gamma^*(\kappa)$ of the statement. Since $\gamma^*(\kappa)$ does not depend on γ, the condition $\gamma \geq \gamma^*(\kappa)$ is preserved under increases in γ, so complete compression once attained is not reversed. Holding γ fixed and solving for a_2, gives an interval of width $b\gamma P_1$, which is the pooling interval at P_1 identified in Lemma 1(ii); the corresponding interval at R_2 has width $b\kappa R_2$.

Taking expectations of the period-2 rule over ε defines the personal-equilibrium operator $\Psi(R_2) = E[P_2^*(R_2, P_1, \varepsilon)]$; its fixed point $R_2^* = \Psi(R_2^*)$, together with existence, uniqueness and global stability, is established in Lemma A.3. ■

A.2.3 Proof of Corollary 1 (Reference-Dependent Price Smoothing)

In the behavioural loss region, $P_2 > R_2^*$, the perceived price for behavioural consumers is

$$\hat{P}_2 = P_2 + (\lambda - 1)(P_2 - R_2^*).$$

Aggregate behaviourally adjusted demand is therefore

$$\widetilde{D}_2(P_2) = (1 - \pi)D(P_2) + \pi D(\hat{P}_2).$$

The firm chooses P_2 to maximise

$$P_2\widetilde{D}_2(P_2).$$

The first-order condition in the behavioural loss region is

$$(1 - \pi)\{D(P_2) + P_2D'(P_2)\} + \pi\{D(\hat{P}_2) + \lambda P_2D'(\hat{P}_2)\} = 0. \qquad (A.1)$$

Let $P_2^0(\varepsilon)$ denote the benchmark optimum without behavioural effects. Suppose, for contradiction, that

$$P_2^*(\varepsilon) > P_2^0(\varepsilon).$$

By Assumption 1, the benchmark revenue function without behavioural effects is strictly concave, so

$$D(P_2^*) + P_2^* D'(P_2^*) < 0.$$

Moreover, in the behavioural loss region,

$$\hat{P}_2 > P_2^*,$$

and the behaviourally sensitive consumers marginal-revenue term is strictly below the corresponding benchmark marginal-revenue term without behavioural effects. Hence both terms in (A.1) are non-positive, with at least one strictly negative. This contradicts the first-order condition. Therefore

$$P_2^*(\varepsilon) \leq P_2^0(\varepsilon),$$

with strict inequality whenever

$$P_2^0(\varepsilon) > R_2^*.$$

This proves part (a).

For linear demand,

$$D(P_2) = a + \varepsilon - bP_2,$$

the first-order condition becomes

$$(1-\pi)(a + \varepsilon - 2bP_2) + \pi\{a + \varepsilon - b\hat{P}_2 - \lambda bP_2\} = 0.$$

Using

$$\hat{P}_2 = \lambda P_2 - (\lambda - 1)R_2^*,$$

this simplifies to

$$a + \varepsilon - 2b\{1 + \pi(\lambda - 1)\}P_2 + \pi b(\lambda - 1)R_2^* = 0.$$

Hence

$$P_2^*(\varepsilon) = \frac{a + \varepsilon + \pi b(\lambda - 1)R_2^*}{2b\{1 + \pi(\lambda - 1)\}}.$$

The benchmark optimum is

$$P_2^0(\varepsilon) = \frac{a + \varepsilon}{2b}.$$

Therefore

$$\frac{\partial P_2^*}{\partial \varepsilon} = \frac{1}{2b\{1 + \pi(\lambda - 1)\}},$$

while

$$\frac{\partial P_2^0}{\partial \varepsilon} = \frac{1}{2b}.$$

The attenuation factor is therefore

$$\frac{\partial P_2^* / \partial \varepsilon}{\partial P_2^0 / \partial \varepsilon} = \frac{1}{1 + \pi(\lambda - 1)} = \frac{1}{1 + \kappa}.$$

This proves part (b).

Finally, $P_2^*(\varepsilon)$ coincides with $P_2^0(\varepsilon)$ in the no-loss region, equals R_2^* in the complete-smoothing region, and rises less steeply than $P_2^0(\varepsilon)$ in the behavioural loss region. Hence the behavioural pricing rule is a compression of the benchmark pricing rule. Therefore

$$\mathrm{Var}_\varepsilon(P_2^*) < \mathrm{Var}_\varepsilon(P_2^0)$$

whenever the demand-shock distribution assigns positive probability to outcomes for which the benchmark optimum lies in the behavioural loss region. ■

A.2.4 Proof of Corollaries 2 and 3 (Comparison-Based Compression and Complete Compression)

Set $\kappa = 0$ throughout. For $P_2 > P_1$ the behaviourally sensitive consumer's perceived price is $\tilde{P}_2 = (1 + \alpha)P_2 - \alpha P_1 > P_2$, aggregate demand is $D_2(P_2; P_1) = \pi D(\tilde{P}_2) + (1 - \pi)D(P_2)$, and the firm maximises $\Pi_2(P_2) = P_2 D_2(P_2; P_1)$. The first-order condition is

$$\Phi(P_2;\alpha,\pi) \equiv \left[\pi D\left(\tilde{P}_2\right) + (1-\pi)D(P_2)\right] + P_2\left[\pi(1+\alpha)D'\left(\tilde{P}_2\right) + (1-\pi)D'(P_2)\right] = 0. \qquad (A.2)$$

By Lemma A.1 the objective is strictly concave in the dispersion region, so $\partial\Phi/\,\partial P_2 < 0$ and (A.2) characterises the unique interior optimum whenever the dispersion-region programme applies.

Step 1: the upper bound. Let P_2^0 satisfy $D(P_2^0) + P_2^0 D'(P_2^0) = 0$; by strict concavity of benchmark revenue, $D(P) + PD'(P) < 0$ for all $P > P_2^0$. Suppose for contradiction $P_2^* \geq P_2^0$. The behaviourally neutral term is then non-positive. For the behavioural term, $\tilde{P}_2^* = P_2^* + \alpha(P_2^* - P_1) > P_2^* \geq P_2^0$, and $(1+\alpha)P_2^* - \tilde{P}_2^* = \alpha P_1 > 0$, so with $D'\left(\tilde{P}_2^*\right) < 0$,

$$D\left(\tilde{P}_2^*\right) + (1+\alpha)P_2^* D'\left(\tilde{P}_2^*\right) < D\left(\tilde{P}_2^*\right) + \tilde{P}_2^* D'\left(\tilde{P}_2^*\right) < 0,$$

the last inequality because $\tilde{P}_2^* > P_2^0$. Hence $\Phi(P_2^*;\alpha,\pi) < 0$, contradicting (A.2). Therefore $P_2^* < P_2^0$ and, P_1 being fixed, $\Delta^* < P_2^0 - P_1$.

Step 2: the lower bound and the threshold. Positivity of Δ^* does not follow from $P_2^0 > P_1$. By Lemma 1 the firm prices exactly at P_1 whenever the left-derivative of profit there is non-negative and the right-derivative non-positive. Evaluating (A.2) at $P_2 = P_1$, where $\tilde{P}_2 = P_1$, gives the right-derivative

$$\Phi(P_1;\alpha,\pi) = D(P_1) + (1+\gamma)P_1 D'(P_1), \qquad \gamma = \pi\alpha,$$

while the left-derivative is $D(P_1) + P_1 D'(P_1)$. Hence $\Delta^* > 0$ if and only if $\Phi(P_1;\alpha,\pi) > 0$, which under linear demand reads $a_2 > bP_1(2+\gamma)$. This is strictly stronger than the benchmark condition $a_2 > 2bP_1$; on the intermediate set $2bP_1 < a_2 \leq bP_1(2+\gamma)$ the left-derivative is non-negative and the right-derivative non-positive, so $P_2^* = P_1$ exactly and $\Delta^* = 0$. This establishes part (a) of Corollary 2 and part (b) of Corollary 3.

