Optimal Promotion Frequency and Arrangement with Demand Memory

Weiguang Liu[1] Will Ma[2] Chung-Piaw Teo[3] Lan Wu[4]

[1]University College London
[2]Columbia Business School, Columbia University
[3]NUS Business School, National University of Singapore
[4]Shenzhen Research Institute of Big Data

Abstract

We study promotion frequency and arrangement that maximize long run average revenue with two given prices. Customers purchase repeatedly, and demand depends on age, current price, and past promotion counts. A telescoping representation of marginal promotion effects yields a revenue decomposition into moving window contributions. Concave contributions make balanced arrangements optimal at fixed cycle length and promotion count. Frequency optimization then reduces to a finite search over a concave, piecewise linear function. Convex contributions make a single promotion block optimal at fixed counts and a constant price optimal when frequency is unrestricted. Additively separable and multiplicative demand give economic interpretations through engagement, reference prices, and retention. We derive revenue loss bounds for balanced and single block policies under mixed curvature and for memory truncation with general demand. Numerical experiments illustrate implementation, sensitivity to price order, and the gap between policy performance and bound tightness.

Keywords promotion planning, customer memory, engagement channel, reference price channel, balanced arrangements

1 Introduction

Promotion planning is a central problem for businesses, such as retailers, online platforms, and service providers. Consumer goods companies in an industry study devoted nearly 20% of their gross revenue to supporting retail promotions [31]. Better promotion decisions can also yield substantial revenue gains. In an evaluation using data from large retailers, Cohen-Hillel et al. [11] report average improvements ranging from 5.1% to 15.6% relative to existing practices. Promotion planning involves two key decisions. Frequency determines how often discounts are offered, while arrangement determines how the promotional periods are arranged in time. Businesses use both dispersed and concentrated patterns in practice. The movie theater chain Cinemark offers discounted tickets on Tuesdays at participating theaters [8], while online retailer Amazon held its 2025 Prime Day event over four consecutive days [4].

While the frequency of promotions has a clear impact on demand, arrangements of promotions can also substantially affect demand and revenue. Research shows that promotional patterns influence consumers' perceptions of deal frequency and that these perceptions are associated with willingness to pay [19]. Laboratory evidence also shows that purchase quantities vary across promotional patterns [20]. Consider a seller planning four promotion days in a month. Spreading them across four weeks exposes customers to a different price history than grouping them as a campaign at the beginning of the month. Also, different arrival times of customers will lead to different experiences of promotions. Under the concentrated

arrangement, early arrivers may experience several consecutive promotions, while later arrivers may encounter none during the remainder of the month, by which time they may already have decided to drop out.

Furthermore, promotions within customer memory can affect future demand. Empirical research documents persistent effects on price sensitivity and subsequent demand, with their direction depending on the customer and setting [5, 23, 25]. Offering one additional promotion may have a diminishing or increasing effect on future demand. We hence note the following two channels. In the *reference price channel*, past discounts lower the benchmark against which customers evaluate the current price, potentially reducing subsequent demand at a given price [22, 34]. In the *engagement channel*, promotions encourage customers to return, browse, or purchase again and may lead to higher demand in the future. A randomized experiment on the online retailer Alibaba documents increases in subsequent product views and purchase incidence [35], while a supermarket field experiment finds that coupon redemption can reduce customer churn and increase the promoted product's customer lifetime value [32]. Arrangement decisions therefore need to take these into account as well.

The central question of this paper is how a seller should choose both promotion frequency and arrangement to maximize long run average revenue when demand depends on the current price and remembered past promotions. Our analysis provides sufficient conditions that make both decisions solvable under finite customer memory.

This paper is related to several research areas. The literature on reference prices establishes how historical price perceptions influence promotion profitability and optimal price dynamics. Greenleaf [14] studies profitable recurring promotions, and Kopalle et al. [18] derive conditions for constant or cyclic pricing under asymmetric responses to gains and losses. Fibich et al. [13] obtain explicit solutions for nonsmooth reference price models, while Popescu and Wu [26] characterize pricing dynamics under general nonlinear reference effects. More recently, Agrawal and Tang [1] establish near optimal markdown policies when reference prices reflect average past prices. These results depend on the reference formation mechanism and the response to perceived gains and losses. We allow for engagement effects as well as reference-price effects, and aim to characterize optimal arrangement and frequency when demand responds to the number of remembered promotions in complex ways.

There is also literature on applied promotion planning research which develops computational methods for optimizing promotion timing when demand depends on past prices. Silva-Risso et al. [30] combine demand estimation, market simulation, and optimization to plan manufacturers' promotions. Cohen et al. [9] develop a linear approximation that sums the profit increments from individual promotions, while Cohen et al. [10] incorporate substitution and complementarity across items. Baardman et al. [6] study promotional vehicle scheduling under capacity constraints. These studies also establish exact solutions in specified cases, including sufficiently separated promotions in Cohen et al. [9] and an unconstrained setting with two prices and substitutable items in Cohen et al. [10]. Our analysis allows for general market share functions and derives sufficient conditions for optimal arrangement and frequency choices.