Moreover $\partial\Phi(P_1)/\,\partial\gamma = P_1 D'(P_1) < 0$, so $\Phi(P_1;\cdot)$ is strictly decreasing in γ and crosses zero at most once. The crossing point is the threshold $\gamma^*(0)$; under linear demand, solving $a_2 - bP_1(2+\gamma) = 0$ gives

$$\gamma^*(0) = \frac{a_2 - 2bP_1}{bP_1},$$

and $\Phi(P_1) > 0$ if and only if $\gamma < \gamma^*(0)$. Complete compression therefore holds on the half-line $[\gamma^*(0), \infty)$ rather than at a single point, and by Lemma 1(ii) it holds, for fixed γ, on an interval of intercepts of width $b\gamma P_1$. This establishes Corollary 3(b) and (c).

Step 3: monotonicity on the interior. Suppose $\gamma < \gamma^*(0)$, so the optimum is interior. By the implicit function theorem $\partial P_2^*/\,\partial\alpha = -(\partial\Phi/\,\partial\alpha)/(\partial\Phi/\,\partial P_2)$ with negative denominator. Using $\partial\tilde{P}_2/\,\partial\alpha = P_2 - P_1 > 0$,

$$\frac{\partial\Phi}{\partial\alpha} = \pi D'\left(\tilde{P}_2\right)(P_2 - P_1) + \pi P_2 D'\left(\tilde{P}_2\right) + \pi(1+\alpha)P_2 D''\left(\tilde{P}_2\right)(P_2 - P_1).$$

The first two terms are strictly negative because $D' < 0$; the third is non-positive under the maintained curvature condition. Hence $\partial \Phi / \partial \alpha < 0$ and $\partial P_2^* / \partial \alpha < 0$. Similarly

$$\frac{\partial \Phi}{\partial \pi} = \left[D\left(\tilde{P}_2\right) - D(P_2)\right] + P_2\left[(1+\alpha)D'\left(\tilde{P}_2\right) - D'(P_2)\right] < 0,$$

the first bracket being negative since $\tilde{P}_2 > P_2$ and the second by the marginal-revenue comparison of Step 1. Since P_1 is fixed, $\partial \Delta^* / \partial \alpha < 0$ and $\partial \Delta^* / \partial \pi < 0$, so Δ^* is strictly decreasing in $\gamma = \pi\alpha$. This establishes Corollary 2(b) and Corollary 3(a).

Step 4: linear demand. With $D(P) = a_2 - bP$, aggregate demand is $D_2 = a_2 - bP_2(1+\gamma) + b\gamma P_1$ and (A.2) gives $P_2^*(\gamma) = [a_2 + \gamma b P_1]/[2b(1+\gamma)]$, hence

$$\Delta^*(\gamma) = \frac{a_2 - bP_1(2+\gamma)}{2b(1+\gamma)}, \qquad \frac{\partial \Delta^*}{\partial \gamma} = \frac{bP_1 - a_2}{2b(1+\gamma)^2}.$$

In the interior region $a_2 > bP_1(2+\gamma) \geq bP_1$, so the derivative is strictly negative; differentiating through $\gamma = \pi\alpha$ gives $\partial \Delta^* / \partial \alpha = \pi(bP_1 - a_2)/[2b(1+\gamma)^2] < 0$ and $\partial \Delta^* / \partial \pi = \alpha(bP_1 - a_2)/[2b(1+\gamma)^2] < 0$. At $\gamma = 0$, $\Delta^* = P_2^0 - P_1$, confirming the bound of Step 1. Substituting $\gamma = \gamma^*(0)$ gives $\Delta^* = 0$, confirming Step 2. ■

Notation note. The existing proof of Corollary 3 uses $\hat{P}_2$ for the comparison-adjusted perceived price. Section 3.2.1 reserves $\hat{P}_2$ for the expectation-adjusted price and $\tilde{P}_2$ for the comparison-adjusted price. The merged proof above uses $\tilde{P}_2$ throughout.

A.2.5 Proof of Proposition 2 (Identity Amplification)

Identity salience affects the period-2 optimisation problem only through the effective inequity parameter

$$\alpha_{eff} = \alpha(1 + \sigma\varphi)(1 - J), \qquad \text{so that } \Gamma = \pi\alpha_{eff} = \gamma(1 + \sigma\varphi)(1 - J).$$

Conditional on α_{eff}, the optimisation problem is identical to that analysed in Corollary 3. Hence

$$\Delta^* = \Delta^*(\alpha_{\text{eff}})$$

with

$$\frac{\partial \Delta^*}{\partial \alpha_{\text{eff}}} < 0$$

whenever the dispersion region is reached.

Similarly,

$$\frac{\partial \Delta^*}{\partial \sigma} = \frac{\partial \Delta^*}{\partial \alpha_{\text{eff}}} \frac{\partial \alpha_{\text{eff}}}{\partial \sigma} = \frac{\partial \Delta^*}{\partial \alpha_{\text{eff}}} \alpha\varphi,$$

so

$$\frac{\partial \Delta^*}{\partial \sigma} \leq 0,$$

with strict inequality whenever

$$\alpha > 0,\ \varphi > 0,\ \pi > 0,$$

and the dispersion region is reached.

Finally, the composite intensity is $\Gamma = \gamma(1 + \sigma\varphi)(1 - J)$. By Corollary 3, Δ^* is strictly decreasing in Γ for $\Gamma < \Gamma^*(\kappa)$ and identically zero for $\Gamma \geq \Gamma^*(\kappa)$. The chain-rule derivatives $\partial\Delta^*/\partial\varphi = (\partial\Delta^*/\partial\Gamma)\gamma\sigma(1 - J)$ and $\partial\Delta^*/\partial\sigma = (\partial\Delta^*/\partial\Gamma)\gamma\varphi(1 - J)$ are therefore strictly negative when $\Gamma < \Gamma^*(\kappa)$ and zero otherwise, which is the strictness qualification of parts (a) and (b). For part (c), $\sigma_A\varphi_A > \sigma_B\varphi_B$ with common π, α, J implies $\Gamma_A > \Gamma_B$; if $\Gamma_A < \Gamma^*(\kappa)$ then both sectors are interior and $\Delta_A^* < \Delta_B^*$, whereas if $\Gamma_B \geq \Gamma^*(\kappa)$ both are at zero and the comparison is degenerate. ■

A.2.6 Proof of Proposition 3 (Procedural Legitimacy)

By construction,

$$\Gamma = \gamma(1 + \sigma\varphi)(1 - J),$$

so

$$\frac{\partial \Gamma}{\partial J} = -\gamma(1 + \sigma\varphi) < 0.$$

Applying the chain rule, $\partial\Delta^*/\partial J = (\partial\Delta^*/\partial\Gamma)(\partial\Gamma/\partial J)$. For $\Gamma < \Gamma^*(\kappa)$ the first factor is strictly negative by Corollary 3(a) and the second is strictly negative, so $\partial\Delta^*/\partial J > 0$, establishing (a). For $\Gamma \geq \Gamma^*(\kappa)$ Corollary 3(b) gives $\Delta^* = 0$ identically, so $\partial\Delta^*/\partial\Gamma = 0$ and $\partial\Delta^*/\partial J = 0$, establishing (b).

For (c), $J \to 1$ implies $\Gamma \to 0$, so the comparison kink becomes inert and the period-2 problem reduces to the expectation-only problem of Corollary 1. Its solution in the region where the expectation reference binds is $P_2^* = [a_2 + b\kappa R_2]/[2b(1 + \kappa)]$, so $\Delta^* \to \Delta^\kappa \equiv [a_2 + b\kappa R_2]/[2b(1 + \kappa)] - P_1$. Since

$$\Delta^0 - \Delta^\kappa = \frac{a_2}{2b} - \frac{a_2 + b\kappa R_2}{2b(1 + \kappa)} = \frac{a_2(1 + \kappa) - a_2 - b\kappa R_2}{2b(1 + \kappa)} = \frac{\kappa(a_2 - bR_2)}{2b(1 + \kappa)} > 0$$

$$\text{whenever } \kappa > 0 \text{ and } a_2 > bR_2,$$

the last condition holding throughout the region in which the expectation reference binds, where $a_2 > 2bR_2$.complete procedural legitimacy restores the comparison-free level of intertemporal price differentiation but not the unconstrained benchmark, the residual gap being the expectation distortion. The two coincide if and only if $\kappa = 0$. ■

A.2.7 Proof of Theorem 1 (Forward-Looking Behavioural Pricing)

The result follows from the effect of the first-period price on the behavioural reference point used in period-2 inequity comparisons. The forward-looking loss-aversion reference R_2^* is determined in personal equilibrium and does not vary directly with the realised first-period price. By contrast, the inequity benchmark is the observed first-period price P_1. Hence increasing P_1 raises the comparison point against which period-2 prices are evaluated.