Another related body of literature establishes results on which features of past prices influence demand under customer memory. Nasiry and Popescu [24] characterize optimal pricing when the reference price combines the lowest and most recent prices. Cohen-Hillel et al. [11] establish conditions for optimal high and low pricing and provide an exact polynomial dynamic program for a single product with demand based on peak and recent prices. Baek et al. [7] characterize optimal cycles when customers remember the best recent offer and revenue satisfies a monotonicity condition. Our model allows for demand to depend on the number or counts of promotional periods encountered by a customer and characterizes the optimal arrangement and frequency. We also show how to handle the case where promotional effects depend on how recent the promotions are, without prescribing a particular functional form in the affine case in Section 4.1.2.

Optimal sequencing literature provides the main mathematical precedent for our arrangement compar-

isons. Hajek [15] and Altman et al. [3] establish optimality of regular sequences under multimodular costs, while Altman et al. [2] also connect balanced routing to activity rate optimization. Their arguments link discrete convexity to the benefits of distributing repeated decisions evenly. However, their methods cannot be directly applied in our settings and we need to develop new techniques to characterize the optimal arrangements.

Technical difficulties of the promotion planning problem include the fact that each promotion enters several customer memory windows within one common price sequence. Also, nonlinear market share functions create interactions among promotions within each history, and transaction prices determine how those interactions contribute to revenue. A standard finite graph formulation establishes the existence of an optimal periodic policy, but its state space grows exponentially with memory depth, making computation expensive for long memories. Hence, it is an important and challenging question to characterize the optimal arrangements and frequency and find a strategy to solve the problem efficiently.

Our main contribution is a set of sufficient conditions for globally optimal promotion arrangements and frequencies. We use a finite telescoping representation of market share, with increments given by the marginal effect $\Delta^{j}_{a,k}$ of increasing remembered promotions from $k-1$ to k, holding age and current price fixed. These marginal effects yield a decomposition of cycle revenue into a component fixed by cycle length and number of promotions and a sum of window contributions F_b. Concave contributions favor balanced arrangements, in which the number of promotions differs by at most one across windows of the same length. Convex contributions favor a single promotion block. These conditions account for nonlinear demand, customer age, and transaction prices.

We can then solve the promotion frequency optimization problem under the same conditions. With concave window revenue contributions, revenue under balanced cycles is a concave, piecewise linear function of promotion frequency. An optimal frequency can be selected from an explicit finite set, yielding a policy that is globally optimal over all feasible infinite price paths and has a cycle of length at most maximum active age plus one. With convex window revenue contributions, a single block is optimal at a prescribed cycle length and number of promotions, while unrestricted frequency optimization yields a constant price policy.

We apply the general results to additively separable and multiplicative market share, two special cases with clearer economic interpretations. For additive demand, we distinguish engagement from reference price effects. Diminishing marginal engagement in exposure and age gives balanced optimality, with affine demand as a special case and a further extension allowing effects to vary across lags. Linear, loss neutral reference price effects give a constant price benchmark. Multiplicative demand models retention and churn through a Markov chain, with price and retention ratios determining arrangement and frequency regimes.

In addition to the exact solutions, for the mixed-curvature scenarios, we provide two computational approximations with revenue guarantees. Finite searches over balanced and single block policies have loss bounds controlled by violations of the corresponding curvature conditions. Memory truncation reduces the graph and bounds loss through demand errors caused by omitted histories.

We examine implementation, sensitivity to demand assumptions, and computational approximation through three numerical experiments. We first illustrate how demand estimation and structural checks guide the selection of a pricing cycle. We then use reference-price demand to assess the consequences of ignoring the order of past prices. Finally, we evaluate balanced and single block policies and memory truncation alongside their revenue loss bounds to identify when these approximations are useful and how well the guarantees reflect policy performance.

The remainder of the paper is organized as follows. Section 2 presents the model and its periodic reformulation. Section 3 develops the revenue decomposition, arrangement conditions, and frequency optimization. Sections 4 and 5 examine additively separable and multiplicative market share functions. Section 6 develops simple policy and memory truncation guarantees under mixed curvature. Section 7 examines

implementation, sensitivity to price order, and approximation performance. Section 8 concludes. The supplementary appendix contains additional proofs and supplementary numerical results.

2 Model with Finite Customer Memory

2.1 Model setup

We consider the following model. A seller offers a single product and chooses its posted price in each period to maximize long run average revenue. That is, in period $t = 0, 1, 2, \ldots$, the seller posts the same price to all active customers,

$$p_t \in \{p^L, p^H\}, \qquad 0 < p^L < p^H,$$

where p^L and p^H are the given promotional and regular prices. Let $\Delta p = p^H - p^L$ denote the price gap.

A unit mass of customers arrives in each period. Customers arriving in the same period form a cohort, whose age a is the number of periods since arrival. Each customer can purchase one unit of the product in each active period and can therefore purchase repeatedly. We suppose that a cohort can remain active at ages $a = 0, \ldots, \mathscr{A}$, where the nonnegative integer $\mathscr{A}$ denotes the maximum active age.

Cohort market share functions depend on the prices observed since arrival. An age a cohort has observed a past prices, all of which can affect current demand. The corresponding window including the current period has length $a + 1$. Thus, $\mathscr{A}$ determines how long customers remain active, and customers remember their past price exposure throughout this active lifetime.

We characterize demand through the expected proportion of each cohort that buys in the current period. Define

$$d_t = \mathbf{1}\{p_t = p^L\}, \qquad y_{a,t} = \sum_{j=1}^{a} d_{t-j},$$

where d_t indicates a promotion at time t and $y_{a,t}$ counts promotions in the cohort's remembered past periods. The number of past promotions is zero when $a = 0$.

The first assumption we make is that the expected demand from a cohort of age a depends on the current price p and the number of promotions encountered in the remembered price histories.

Assumption 2.1. For each age a, the expected demand at time t is $M_a(p_t, y_{a,t})$, where the common *market share function*

$$\begin{aligned} M_a \; &: \{p^L, p^H\} \times \{0, \ldots, a\} \longrightarrow [0, 1], \\ &(p_t, y_{a,t}) \longmapsto M_a(p_t, y_{a,t}), \end{aligned}$$

gives the expected proportion of the cohort that will purchase.

Expected revenue in a period equals the posted price multiplied by total expected demand across active cohorts. Let

$$\Pi = \left\{\pi = (p_t)_{t \geq 0} \,:\, p_t \in \{p^L, p^H\}\right\}$$

be the set of feasible price paths. Fix an initial price history $(p_{-\mathscr{A}}, \ldots, p_{-1})$. For $\pi \in \Pi$, define

$$\bar{R}(\pi) = \liminf_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} p_t \sum_{a=0}^{\mathscr{A}} M_a(p_t, y_{a,t}), \qquad \bar{R}^* = \sup_{\pi \in \Pi} \bar{R}(\pi).$$

Starting with an empty population or a different initial price history changes only finitely many initial revenues and leaves the long run objective unchanged. Each price decision affects current revenue through the transaction price and current market share, and affects subsequent revenue through customers' remembered promotional exposure.

2.2 Motivation for the modeling assumptions

We study a high-low pricing set up which is an established framework in the classical promotion literature [21, 27]. Such policies can be optimal under conditions on price dispersion and the relative strength of current and historical price effects [11]. They are straightforward to implement through a price path specifying when to switch between two preset prices.

We assume there is a maximum active age $\mathscr{A}$ which keeps the active population bounded as new cohorts arrive. Models of customer dropout and evidence on relationship duration motivate finite customer activity [12, 28, 29]. Observed customer stay lengths can inform this cutoff. Customers remember prices observed since arrival throughout their active lifetime, so cohorts of different ages have different history lengths. Finite price histories also appear in promotion planning models such as Cohen et al. [9].

The market share function $M_a(p, y)$ separates the current price from a summary of historical exposure, following reference price models such as Greenleaf [14], Kopalle et al. [18], and Popescu and Wu [26]. Average past prices and the cumulative number of promotions provide examples of scalar history summaries [1, 25]. Assumption 2.1 additionally treats remembered past periods symmetrically. Under binary pricing, the number of past promotions determines the average remembered price. For $a \geq 1$,

$$r_{a,t} = \frac{1}{a}\sum_{j=1}^{a} p_{t-j} = p^H - \frac{\Delta p}{a} y_{a,t}.$$

For a given age, the historical average and the number of promotions therefore contain the same information. A market share function of the current price and a reference price that is the average of past prices can consequently be incorporated in our market share function $M_a(p, y)$. Holding the current price fixed, demand may decrease with past promotions through the reference price channel or increase through engagement. We also do not restrict how the market share function depends on the number of past promotions, allowing for nonlinear effects and interactions among multiple past promotions.

The affine extensions in Sections 4.1.2 and 4.2 allow historical effects to vary across lags. Historical coefficients can be positive under engagement and negative under reference price effects, while the current promotion retains a separate coefficient.

2.3 Cyclic formulation

In this section we show that an optimal pricing policy can be chosen periodic. This result reduces the seller's problem to selecting a finite cycle of prices while preserving the optimal value over all feasible infinite price paths. Finite memory and the stable age composition make current revenue a fixed function of the current price and a finite history of past prices. We use this property to represent the pricing problem on a finite graph.

We define state s_t, which records the price history needed to calculate current revenue and update customers' remembered exposure, that is,

$$s_t = (p_{t-\mathscr{A}}, p_{t-\mathscr{A}+1}, \dots, p_t) \in \{p^L, p^H\}^{\mathscr{A}+1}.$$

The last entry is the current price. For a state $s = (q_0, \dots, q_{\mathscr{A}})$, the number of past promotions remembered by an age a cohort is

$$y_a(s) = \sum_{j=1}^{a} \mathbf{1}\{q_{\mathscr{A}-j} = p^L\},$$

with the empty sum equal to zero. The revenue associated with this state is

$$w(s) = q_{\mathscr{A}} \sum_{a=0}^{\mathscr{A}} M_a\big(q_{\mathscr{A}}, y_a(s)\big).$$

Although current demand depends on the number of past promotions, the remembered price sequence determines how that number changes as prices leave the memory window.