Let the period-2 value function be

$$V_2(P_1, \varepsilon) = \max_{P_2} \; P_2\big[\pi D(\breve{P}_2) + (1-\pi)D(P_2)\big],$$

where

$$\breve{P}_2 = P_2 + (\lambda - 1)(P_2 - R_2) + \alpha(P_2 - P_1),$$

or explicitly inside the dispersion region,

$$\breve{P}_2 = (\lambda + \alpha)P_2 - (\lambda - 1)R_2 - \alpha P_1.$$

Within the dispersion region both behavioural reference objects are active. Because the expectation reference is pinned down by personal equilibrium, variations in the first-period price affect the continuation value only through the comparison component of $\breve{P}_2$, so $\partial \breve{P}_2 / \partial P_1 = -\alpha$.

By the envelope theorem,

$$\frac{\partial V_2}{\partial P_1} = \pi P_2^* D'(\breve{P}_2)\frac{\partial \breve{P}_2}{\partial P_1} = -\pi\alpha P_2^* D'(\breve{P}_2).$$

Since

$$D'(\breve{P}_2) < 0,$$

we have

$$\frac{\partial V_2}{\partial P_1} > 0$$

whenever the dispersion region is reached. Raising P_1 therefore relaxes the future inequity constraint and increases the continuation value.

The firm's period-1 objective is

$$P_1 D(P_1) + \beta \mathbb{E}_\varepsilon[V_2(P_1, \varepsilon)].$$

Differentiating the period-1 objective and taking the total derivative of the continuation value with respect to P_1 gives

$$D(P_1) + P_1 D'(P_1) + \beta\, \mathbb{E}_\varepsilon\big[b\gamma P_2^*(\varepsilon)\mathbf{1}_{disp} + (a_2 - 2bP_1)\mathbf{1}_{pool}\big] = 0, \qquad (\mathrm{A}.3)$$

where $\mathbf{1}_{disp}$ indicates states in which the period-2 optimum lies strictly above the comparison reference, and $\mathbf{1}_{pool}$ states in which it lies exactly at it, that is $\bar{A} < a_2 \leq \bar{A} + b\gamma P_1$ in the notation of Table 4. The two terms are the two channels through which the first-period price relaxes the period-2 behavioural constraint. In the dispersion region the envelope theorem applies and the effect operates through the wedge, $\partial \check{P}_2 / \partial P_1 = -\alpha$. In the pooling region the period-2 optimum is pinned to the kink and moves one for one with P_1, so the envelope theorem does not apply and the derivative must be taken directly; period-2 revenue is then $P_1 D(P_1)$ and its derivative $a_2 - 2P_1$, strictly positive throughout the pooling region since $a_2 > 2bP_1$ there. Both terms are non-negative and at least one is strictly positive whenever $\gamma > 0$ and either region has positive probability. The final term is strictly positive whenever

$$\beta > 0, \qquad \pi\alpha > 0,$$

and the period-2 optimisation reaches the dispersion region with positive probability.

The static monopoly benchmark P_1^B satisfies

$$D(P_1^B) + P_1^B D'(P_1^B) = 0.$$

Since the additional continuation-value term in $(\mathrm{A}.3)$ is strictly positive, the period-1 first-order condition can be restored only at a higher price. Strict concavity of period-1 revenue therefore implies

$$P_1^* > P_1^B.$$

This establishes the first claim.

Non-monotonicity in γ. From the linear solution,

$$P_1^* - P_1^B = \frac{\beta\gamma}{2}\,\mathbb{E}_\varepsilon\big[P_2^*(\varepsilon)\mathbf{1}_{disp}\big],$$

three distinct effects of γ operate. The composite intensity enters multiplicatively, raising the wedge. Against this, $P_2^*(\varepsilon)$ is strictly decreasing in γ on the dispersion region by Proposition 1(ii); and $\mathbf{1}_{disp}$ is non-increasing in γ, since by Proposition 1(iii) the dispersion region is $\{a_2 > \bar{A} + b\gamma P_1\}$, whose lower boundary is strictly increasing in γ at rate bP_1.

At $\gamma = 0$ the comparison channel is inert, $\partial\breve{P}_2/\,\partial P_1 = 0$, and the wedge is zero.

Consider next the behaviour of the wedge as γ grows. The lower boundary of the dispersion region,

$$\bar{A} + b\gamma P_1 = 2b(1+\kappa)P_1 - b\kappa R_2 + b\gamma P_1,$$

is strictly increasing in both γ and P_1, so the dispersion region shrinks in each. The equilibrium expectation reference satisfies $R_2^* = \mathbb{E}[P_2^*] \le \mathbb{E}[P_2^0] = a_1/(2b) = P_1^B$, by Lemma A.3 and Corollary 1(a). Hence at any candidate $P_1 \ge P_1^B$ the boundary is bounded below by $bP_1(2+\kappa+\gamma)$, and since ε has compact support $[-d, d]$ the dispersion region is empty whenever

$$\gamma \ge \bar{\gamma}(P_1) \equiv \max\left\{0,\ \frac{a_1 + d - b(2+\kappa)P_1}{bP_1}\right\}.$$

Because $\bar{\gamma}$ is strictly decreasing in P_1 and P_1^* is non-decreasing in γ, the equation $\gamma = \bar{\gamma}\big(P_1^*(\gamma)\big)$ has a unique solution $\hat{\gamma}$.

For $\gamma \ge \hat{\gamma}$ the first term of (A.3) vanishes identically, and the first-order condition reduces to

$$D(P_1) + P_1 D'(P_1) + \beta\,\mathbb{E}_\varepsilon[(a_2 - 2P_1)\mathbf{1}\{a_2 > 2bP_1\}] = 0,$$

which does not involve γ. The launch price is therefore constant on $[\hat{\gamma}, \infty)$, and strictly above P_1^B because the remaining expectation is strictly positive whenever $a_1 + d > 2bP_1^B$. For $0 < \gamma < \hat{\gamma}$ both terms of (A.3) are present and non-negative, and the first is strictly positive because the dispersion region has positive probability by construction of $\hat{\gamma}$. The wedge is therefore strictly positive on $(0, \hat{\gamma})$. Its monotonicity there is not established analytically: γ enters the first term multiplicatively, which raises the wedge, but P_2^* is strictly decreasing in γ on the dispersion region by Proposition 1(ii) and the region itself shrinks at rate bP_1, so the product is not signed without further restriction. The wedge is strictly increasing in γ at every parameter combination reported in Table 5 and in panel (iii) of Table 8, and no case in which it declines was encountered over the ranges examined in Appendix A.4. Since the plateau above $\hat{\gamma}$ and the

strict positivity below it are what the empirical content of part (iii) requires, the analytical statement is confined to those. This establishes part (iii). ■

Economically, a higher comparison intensity makes an elevated benchmark more valuable per state in which it binds, but simultaneously reduces the number of such states and the price charged in them. The first effect dominates numerically while the dispersion region remains non-empty; once it is exhausted, neither effect operates and the premium is sustained entirely by the level at which the firm pools.

Monotonicity in β. Write the period-1 objective as

$$\Pi(P_1, \beta) = P_1 D(P_1) + \beta \mathbb{E}_\varepsilon[V_2(P_1, \varepsilon)].$$

Its cross-partial derivative is

$$\frac{\partial^2 \Pi}{\partial P_1 \partial \beta} = \mathbb{E}_\varepsilon[b\gamma P_2^*(\varepsilon)\mathbf{1}_{disp} + (a_2 - 2bP_1)\mathbf{1}_{pool}],$$

which is the continuation term of (A.4) divided by β. Both summands are non-negative, and at least one is strictly positive for every $\gamma > 0$: below $\hat{\gamma}$ the dispersion region has positive probability, and at every $\gamma > 0$ the pooling region has positive probability whenever $a_1 + d > 2bP_1^B$. The objective therefore exhibits strictly increasing differences in (P_1, β).

By Topkis's monotone comparative statics theorem (Topkis, 1978), the set of period-1 optima is non-decreasing in β in the strong set order. This requires no concavity of Π in P_1 and is therefore valid notwithstanding Remark 4. Where the optimum is unique and interior, the second-order condition $\Pi_{P_1P_1} < 0$ holds at that point, and implicit differentiation of the first-order condition gives

$$\frac{dP_1^*}{d\beta} = -\frac{\Pi_{P_1\beta}}{\Pi_{P_1P_1}} > 0.$$

Since P_1^B is independent of β, the wedge inherits this monotonicity. The argument applies on the whole range $\gamma > 0$, including $\gamma \geq \hat{\gamma}$, where only the pooling term is present but is likewise proportional to β. This establishes part (ii). ■

A.2.8 Proof of Proposition 4 (Reference Ordering)

Fix the comparison reference at the equilibrium first-period price P_1^* and recall the personal-equilibrium operator $\Psi(R_2) = \mathbb{E}[P_2^*(R_2, P_1^*, \varepsilon)]$ of Lemma A.3, whose unique fixed point is R_2^*. The argument proceeds in five steps.

Step 1 (single crossing). By Lemma A.3, $\Psi'(R_2) < 1$ at every admissible reference, so the map $g(R_2) := \Psi(R_2) - R_2$ is strictly decreasing. Since $g(R_2^*) = 0$, for any admissible price x we have $g(x) < 0$ if and only if $x > R_2^*$. Evaluating at $x = P_1^*$ gives

$$P_1^* > R_2^* \quad \Leftrightarrow \quad \Psi(P_1^*) < P_1^*. \tag{P4.1}$$

It therefore suffices to sign $\Psi(P_1^*) - P_1^*$. Only the contraction bound $\Psi' < 1$ is used here, not monotonicity of Ψ.

Step 2 (coincident kinks). Evaluating Ψ at $R_2 = P_1^*$ places both behavioural kinks of the period-2 objective at the same point P_1^*: the expectation reference and the comparison reference coincide. Both enter the perceived price only through the upward terms $(\lambda - 1) \cdot max\{0, P_2 - P_1^*\}$ and $\alpha \cdot max\{0, P_2 - P_1^*\}$, so the behavioural adjustment weakly raises the perceived price when $P_2 > P_1^*$ and is inert otherwise.

Step 3 (state-by-state compression). Let $P_2^0(\varepsilon)$ denote the benchmark period-2 optimum without behavioural effects. For every ε,

$$P_2^*(P_1^*, P_1^*, \varepsilon) \leq P_2^0(\varepsilon), \tag{P4.2}$$

with strict inequality whenever $P_2^0(\varepsilon) > P_1^*$. If $P_2^0(\varepsilon) \leq P_1^*$ neither kink binds and $P_2^* = P_2^0$; if $P_2^0(\varepsilon) > P_1^*$ the upward kinks lower behaviourally sensitive demand above P_1^*, and hence marginal revenue there, so by strict concavity (Lemma A.1) the constrained optimum lies strictly below $P_2^0(\varepsilon)$, either pooling at P_1^* or in the attenuated interior region. Under linear demand $D(P) = a_2 - bP$ the both-bind optimum of Proposition 1 gives

$$P_2^*(P_1^*, P_1^*, \varepsilon) = \frac{[a_2 + b(\kappa + \gamma)P_1^*]}{[2b(1 + \kappa + \gamma)]} \leq \frac{a_2}{2b} = P_2^0(\varepsilon) \quad \Leftrightarrow \quad bP_1^* \leq a_2,$$

which holds throughout the dispersion region.

Step 4 (expectations). Integrating (P4.2) over ε,

$$\Psi(P_1^*) = \mathbb{E}[P_2^*(P_1^*, P_1^*, \varepsilon)] \leq \mathbb{E}[P_2^0(\varepsilon)].$$

Under linear demand with $\mathbb{E}[\varepsilon] = 0$ and $a_2 = a_1 + \varepsilon$, the benchmark optimum is $P_2^0(\varepsilon) = a_2/(2b)$, so $\mathbb{E}[P_2^0(\varepsilon)] = a_1/(2b) = P_1^B$, the static monopoly benchmark.

Step 5 (conclusion). By Theorem 1, $P_1^B < P_1^*$ whenever $\beta > 0, \gamma = \pi\alpha > 0$ and the dispersion region is reached with positive probability. Hence

$$\Psi(P_1^*) \leq P_1^B < P_1^*,$$

so $\Psi(P_1^*) < P_1^*$, and (P4.1) gives $P_1^* > R_2^*$. ■

Remark (general demand). Step 3 uses only Assumption 1, so the compression (P4.2) holds for any admissible demand; only the identity $\mathbb{E}[P_2^0] = P_1^B$ in Step 4 is specific to linear demand and

a mean-zero shock. Under general demand the ordering $P_1^* > R_2^*$ holds whenever $P_1^* \geq \mathbb{E}[P_2^0(\varepsilon)]$, that is, whenever the forward-looking elevation of Theorem 1 is at least the curvature gap between the mean benchmark price and the static benchmark; this condition is satisfied throughout the calibration of Section 4.

A.2.9 Proof of Proposition 5 (Launch Premium at General T)

Notation. Let m denote the prevailing floor, that is the lowest price paid by any earlier cohort, and let k denote remaining inventory. Write the wedge faced by a behaviourally sensitive consumer in epoch t as the composite fairness intensity multiplied by the excess of the current price over the floor, truncated below at zero, so the wedge is zero at any price weakly below the floor. Let the per-epoch payoff under price P, floor m, inventory k and shock e be the product of the price and realised sales, where sales are the lesser of behaviourally adjusted demand and remaining inventory. Let W denote the value of the firm's problem from the second epoch onward under optimal play, as a function of the floor and of initial inventory.

Step 1. The per-epoch payoff is weakly increasing in the floor. Fix an epoch, a price P, an inventory level k and a shock realisation e. If P is weakly below the floor the wedge is zero and the payoff does not depend on the floor. If P exceeds the floor, the wedge is strictly decreasing in the floor, so the behaviourally adjusted price is strictly decreasing in the floor, so behaviourally adjusted demand is weakly increasing in the floor. Sales are the minimum of that demand and remaining inventory, and the minimum of a weakly increasing function and a constant is weakly increasing. Revenue at a fixed price is therefore weakly increasing in the floor, and strictly increasing whenever demand is interior, that is whenever the capacity constraint does not bind and demand is positive.

Step 2. A higher floor cannot reduce continuation value through faster inventory depletion. This is the only channel by which the monotonicity of Step 1 could be overturned, since a higher floor raises sales and so draws down perishable inventory more quickly. Lemma A.4 rules it out: by mimicking, the firm at the higher floor can reproduce the entire path of demand, sales and inventory generated at the lower floor while posting weakly higher nominal prices at every date, so the value function is weakly increasing in the floor irrespective of the shadow price of capacity.

Step 3. Monotonicity of the value function. Proceed by backward induction on the epoch index. In the terminal epoch there is no continuation, and the value is the maximum over prices of the per-epoch payoff, which by Step 1 is the maximum of a family of functions each weakly increasing in the floor, hence weakly increasing. Suppose the value function in epoch t plus one is weakly increasing in the floor. In epoch t the objective is the per-epoch payoff plus the discounted continuation value evaluated at the transitioned floor and inventory. By Step 1 the first term is weakly increasing in the floor; by Step 2 and the induction hypothesis the second is weakly increasing. The maximum over prices of a family of functions each weakly increasing in a parameter is weakly increasing in that parameter. The expectation over shocks preserves this. Hence the value function in epoch t is weakly increasing in the floor, and W is weakly increasing.