Feasible pricing decisions determine transitions between states. Let $\mathcal{G}_{\mathscr{A}}$ be the directed graph whose nodes are the states defined above, with an edge

$$(q_0, \dots, q_{\mathscr{A}}) \longrightarrow (q_1, \dots, q_{\mathscr{A}}, q) \qquad \text{for each } q \in \{p^L, p^H\}.$$

Each transition removes the oldest price and appends the next posted price. A feasible price path therefore traces a walk through the graph, with node weights equal to period revenues. Repeated traversal of a directed cycle generates a periodic price path, whose average revenue equals the mean node weight of the cycle.

Proposition 2.1 (Existence of an optimal periodic policy)**.** *The optimal long run average revenue equals the maximum mean cycle value in $\mathcal{G}_{\mathscr{A}}$ with node weights $w(s)$. In particular, there exists an optimal periodic promotion policy.*

We can therefore formulate the remaining analysis in terms of finite price cycles. Write

$$C = (p_1(C), \dots, p_{\mathscr{T}}(C)), \qquad p_t(C) \in \{p^L, p^H\},$$

where $\mathscr{T} \geq 1$ is the cycle length. Cycle indices may be any integer, including values below 1 or above $\mathscr{T}$, and are interpreted modulo $\mathscr{T}$. Specifically, we identify t with $1 + ((t-1) \bmod \mathscr{T})$ and use this convention throughout the paper without further mention. Define

$$d_t(C) = \mathbf{1}\{p_t(C) = p^L\}, \qquad y_{a,t}(C) = \sum_{j=1}^{a} d_{t-j}(C).$$

These quantities give the current promotion indicator and the number of past promotions in the periodic extension of C. When $a > \mathscr{T}$, a memory window spans multiple repetitions of the cycle, and each occurrence represents a distinct period of exposure. Total and average cycle revenues are

$$R(C) = \sum_{t=1}^{\mathscr{T}} p_t(C) \sum_{a=0}^{\mathscr{A}} M_a\big(p_t(C), y_{a,t}(C)\big), \qquad \bar{R}(C) = \frac{R(C)}{\mathscr{T}}.$$

Let

$$m = \sum_{t=1}^{\mathscr{T}} d_t(C), \qquad \rho = \frac{m}{\mathscr{T}}$$

denote the number of promotions per cycle and their frequency. The seller's problem is to choose the cycle length, number of promotions, and positions of the promotional periods to maximize $\bar{R}(C)$.

Cyclic rotations change only the starting point of a cycle and have the same average revenue. We therefore treat cycles that differ only by rotation as the same arrangement.

The graph formulation provides a computational benchmark. A maximum mean cycle can be found in time polynomial in the numbers of nodes and edges using the cycle mean algorithm of Karp [17]. However, since $\mathcal{G}_{\mathscr{A}}$ contains $2^{\mathscr{A}+1}$ states, this computation grows exponentially with memory depth.

3 Optimal Promotion Arrangement and Frequency

This section derives sufficient conditions for globally optimal promotion arrangements and frequencies. We first decompose market share function with marginal effects of remembered promotions. These effects

determine the revenue contribution F_b of the number of promotions in a window of length $b+1$. Concavity or convexity of these contributions determines optimal arrangements at fixed cycle length and number of promotions. In the concave case, the optimal promotion frequency can then be found by a finite search over a concave, piecewise linear function. In the convex case, we show that a constant price policy is optimal.

3.1 Representing market share functions

For a cohort of age a, write $M_a^j(y) = M_a(p^j, y)$, where $j \in \{H, L\}$ and $y = 0, \dots, a$. Define the baseline and marginal promotion effects by

$$\alpha_a^j = M_a^j(0), \qquad \Delta_{a,k}^j = M_a^j(k) - M_a^j(k-1), \qquad k = 1, \dots, a.$$

Here, α_a^j denotes the baseline market share, and $\Delta_{a,k}^j$ is the change in purchase probability when remembered promotions increase from $k-1$ to k, holding customer age and the current price fixed.

The following simple lemma will be useful for the revenue decomposition in the next section. It shows that each market share function can be represented as a finite telescoping sum of marginal effects. The proof is immediate from the definition of α_a^j and $\Delta_{a,k}^j$, which can be found in the supplementary appendix, Section B.

Lemma 3.1 (Finite telescoping representation of market share)**.** *For each age a and current price $j \in \{H, L\}$, every conditional market share function $M_a^j : \{0, \dots, a\} \to \mathbb{R}$ has the unique representation*

$$M_a^j(y) = \alpha_a^j + \sum_{k=1}^{a} \Delta_{a,k}^j \mathbf{1}\{y \geq k\} = \alpha_a^j + \sum_{k=1}^{y} \Delta_{a,k}^j, \qquad y = 0, \dots, a,$$

with $\alpha_a^j = M_a^j(0)$ and $\Delta_{a,k}^j = M_a^j(k) - M_a^j(k-1)$.