Step 4. Strictness. Suppose the dispersion region has positive probability, that is, with positive probability the optimal price in some epoch strictly exceeds the prevailing floor with interior demand. On that event the inequality in Step 1 is strict, because the composite fairness

intensity is strictly positive. Since the event has positive probability, the expectation is strictly increased, so W is strictly increasing in the floor at the relevant point.

Step 5. The first-order condition. Because the launch cohort has no predecessor, its wedge is zero and launch demand is standard. Total value as a function of the launch price is the sum of launch revenue, weighted by the cohort mass one over T, and the continuation value W evaluated at a floor equal to the launch price. The static monopoly benchmark maximises launch revenue, so the derivative of the first term vanishes there. By Step 4 the derivative of the second term is strictly positive there. The derivative of total value at the static monopoly benchmark is therefore strictly positive, so the benchmark cannot be optimal and the optimal launch price lies strictly above it. Neither linearity of demand nor slackness of capacity was used. Setting T equal to two recovers Theorem 1. This completes the proof.

A.2.10 Proof of Proposition 6 (Refinement Monotonicity of the Continuation Value)

Part (i) is Lemma A.5.

For part (ii), write the launch weight as $\theta \equiv 1/T \in (0,1]$ and total value at launch price P, holding the continuation problem fixed, as

$$V(P;\theta) = \theta R(P) + W(P), \qquad R(P) \equiv P\,D(P),$$

where R is launch revenue at unit cohort mass and W the continuation value evaluated at a floor equal to the launch price. By Assumption 1, R is strictly concave with maximiser P_1^B, so $R'(P) < 0$ for every $P > P_1^B$. Hence

$$\frac{\partial^2 V}{\partial P\ \partial \theta} = R'(P) < 0 \qquad \text{for all } P > P_1^B,$$

so V exhibits decreasing differences in (P, θ) on $[P_1^B, \infty)$. By Topkis's monotone comparative statics theorem (Topkis, 1978) the argmax correspondence is non-increasing in θ in the strong set order, and therefore non-decreasing in T. The restriction to $[P_1^B, \infty)$ is essential rather than cosmetic. Since R is strictly concave with maximiser P_1^B, the cross-partial $\frac{\partial^2 V}{\partial P\ \partial \theta} = R'(P)$ is strictly positive be P_1^B, so V exhibits *increasing* differences there and the comparative static reverses. The argument is nonetheless without loss: Proposition 5 establishes that every optimum lies strictly above P_1^B at any T ≥ 2 and any capacity, so the region on which the sign reverses contains no optimum and the argmax correspondence over the restricted domain coincides with the argmax over the whole of it.

Part (iii) restates Proposition 5, which holds for every $T \geq 2$ and every capacity K and is therefore uniform in the repricing frequency.

What is not claimed. Parts (i) and (ii) do not combine into a monotone comparative static of the launch premium in T. Part (ii) holds the continuation problem fixed, whereas refining the window changes it; part (i) signs the level of W_T but not its derivative $W_T{}'$ in the floor, and it is $W_T{}'$ that enters the launch first-order condition

$$\theta R'(P) + W_T{}'(P) = 0.$$

A sufficient condition for monotonicity of the launch price in T along refinements is that $W_{T'}{}'(P) \geq W_T{}'(P)$ for $P \geq P_1^B$ whenever T divides T', that is, that refinement raises the *marginal* as well as the total value of an elevated floor. This is satisfied in every numerical configuration reported in Table 5 but is not established here: with two state variables and a kinked payoff, W_T does not admit a closed form and its cross-partial cannot be signed. The monotonicity in T and the existence of a limit as the window is refined without bound are accordingly reported as numerical findings in Section 4.3 and qualified in Section 4.6. ■

Appendix A.3: Robustness

This appendix shows that the principal comparative statics are robust to relaxing the assumption of linear demand used in the main analysis. Perishable capacity is treated in the main text, in Section 4.4.

A.3.1 General-Demand Comparative Statics

The behavioural comparative statics developed in Section 3 do not depend on the assumption of linear demand. Linear demand is used only to obtain closed-form expressions. This subsection establishes the sign of the period-2 price response under general demand, and gives a curvature condition under which the attenuation result of Proposition 1 also survives.

Throughout, the demand shock is an additive shift,

$$D(P;\varepsilon) = \widetilde{D}(P) + \varepsilon, \qquad \widetilde{D}' < 0,$$

which is the general-demand counterpart of the specification $a_2 = a_1 + \varepsilon$ used in the text. Write $\mu := \lambda + \alpha$ for the pass-through from the posted price to the perceived price in the region where both behavioural mechanisms bind, so that

$$\check{P}_2 = \mu P_2 - (\lambda - 1)R_2^* - \alpha P_1, \qquad \frac{\partial \check{P}_2}{\partial P_2} = \mu.$$

Setting $\alpha = 0$ recovers the expectation-reference boundary case corresponding to Corollary 1, with $\mu = \lambda$.

Sign of the price response

The period-2 first-order condition in the behavioural region is

$$\Phi(P_2, \varepsilon) = (1-\pi)[D(P_2;\varepsilon) + P_2 D'(P_2)] + \pi\left[D\left(\check{P}_2;\varepsilon\right) + \mu P_2 D'\left(\check{P}_2\right)\right] = 0.$$

Implicit differentiation gives

$$\frac{\partial P_2^*}{\partial \varepsilon} = -\frac{\partial \Phi / \partial \varepsilon}{\partial \Phi / \partial P_2}.$$

By Lemma A.1, $\partial \Phi / \partial P_2 < 0$ at the optimum, so the sign depends only on the numerator. Under an additive shift, $D_\varepsilon \equiv 1$ and $D'_\varepsilon \equiv 0$, whence

$$\frac{\partial \Phi}{\partial \varepsilon} = (1-\pi) + \pi = 1 > 0,$$

and therefore $\partial P_2^* / \partial \varepsilon > 0$. A positive demand shock continues to raise the optimal period-2 price under general demand, with no further restriction on $\widetilde{D}$ beyond those already maintained.

The additive form is not merely convenient. Under a multiplicative shock $D(P;\varepsilon) = (1+\varepsilon)\widetilde{D}(P)$, the numerator reproduces the first-order condition itself and vanishes at the optimum, so $\partial P_2^*/\,\partial\varepsilon = 0$ — the familiar neutrality of the unconstrained monopoly price to proportional demand shifts. Some restriction on how the shock moves marginal revenue is therefore unavoidable; the additive shift is the weakest natural one and is what the linear specification implements. More generally, everything below goes through under the assumption that the shock raises marginal revenue, $D_\varepsilon + PD'_\varepsilon > 0$.

Attenuation

The benchmark first-order condition is $\Phi^0(P_2,\varepsilon) = D(P_2;\varepsilon) + P_2D'(P_2) = 0$, with $\partial P_2^0/\,\partial\varepsilon = -1/\Phi_P^0$. Since $\partial P_2^*/\,\partial\varepsilon = -1/\Phi_P$, attenuation is exactly the statement that the behavioural objective is more sharply concave at its optimum, $|\Phi_P| > |\Phi_P^0|$. This does not hold unconditionally. Define the curvature index

$$s(P) := -\frac{PD''(P)}{D'(P)},$$

so that $s \equiv 0$ for linear demand, $s < 0$ for concave demand, and $s > 0$ for convex demand.

Assumption A.2 (bounded curvature). There exists $\bar{s} < 2/(\mu+1)$ such that $s(P) \le \bar{s}$ for all P in the interval spanned by P_2^0, P_2^* and $\check{P}_2^*$.

A.3.2 Lemma A.6 (General-demand attenuation)

Let the demand shock be an additive shift and let Assumption A.2 hold. Then, in the region where the behavioural mechanisms bind,

$$0 < \frac{\partial P_2^*}{\partial\varepsilon} < \frac{\partial P_2^0}{\partial\varepsilon},$$

so the period-2 price responds to demand shocks in the same direction as in the benchmark but by strictly less. Linear demand satisfies Assumption A.2 with $s \equiv 0$ for every $\mu \ge 1$, in which case the ratio of responses is exactly $1/(1+\kappa+\gamma)$, and $1/(1+\kappa)$ when $\alpha = 0$.

Proof. The sign is established above. For the magnitude, differentiate the two first-order conditions:

$$\Phi_P = (1-\pi)[2D'(P_2) + P_2D''(P_2)] + \pi\left[2\mu D'\left(\check{P}_2\right) + \mu^2 P_2 D''\left(\check{P}_2\right)\right],$$

$$\Phi_P^0 = 2D'(P_2) + P_2D''(P_2).$$

Evaluating both at a common price P and substituting $D'' = -sD'/P$,

$$\Phi_P^0 = D'(P)(2-s),$$

$$\Phi_P = (1-\pi)D'(P)(2-s) + \pi D'(P)\mu(2-\mu s),$$

so that

$$\frac{\Phi_P}{\Phi_P^0} = (1-\pi) + \pi\mu\,\frac{2-\mu s}{2-s}.$$

Benchmark concavity requires $s < 2$; behavioural concavity, guaranteed by Lemma A.1, requires the tighter $s < 2/\mu$, which Assumption A.2 implies since $2/(\mu+1) < 2/\mu$ for $\mu > 1$. Both denominators are therefore positive and the ratio exceeds one if and only if

$$\mu(2-\mu s) > 2 - s \Leftrightarrow 2(\mu-1) > s(\mu^2-1) \Leftrightarrow s < \frac{2}{\mu+1},$$

using $\mu > 1$. Assumption A.2 delivers this with strict inequality, so $|\Phi_P| > |\Phi_P^0|$ and hence $\partial P_2^*/\,\partial\varepsilon < \partial P_2^0/\,\partial\varepsilon$. Under linear demand $D'' \equiv 0$, the ratio is

$$(1-\pi) + \pi\mu = 1 + \pi(\lambda-1) + \pi\alpha = 1 + \kappa + \gamma$$

exactly, giving the stated attenuation factor. ■

Remark (the role of curvature). The gap between the concavity requirement $s < 2/\mu$ and the attenuation requirement $s < 2/(\mu+1)$ is not an artefact of the argument. The behavioural wedge multiplies the pass-through from the posted price to the perceived price by μ, which sharpens the slope term ($2\mu D'$ against $2D'$, attenuating) but multiplies the curvature term by μ^2 ($\mu^2 PD''$ against PD''). Where demand is convex the second effect works against concavity and, for $s \in (2/(\mu+1),\ 2/\mu)$, dominates: the behavioural optimum remains well defined but the price response is *amplified* rather than attenuated. With $\lambda = 2$ and $\alpha = 1$, so $\mu = 3$, attenuation requires $s < 0.5$ while concavity permits s up to 0.667. Linear demand is the case in which the curvature term is absent altogether, which is why the closed form is unconditionally clean, and all weakly concave demand ($s \le 0$) satisfies Assumption A.2 for every μ. The condition is therefore restrictive only for demand that is appreciably convex.

Remark (locality). Assumption A.2 bounds curvature on the interval spanned by the three relevant prices, but the proof compares Φ_P and Φ_P^0 at a common point, whereas Φ_P is evaluated with D' and D'' at $\check{P}_2^*$ and Φ_P^0 at P_2^0. The lemma is therefore exact at the linear benchmark and valid to first order in curvature; a global statement requires in addition that $|D'|$ vary boundedly over that interval, since the *levels* of the slope at the distinct evaluation points also enter the comparison. Nothing in Section 4 turns on the global version.

The linear-demand expressions of Section 3 are thus not merely illustrative. They are the case in which attenuation holds for every admissible μ, and the general-demand analysis identifies precisely how far demand may depart from linearity before the ranking is at risk.

Appendix A.4 Numerical Implementation

A.4.1 Two numerical exercises

The paper reports two distinct numerical exercises, and it is worth separating them at the outset because they solve different models and are not directly comparable.

The first computes the two-period unified model of Sections 3.2 to 3.8, in which both reference objects are active and the expectation reference is determined in personal equilibrium. It produces Tables 8 and 9. The second computes the T-period extension of Section 4, in which the selling window is divided into repricing epochs, capacity is finite, and only the comparison channel operates. It produces Tables 5 to 7. Magnitudes from the two exercises should not be read across: they differ in the number of reference objects, in the presence of a capacity constraint, and in the width of the shock support, for reasons given below.

Unless otherwise stated the numerical exercises set the procedural-legitimacy parameter to $J = 0$, so that no legitimacy discount attenuates the comparison channel. The reported distortions are therefore the largest attainable along that margin, and Proposition 3 implies that any positive legitimacy would reduce them.

A.4.2 Environment and calibration

The numerical exercises are intended to illustrate the economic significance of the behavioural mechanisms rather than to provide structural estimates. Demand is linear,

$$D(P) = a_t - bP,$$

with the period-1 intercept normalised to $a_1 = 1$, the slope to $b = 1$, and the discount factor initially set to $\beta = 1$. These normalisations imply the benchmark period-1 optimum

$$P_1^0 = a_1/2b = 1/2,$$

which coincides with the static monopoly benchmark P_1^B because, absent behavioural effects, the continuation value does not depend on the first-period price. The corresponding choke price is $\bar{P} = \frac{a_1}{b} = 1$. Demand uncertainty enters through a uniform shock to the intercept,

$$\varepsilon \sim \mathrm{U}[-d, d], \quad a_t = a_1 + \varepsilon,$$

with the launch period facing no shock, so that the no-behavioural launch price is exactly P_1^B and any departure from 0.5 is attributable to the behavioural mechanism alone.

The width of the shock support differs between the two exercises, and the difference is substantive rather than incidental. The two-period model sets $d = 0.9$; the T-period model sets $d = 0.5$. In the two-period model each period covers half the selling window, so a period shock represents the full variation in demand over that half. In the T-period model each epoch is one of T successive draws over a window of fixed total demand, so an epoch shock represents a correspondingly smaller share of aggregate uncertainty. Holding d fixed across the two would therefore impose more aggregate uncertainty on the T-period model as T rises,

which is precisely the confound the fixed-window construction of Section 4.1 is designed to avoid. Panel (iv) of Table 8 reports the sensitivity of the two-period results to the narrower support.

The behavioural environment is summarised by two composite indices. The comparison channel is governed by

$$\Gamma = \pi\alpha(1 + \sigma\varphi),$$

which combines the share of behaviourally sensitive consumers, the intensity of inequity aversion, and the amplification generated by identity salience. The expectation channel is governed by the smoothing index $\kappa = \pi(\lambda - 1)$, which combines the same behavioural share with the loss-aversion coefficient. The expectation reference is not a free parameter: it is determined endogenously in personal equilibrium as $R_2 = E[P_2]$, so that consumers' price expectations are correct on average. The numerical experiments vary the primitives $(\lambda, \pi, \alpha, \sigma\varphi)$ over representative ranges while preserving the analytical structure of Section 3.

The behavioural intensities adopted here sit at or below the range implied by the estimates of Richards, Liaukonytė and Streletskaya (2016) discussed in Section 4.2, so the calibration is conservative with respect to the magnitude of the reported distortions. The normalisations affect the scale of reported prices but not the proportional distortions that constitute the principal predictions of the theory.

A.4.3 Computing expectations

Every quantity reported in Tables 5, 6, 8 and 9 is an expectation over the demand shock, and the accuracy with which those expectations are computed determines the accuracy of the results. Because ε is continuous, the expectation cannot be evaluated at every realisation. It is instead approximated by a weighted sum,

$$\mathrm{E}[f(\varepsilon)] \approx \sum_i w_i \cdot f(\varepsilon_i),$$

where the ε_i are the points at which the function is evaluated, known as nodes, and the w_i are the weights attached to them. All numerical integration takes this form; methods differ only in how the nodes and weights are chosen.