Positive marginal effects $\Delta_{a,k}^j > 0$ indicate an engagement effect. Decreasing positive effects describe diminishing engagement, while increasing positive effects describe reinforcement. Negative marginal effects, on the other hand, describe a negative historical response, as in a reference-price channel. Marginal effects are defined separately at the two current prices, allowing historical exposure to influence regular-price and promotional-price demand differently.

3.2 Number of promotions and revenue decomposition

Starting from cycle revenue and applying Lemma 3.1 gives

$$R(C) = \sum_{a=0}^{\mathscr{A}} \sum_{t=1}^{\mathscr{T}} \left[p^H \alpha_a^H + D_a(0) d_t + p^H \sum_{k=1}^{a} \Delta_{a,k}^H \mathbf{1}\{y_{a,t} \geq k\} - d_t \sum_{k=1}^{a} r_{a,k} \mathbf{1}\{y_{a,t} \geq k\} \right],$$

where $D_a(0) = p^L \alpha_a^L - p^H \alpha_a^H$ and

$$r_{a,k} = p^H \Delta_{a,k}^H - p^L \Delta_{a,k}^L.$$

The quantity $r_{a,k}$ compares the revenue effects of an additional remembered promotion at the regular and promotional prices.

Define the number of promotions in the window containing period t and the preceding b periods by

$$N_{b,t}(C) = \sum_{i=0}^{b} d_{t-i}(C), \qquad b = 0, \dots, \mathscr{A}.$$

For the derivation, write $x_+ = \max\{x, 0\}$ and

$$V_{b,k}(C) = \sum_{t=1}^{\mathscr{T}} \big(N_{b,t}(C) - k\big)_+.$$

The quantity $V_{b,k}$ sums the excess of the number of promotions in each window above a threshold k. Set $N_{-1,t} = 0$ and $V_{-1,k} = 0$.

The promotion combination counts used below are

$$Q_{b,r}(C) = \sum_{t=1}^{\mathscr{T}} \binom{N_{b,t}(C)}{r}, \qquad b = 0, \dots, \mathscr{A}, \quad r = 0, \dots, b+1.$$

We set $Q_{b,r} = 0$ when $b < 0$, $r < 0$, or $r > b + 1$.

Recall that $y_{a,t} = \sum_{j=1}^{a} d_{t-j}$ counts the past promotions remembered by customers of age a, $y_{a,t} = N_{a-1,t-1}$, cyclic summation gives

$$\sum_t y_{a,t} = am, \qquad \sum_t (y_{a,t} - k)_+ = V_{a-1,k}.$$

Here $m = \sum_{t=1}^{\mathscr{T}} d_t$ is the number of promotional periods in one price cycle, where $d_t = 1$ when the promotional price is charged and 0 otherwise. The corresponding promotion frequency is $\rho = m/\mathscr{T}$.

For binary d_t,

$$d_t \mathbf{1}\{y_{a,t} \geq k\} = (y_{a,t} + d_t - k)_+ - (y_{a,t} - k)_+.$$

Since $y_{a,t} + d_t = N_{a,t}$, summing over the cycle yields

$$\sum_t d_t \mathbf{1}\{y_{a,t} \geq k\} = V_{a,k} - V_{a-1,k}.$$

These identities account for both the promotion history and the price at which its demand effect is monetized.

The regular-price historical response can also be written as

$$\sum_{k=1}^{a} \Delta^H_{a,k} \mathbf{1}\{y \geq k\} = \Delta^H_{a,1} y + \sum_{k=1}^{a-1} (\Delta^H_{a,k+1} - \Delta^H_{a,k})(y - k)_+, \qquad a \geq 1.$$

Substituting this expression and the cycle identities above, then collecting terms with the same $V_{b,k}$, gives the following decomposition. The marginal effects $\Delta^j_{a,k}$ describe demand, while the coefficients ψ below describe changes in the marginal revenue contribution of promotions within a window.

Proposition 3.2 (Revenue decomposition by window contributions)**.** *Under Assumption 2.1, every cycle C with $\mathscr{T}$ periods and m promotions has revenue*

$$R(C) = \underbrace{\mathscr{T} g_0 + m g_1}_{K_{\mathscr{T},m}} + \sum_{b=1}^{\mathscr{A}} \sum_{t=1}^{\mathscr{T}} F_b\big(N_{b,t}(C)\big),$$

where

$$F_b(n) = \sum_{k=1}^{b} (n-k)_+ \psi_{b,k}, \qquad n = 0, \dots, b+1,$$

$$\psi_{b,k} = \begin{cases} p^H(\Delta^H_{b+1,k+1} - \Delta^H_{b+1,k}) + r_{b+1,k} - r_{b,k}, & 1 \le k \le b < \mathscr{A}, \\ -r_{\mathscr{A},k}, & 1 \le k \le b = \mathscr{A}, \end{cases}$$

and

$$g_0 = p^H \sum_{a=0}^{\mathscr{A}} \alpha_a^H,$$

$$g_1 = \sum_{a=0}^{\mathscr{A}} \left(p^H a \Delta^H_{a,1} + p^L \alpha_a^L - p^H \alpha_a^H\right).$$

The term $a\Delta^H_{a,1}$ is interpreted as zero at age zero. In particular,

$$F_b(0) = F_b(1) = 0, \qquad F_b(k+1) - 2F_b(k) + F_b(k-1) = \psi_{b,k}.$$

The function F_b is the promotion interaction contribution of a window of length $b+1$, after the baseline terms related to g_0, g_1 have been collected in the $K_{\mathscr{T},m}$. Concavity of the function F_b means that adding another promotion contributes less to this term when the number of promotions in the window increases. The coefficient $\psi_{b,k}$ measures the change in that additional contribution, and the concavity of F_b is equivalent to $\psi_{b,k} \le 0$ for every k.