The obvious choice is to space the nodes evenly and weight them equally, which amounts to sampling the shock on a uniform grid. This is simple but wasteful, because it forgoes the freedom to choose where the function is evaluated. Gauss–Legendre quadrature exploits that freedom. With n nodes there are $2n$ quantities to choose (n positions and n weights) and they can be selected so that the rule integrates every polynomial of degree up to $2n - 1$ exactly rather than approximately. The nodes that achieve this are not evenly spaced but cluster towards the ends of the interval, where a function has most scope to depart from the behaviour implied by its values elsewhere; the interior is easier to pin down and requires fewer points, and carries correspondingly larger weights.

The intuition is that a well-chosen handful of evaluation points can capture what a large number of arbitrarily placed ones would. Determining a straight line does not require twenty evenly

spaced measurements: two well-chosen points determine it exactly. Gauss–Legendre generalises this principle to higher-degree functions.

The method is well suited to the present model. The period-2 objective is piecewise quadratic, so on each behavioural regime it is a polynomial of degree two, far within the exactness of a rule accurate to degree 29. On each smooth piece the quadrature is therefore not approximating at all but is exact to machine precision, and the only approximation error arises at the kinks where the regimes meet, at which the objective ceases to be polynomial. Fifteen nodes accordingly deliver accuracy that random sampling would require many orders of magnitude more evaluations to match.

The T-period exercises use a fifteen-node rule and the two-period exercises a sixty-node rule, the latter reflecting the personal-equilibrium fixed point, which requires the expectation to be recomputed at every iteration and at every candidate first-period price. Because ε is uniform on $[-d, d]$ its density is constant at $1/2d$, so the standard rule on $[-1, 1]$ is applied after rescaling the nodes by d and normalising the weights to sum to one; the weights are then probabilities and the approximation is a straightforward weighted average.

Two consequences follow. The computation is deterministic: the same nodes and weights are used on every run, so every reported figure is exactly reproducible and none is subject to sampling error. And no convergence criterion is required for the integration itself, since the rule is exact on the polynomial pieces; the only discretisations requiring verification are the price and inventory grids, addressed in A.4.6.

A.4.4 Solving the model

The equilibrium is computed by backward induction. For each node of the shock distribution the period-2 problem is solved on a fine price grid, exploiting the strict concavity established in Lemma A.1 within each behavioural regime and comparing the resulting regime maxima. The order of operations matters and reflects the timing assumption of the model: the optimal price is computed separately at each shock realisation, and the expectation is taken over the resulting profits. Reversing the order (taking the expectation first and then optimising) would describe a firm committing to a price before observing demand, which is not the problem posed in Section 3.1.

The period-1 price is obtained by maximising the sum of first-period profit and the expected continuation value, with the personal-equilibrium expectation reference imposed as a fixed point at each candidate price. The fixed point is computed by iteration, which converges because the mapping is a contraction (Lemma A.3).

The T-period extension of Section 4 is solved by backward induction over a pair of state variables: the running price floor and remaining inventory. The floor lies on the price grid, so its transition (the running minimum of prices posted so far) is exact rather than interpolated. Continuation values are interpolated linearly in inventory between grid points. With capacity K and no replenishment, remaining inventory is confined to the interval $[0, K]$: it begins at K in the launch epoch and is non-increasing thereafter. The inventory grid is therefore constructed on $[0, K]$ rather than on $[0, 1]$, since points above K correspond to states the firm cannot occupy and would reduce the effective resolution by a factor of $1/K$.

The period-2 objective is globally strictly concave (Lemma 1), so the inner problem has a unique solution. The period-1 objective is not established to be concave (Remark 4); single-peakedness was verified at every grid point over the reported parameter range.

A.4.5 Simulated price paths

Table 7 is computed differently from every other table in the paper. The quantities it reports, the spread between the highest and lowest price posted within a selling window, and the standard deviation of prices across that window, are properties of a realised price path rather than of any single epoch. They cannot be recovered from epoch-by-epoch expectations, because the maximum and minimum of a sequence are not functions of the marginal distributions of its elements. Obtaining them requires simulating whole paths.

The optimal policy is first computed by the backward induction of A.4.4, using quadrature as elsewhere. Four thousand selling windows are then simulated under that policy, and the dispersion statistics are computed within each window and averaged across windows. The shocks are drawn from the quadrature nodes with probabilities equal to the quadrature weights, so that the simulation and the optimisation face an identical shock support; sampling from a continuous uniform instead would confront the policy with realisations it was not optimised against.

Table 7 is therefore the only table in the paper subject to sampling error. A fixed random seed is recorded in the code, so the reported figures are reproducible exactly, but they would shift slightly under a different seed. Because the statistics are averages over four thousand paths, that variation is small relative to the differences the table reports: the mean spread falls by more than half between the lowest and highest fairness intensities.

A.4.6 Grid resolution and sensitivity

Two discretisations require verification: the price grid and, for the capacity extension, the inventory grid.

The price grid has width 0.001 over the relevant range. Because the floor transition is exact on this grid and the period-2 objective is piecewise quadratic, refinement affects the reported prices only through rounding at the third decimal.

The inventory grid is the only discretisation introduced by the capacity extension, and its resolution was verified by refinement. At the tightest capacity considered, $K = 0.15$, the fairness premium is +0.039 at 41 nodes, +0.041 at 81 and +0.041 at 161; at $K = 0.25$ it is +0.062 at both 41 and 81 nodes. Consecutive refinements therefore agree to within 0.001 at 81 nodes, which is the resolution adopted for Tables 6 and 7. An internal consistency check confirms the extension: at $K = 0.90$ capacity does not bind, and the fairness premium of +0.063 coincides exactly with the slack-capacity plateau reported in Table 5 at $T = 8$ and with the saturated launch price of Table 7.

Robustness is assessed by varying each behavioural parameter about the central calibration. The direction of every reported comparative static is invariant across the ranges considered: the initial price rises and the equilibrium price difference across cohorts falls as the effective comparison intensity increases, the response to demand shocks is attenuated as the effective expectation intensity increases, and both effects are amplified by social identity and attenuated by procedural legitimacy. Magnitudes scale approximately with the composite indices rather

than with the underlying primitives separately, which is the sense in which the composites are the economically relevant objects. The results are not driven by the linear demand specification, as Appendix A.3 establishes analytically for general demand under a bounded-curvature condition.

A.4.7 Replication

All numerical results reported in the paper are reproducible from the documented parameter values. The accompanying code computes the equilibrium of the two-period model of Sections 3.2 to 3.8 and of the T-period extension of Section 4, and generates the calibration table and disclosure counterfactual reported in Tables 8 and 9 and the launch-premium, capacity and dispersion results reported in Tables 5 to 7. Because expectations are computed by quadrature rather than by sampling, every figure other than those in Table 7 is deterministic and exactly reproducible. Table 7 uses a fixed random seed, recorded in the code. The code and a replication script will be made available on publication.

References

Akerlof, George A. and Rachel E. Kranton (2000), "Economics and Identity," *Quarterly Journal of Economics*, 115(3), pp. 715–753.

Anderson, Eric T. and Duncan I. Simester (2010), "Price Stickiness and Customer Antagonism," *Quarterly Journal of Economics*, 125(2): pp. 729–765.

BBC News (2024), "Oasis Hit Out at Ticketmaster's Dynamic Pricing After Backlash," *BBC News*, September 4.

Bellemare, Charles, Sabine Kröger and Arthur van Soest (2008), "Measuring Inequity Aversion in a Heterogeneous Population Using Experimental Decisions and Subjective Probabilities," *Econometrica*, 76(4): pp. 815–839.

Blake, Tom, Moshary, Sarah, Sweeney, Kane and Steve Tadelis (2021), "Price Salience and Product Choice," *Marketing Science*, 40(4): pp. 619–636.