3.3 Optimal arrangement structure

With the revenue decomposition in Proposition 3.2, we can now analyze the optimal arrangement of promotions for a given cycle length and number of promotions. The following definitions introduce two benchmark arrangements that have nice properties in terms of $N_{b,t}$ they induce and will be used throughout the paper.

Definition 3.1 (Balanced cycle)**.** Let C be a cycle of $\mathscr{T}$ periods with promotion indicators (d_t) and $m = \sum_{t=1}^{\mathscr{T}} d_t$ promotions. The cycle is *balanced* if, for every window length $k = 1, \dots, \mathscr{A}+1$ and every starting index $u \in \mathbb{Z}$,

$$\sum_{t=u+1}^{u+k} d_t \in \left\{ \left\lfloor \frac{km}{\mathscr{T}} \right\rfloor, \left\lceil \frac{km}{\mathscr{T}} \right\rceil \right\},$$

where $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ denote the floor and ceiling functions, respectively.

Definition 3.2 (Single block cycle)**.** Let C be a cycle of $\mathscr{T}$ periods with m promotions. For $1 \le m \le \mathscr{T} - 1$, the cycle is a *single block cycle* if, up to cyclic rotation, $d_1 = \cdots = d_m = 1$ and $d_{m+1} = \cdots = d_{\mathscr{T}} = 0$. Thus, the promotional periods form a single promotion block and the regular price periods form a single block. The constant cycles with $m = 0$ or $m = \mathscr{T}$ are included as boundary cases.

Balance requires each window count to equal either the floor or the ceiling of its mean $km/\mathscr{T}$. Since for any window of length k, the number of promotions in that window differs by at most one as we move the window along the cycle, the balanced cycle spreads promotions as evenly as possible across the cycle relevant to customer memory. On the other hand, a single block cycle places all promotions in a single block for a fixed cycle length and number of promotions.

Example 3.1 (Balanced and single block cycles)**.** Let $\mathscr{T} = 5$, $m = 2$, and $\mathscr{A} = 3$. The following cycles have the same promotion frequency $2/5$.

Arrangement	Period 1	Period 2	Period 3	Period 4	Period 5
Balanced	p^H	p^H	p^L	p^H	p^L
Single block	p^L	p^L	p^H	p^H	p^H

In the balanced cycle, windows of length one or two contain either zero or one promotion, while windows of length three or four contain either one or two promotions. Thus, for each length up to $\mathscr{A} + 1 = 4$, the number of promotions differs across windows by at most one. In the single block cycle, windows of length two can contain zero, one, or two promotions, reflecting the concentration of exposure within the campaign.

It is clear that a single block cycle always exists. Next we show by construction that a balanced cycle also exists for every feasible cycle length and number of promotions.

Lemma 3.3 (Existence of balanced cycles)**.** *For every integer $\mathscr{T} \geq 1$ and $m \in \{0, \dots, \mathscr{T}\}$, the binary sequence*

$$d_t = \left\lfloor \frac{tm}{\mathscr{T}} \right\rfloor - \left\lfloor \frac{(t-1)m}{\mathscr{T}} \right\rfloor, \qquad t \in \mathbb{Z},$$

is periodic with period $\mathscr{T}$, is balanced, and contains exactly m promotions per cycle.

For each $(\mathscr{T}, m)$, the single block arrangement is unique up to rotation. The balanced arrangement is also unique up to rotation when $\mathscr{T} \leq \mathscr{A} + 1$. For longer cycles, distinct arrangements can satisfy the same balance requirements over the relevant windows. Proofs of the existence and uniqueness results appear in Section B of the supplementary appendix.

The following lemma compares sums of convex or concave functions of the numbers of promotions in moving windows under these two arrangements.

Lemma 3.4. *Fix $(\mathscr{T}, m)$. For every depth $b = 0, \dots, \mathscr{A}$ and every discretely convex function ϕ on $\{0, \dots, b+1\}$, every balanced cycle minimizes*

$$\sum_{t=1}^{\mathscr{T}} \phi\big(N_{b,t}(C)\big),$$

and a single block cycle maximizes it. For a discretely concave function the comparisons are reversed. In particular, these comparisons apply to the window revenue contribution functions F_b.

The proof appears in the supplementary appendix, Section B.

Substituting these simultaneous extrema into the revenue decomposition gives sufficient conditions for optimal arrangements.