Bolton, Gary E. and Axel Ockenfels (2000), "ERC: A Theory of Equity, Reciprocity, and Competition," *American Economic Review*, 90(1): pp. 166–193.

Chen, Yan and Sherry Xin Li (2009), "Group Identity and Social Preferences," *American Economic Review*, 99(1): pp. 431–457.

CMA (2024), "CMA launches investigation into Ticketmaster over Oasis concert sales," Press Release (September 5).

CMA (2025), "CMA seeks changes to the way Ticketmaster labels tickets and provides pricing information to fans," Press Release (March 25).

Cohen, Maxime C., Sentao Miao and Yining Wang (2025), "Dynamic Pricing with Fairness Constraints," *Operations Research*, 73(6): pp. 3027–3043.

Courty, Pascal (2003), "Some Economics of Ticket Resale," *Journal of Economic Perspectives*, 17(2): pp. 85–97.

DellaVigna, Stefano (2009), "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, 47(2): pp. 315–372.

Drayer, Joris, Stephen L. Shapiro and Soonhwan Lee (2012), "Dynamic Ticket Pricing in Sport: An Agenda for Research and Practice," *Sport Marketing Quarterly*, 21(3): pp. 184–194.

Eyster, Erik, Kristóf Madarász and Pascal Michaillat (2021), "Pricing under Fairness Concerns," *Journal of the European Economic Association*, 19(3): pp. 1853–1898.

Fehr, Ernst and Klaus M. Schmidt (1999), "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, 114(3): pp. 817–868.

Ferreira, Kris Johnson, Bin Hong Alex Lee and David Simchi-Levi (2016), "Analytics for an Online Retailer: Demand Forecasting and Price Optimization," *Manufacturing & Service Operations Management*, 18(1), pp. 69–88.

Gallego, Guillermo and Garrett van Ryzin (1994), "Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons," *Management Science*, 40(8): pp. 999–1020.

Guo, Liang (2015), “Inequity Aversion and Fair Selling,” *Journal of Marketing Research*, 52(1): pp. 77–89.

Hardie, Bruce G. S., Eric J. Johnson and Peter S. Fader (1993), “Modeling Loss Aversion and Reference Dependence Effects on Brand Choice,” *Marketing Science*, 12(4): pp. 378–394.

Haws, Kelly L. and William O. Bearden (2006), “Dynamic Pricing and Consumer Fairness Perceptions,” *Journal of Consumer Research*, 33(3): pp. 304–311.

Heidhues, Paul and Botond Kőszegi (2008), “Competition and Price Variation When Consumers Are Loss Averse,” *American Economic Review*, 98(4): pp. 1245–1268.

Kahneman, Daniel and Amos Tversky (1979), “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica*, 47(2): pp. 263–291.

Kahneman, Daniel, Jack L. Knetsch and Richard Thaler (1986), “Fairness as a Constraint on Profit Seeking: Entitlements in the Market,” *American Economic Review*, 76(4): pp. 728–741.

Kalyanaram, Gurumurthy and Russell S. Winer (1995), “Empirical Generalizations from Reference Price Research,” *Marketing Science*, 14(3): pp. G161–G169.

Kimes, Sheryl E. (1994), “Perceived Fairness of Yield Management,” *Cornell Hotel and Restaurant Administration Quarterly*, 35(1): pp. 22–29.

Kőszegi, Botond and Matthew Rabin (2006), “A Model of Reference-Dependent Preferences,” *Quarterly Journal of Economics*, 121(4): pp. 1133–1165.

Krasikov, Ilia and Rohit Lamba (2026), “On Dynamic Pricing,” *American Economic Journal: Microeconomics*, 18(1): pp. 174–207.

Leasca, Stacey (2024), “Wendy’s Still Plans to Introduce Dynamic Pricing in 2025 — Here’s Why It May Not Be a Great Idea,” *Food & Wine*, (February 29).

Li, Krista J. and Sanjay Jain (2016), “Behavior-Based Pricing: An Analysis of the Impact of Peer-Induced Fairness,” *Management Science*, 62(9): pp. 2705–2721.

Li, Wenjing, David M. Hardesty and Adam W. Craig (2018), “The Impact of Dynamic Bundling on Price Fairness Perceptions,” *Journal of Retailing and Consumer Services*, 40: pp. 204–212.

London, Lela (2023), “British Pubs Will Charge More To Drink At ‘Peak’ Times,” *Forbes*, (September 12).

Mazumdar, Tridib, S. P. Raj and Indrajit Sinha (2005), “Reference Price Research: Review and Propositions,” *Journal of Marketing*, 69(4): pp. 84–102.

Nasiry, Javad and Ioana Popescu (2011), “Dynamic Pricing with Loss-Averse Consumers and Peak-End Anchoring,” *Operations Research*, 59(6), pp. 1361–1368.

Phillips, Robert L. (2005), *Pricing and Revenue Optimization*. Stanford University Press.

Pope, Devin G. and Maurice E. Schweitzer (2011), “Is Tiger Woods Loss Averse? Persistent Bias in the Face of Experience, Competition, and High Stakes,” *American Economic Review*, 101(1): pp. 129–157.

Popescu, Ioana and Yaozhong Wu (2007), “Dynamic Pricing Strategies with Reference Effects,” *Operations Research*, 55(3): pp. 413–429.

Rabin, M. (1993), “Incorporating fairness into game theory and economics,” *American Economic Review* 83(5), pp. 1281–1302.

Richards, Timothy J., Jura Liaukonyte and Nadia A. Streletskaya (2016), “Personalized Pricing and Price Fairness,” *International Journal of Industrial Organization*, 44: pp. 138–153.

Selove, Matthew (2019), “Dynamic Pricing with Fairness Concerns and a Capacity Constraint,” *Quantitative Marketing and Economics*, 17(4): pp. 385–413.

Shapiro, Stephen L. and Joris Drayer (2012), “A New Age of Demand-Based Pricing: An Examination of Dynamic Ticket Pricing and Secondary Market Prices in Major League Baseball,” *Journal of Sport Management*, 26(6): pp. 532–546.

Smith, Barry C., John F. Leimkuhler and Ross M. Darrow (1992), “Yield Management at American Airlines,” *Interfaces*, 22(1): pp. 8–31.

Stokey, Nancy L. (1979), “Intertemporal Price Discrimination,” *Quarterly Journal of Economics*, 93(3): pp. 355–371.

Tajfel, Henri and John C. Turner (1979), “An Integrative Theory of Intergroup Conflict,” in W. G. Austin and S. Worchel (eds.), *The Social Psychology of Intergroup Relations*, Brooks/Cole, Monterey, CA, pp. 33–47.

Talluri, Kalyan T. and Garrett J. van Ryzin (2004), *The Theory and Practice of Revenue Management*, Springer.

Topkis, D. M. (1978), “Minimizing a submodular function on a lattice,” *Operations Research* 26(2), pp. 305–321

Tversky, Amos and Daniel Kahneman (1991), “Loss Aversion in Riskless Choice: A Reference-Dependent Model,” *Quarterly Journal of Economics*, 106(4): pp. 1039–1061.

Tversky, Amos and Daniel Kahneman (1992), “Advances in Prospect Theory: Cumulative Representation of Uncertainty,” *Journal of Risk and Uncertainty*, 5(4): pp. 297–323.

Williams, Kevin R. (2022), “The Welfare Effects of Dynamic Pricing: Evidence from Airline Markets,” *Econometrica*, 90(2): pp. 831–858.

Winer, Russell S. (1986), “A Reference Price Model of Brand Choice for Frequently Purchased Products,” *Journal of Consumer Research*, 13(2): pp. 250–256.

Witts, Sophie (2023), “Unhappy hour? Stonegate’s decision to launch dynamic pricing sparks debate,” *The Caterer*, (September 12).

Xia, Lan, Kent B. Monroe and Jennifer L. Cox (2004), “The Price Is Unfair! A Conceptual Framework of Price Fairness Perceptions,” *Journal of Marketing*, 68(4): pp. 1–15.