Theorem 3.5 (Optimal arrangement from window contributions)**.** *Fix $(\mathscr{T}, m)$ and compare cycles with $\mathscr{T}$ periods and m promotions.*

(i) If every F_b is discretely concave, equivalently $\psi_{b,k} \leq 0$ for $1 \leq k \leq b \leq \mathscr{A}$, then every balanced cycle is optimal. If additionally $\mathscr{T} \leq \mathscr{A} + 1$ and $\psi_{b,k} < 0$ for $1 \leq k \leq b, \quad b = 1, \dots, \mathscr{T} - 2$, the balanced arrangement is uniquely optimal up to cyclic rotation.

(ii) If every F_b is discretely convex, equivalently $\psi_{b,k} \geq 0$ for every such depth and exposure index, then a single block cycle is optimal.

Proof. At fixed $(\mathscr{T}, m)$, arrangement affects revenue only through $\sum_{b,t} F_b(N_{b,t})$. A balanced cycle maximizes the sum for each concave F_b, while a single block cycle maximizes it for each convex F_b. Lemma 3.4 therefore

establishes parts (i) and (ii). For uniqueness under part (i), suppose $\mathscr{T} \leq \mathscr{A} + 1$ and the stated strict inequalities hold. As shown in the proof of Lemma 3.4 in Section B of the supplementary appendix, every nonbalanced cycle has two windows of some length $b + 1$, with $1 \leq b \leq \mathscr{T} - 2$, whose numbers of promotions differ by at least two. Strict concavity of the corresponding F_b makes its total contribution strictly smaller than in a balanced cycle, while concavity gives weak comparisons at all other depths. Hence every nonbalanced cycle has strictly lower revenue. Uniqueness of balanced cycles up to rotation completes the argument. □

Table 1 summarizes how the shape of F_b determines arrangement and frequency.

Table 1: Structural implications of the window revenue contributions.

Window contributions	Arrangement at fixed $(\mathscr{T}, m)$	Joint arrangement and frequency choice
All concave	Balanced	Concave frequency objective with an optimal period of at most $\mathscr{A} + 1$
All convex	A single promotion block	A constant price is optimal
Neither condition	Tradeoffs among window contributions	Balanced and single block policies have loss bounds, and memory truncation approximates the optimal policy

When the concavity or convexity conditions fail, several separated promotion blocks can be optimal. Set $\mathscr{T} = 12$, $m = 4$, $\mathscr{A} = 5$, $p^L = 1$, and $p^H = 1.8$, and consider the threshold market share function

$$M_a(p, y) = 0.04 + 0.48(0.85)^a \mathbf{1}\{y + \mathbf{1}\{p = p^L\} \geq 2\}.$$

Enumerating all 495 feasible cycles yields an optimal cycle

$$d = (0, 0, 0, 0, 1, 1, 0, 0, 0, 0, 1, 1).$$

Its promotions form two pairs separated by four regular price periods, and it earns more revenue than either benchmark arrangement. The complete comparison across five market share functions appears in Section B of the supplementary appendix.

3.4 Optimal frequency for balanced arrangements

Balanced cycles with the same promotion frequency generate the same average revenue. When every F_b is concave, this property reduces the seller's problem to maximizing a concave function of frequency over a finite candidate set.

For $\rho \in [0, 1]$, define

$$n_b(\rho) := \lfloor (b + 1)\rho \rfloor, \qquad \xi_b(\rho) := (b + 1)\rho - n_b(\rho).$$

In a balanced cycle with rational frequency $\rho = m/\mathscr{T}$, windows of length $b + 1$ contain either $n_b(\rho)$ or $n_b(\rho) + 1$ promotions. Their average count is $(b + 1)\rho$, so the proportion containing the larger count is $\xi_b(\rho)$. For example, an average count of 2.3 means that 70% of the windows contain two promotions and 30% contain three.

For each admissible depth b and order $r \geq 2$, define

$$q_{b,r}(\rho) := \big(1 - \xi_b(\rho)\big)\binom{n_b(\rho)}{r} + \xi_b(\rho)\binom{n_b(\rho) + 1}{r}.$$

At a rational frequency, this expression gives the average promotion combination count of any balanced cycle,

$$q_{b,r}(\rho) = \frac{Q_{b,r}(C^{\mathrm{bal}})}{\mathscr{T}}.$$

The same formula defines a continuous interpolation over all $\rho \in [0,1]$.

Proposition 3.6 (Frequency representation for balanced cycles)**.** *For any balanced cycle C^{bal} with $\mathscr{T}$ periods and m promotions, average revenue depends on $(\mathscr{T}, m)$ only through $\rho = m/\mathscr{T}$. Specifically,*

$$\bar{R}(C^{\mathrm{bal}}) = g(\rho), \qquad g(\rho) = g_0 + g_1\rho + \sum_{b=1}^{\mathscr{A}} \widetilde{F}_b\big((b+1)\rho\big), \tag{1}$$

where $\widetilde{F}_b$ is the linear interpolation of F_b between consecutive integers, and g_0, g_1 are defined in Proposition 3.2.

The constant term g_0 is average revenue under the regular price. The term $g_1\rho$ collects the contributions linear in promotion frequency, while the remaining terms capture the window revenue contributions. When all F_b's are concave, Theorem 3.5 makes a balanced cycle optimal at every fixed $(\mathscr{T}, m)$. Maximizing g therefore determines both an optimal frequency and an optimal arrangement.

Proposition 3.7 (Concavity and short cycles)**.** *Suppose F_b is discretely concave on $\{0, \dots, b+1\}$ for every $b = 1, \dots, \mathscr{A}$. Then g is continuous, piecewise linear, and concave on $[0,1]$. A maximizer of g can be chosen from*

$$\mathscr{B} = \{0,1\} \cup \left\{ \frac{j}{b+1} \,:\, b = 1, \dots, \mathscr{A},\ j = 1, \dots, b \right\}.$$

Consequently, there exists a balanced cycle with period at most $\mathscr{A}+1$ that is optimal over all feasible infinite price paths.

The result gives a procedure for selecting the cycle. Evaluate g at every frequency in $\mathscr{B}$, find a maximizer $\rho^\star = m^\star/\mathscr{T}^\star$ and construct the corresponding balanced cycle using Lemma 3.3. Given the window contribution tables, the direct search takes $O(\mathscr{A}^3)$ arithmetic operations and $O(\mathscr{A}^2)$ space for $\mathscr{A} \geq 1$. These are unit cost arithmetic bounds. Algorithm 1 in the supplementary appendix gives the procedure.

3.5 Constant price optimality under convex window contributions

When every F_b is convex, Theorem 3.5 makes a single promotion block optimal at prescribed cycle length and number of promotions. Allowing the number of promotions to vary gives the following constant-price result.

Under the constant regular price, every cohort remembers zero promotions. Under the constant promotional price, an age a cohort remembers a promotions. Let D denote the average revenue under the constant promotional price minus that under the constant regular price,

$$D := p^L \sum_{a=0}^{\mathscr{A}} M_a^L(a) - p^H \sum_{a=0}^{\mathscr{A}} M_a^H(0).$$

Proposition 3.8 (Constant price optimality)**.** *Suppose every F_b is discretely convex, equivalently $\psi_{b,k} \geq 0$ for $1 \leq k \leq b \leq \mathscr{A}$. If $D \leq 0$, the constant regular price is optimal over all feasible infinite price paths. If $D \geq 0$, the constant promotional price is optimal over all feasible infinite price paths. Both constant prices are optimal when $D = 0$.*

The result follows from an upper bound on the revenue of every cycle. Discrete convexity bounds each window contribution by the chord joining its endpoint values,

$$F_b(n) \le \left(1 - \frac{n}{b+1}\right) F_b(0) + \frac{n}{b+1} F_b(b+1).$$

Averaging and using $\sum_t N_{b,t} = (b+1)m$ gives

$$\bar{R}(C) \le (1-\rho)p^H \sum_{a=0}^{\mathscr{A}} M_a^H(0) + \rho p^L \sum_{a=0}^{\mathscr{A}} M_a^L(a).$$

The right side is a weighted average of the revenues under the two constant prices. It is therefore bounded above by the revenue of the better constant price. Proposition 2.1 extends this comparison to all feasible infinite price paths. The full proof appears in Section B.

When $D \neq 0$, every cycle with $0 < \rho < 1$ earns strictly less than the better constant price. When $D = 0$, both constant prices are optimal and a nonconstant cycle cannot exceed the constant price revenue.

3.6 Heterogeneous active lifetimes

The revenue decomposition and optimality conditions can be extended to the case of customer heterogeneity through population weighted coefficients. Suppose each arriving cohort contains fixed proportions of finitely many customer types with different active lifetimes and market share functions, with full-history memory for every type satisfying Assumption 2.1. All types face the same price path. Aggregate revenue retains the decomposition in Proposition 3.2, with each coefficient averaged across types using their cohort proportions. The arrangement and frequency results therefore apply whenever the corresponding concavity or convexity conditions hold for the population weighted window contributions, even when those conditions fail for some individual types. Under concave aggregate window contributions, an optimal balanced cycle has period at most one plus the largest maximum active age across types. The formal statement and proof appear in Section B of the supplementary appendix.

4 Additively Separable Market Share

Additively separable market share distinguishes the effect of the current promotion from the response to remembered past promotions. For each age a, consider

$$M_a(p, y) = \alpha_a + \eta_a \mathbf{1}\{p = p^L\} + f_a(y), \qquad y = 0, \dots, a, \qquad f_a(0) = 0. \tag{2}$$

Here α_a is market share at the regular price with no remembered promotions, and $\eta_a \ge 0$ is the current promotion effect. The historical response f_a is the same at either current price and may be nonlinear. Throughout this section, parameters give market shares in $[0, 1]$ for every feasible history.

We organize the analysis around two mechanisms. The engagement channel raises demand through remembered promotions, while the reference price channel can lower demand because past discounts reduce the benchmark used to evaluate the current price. Section 4.1 gives sufficient conditions for balanced optimality under additive engagement, illustrates them with simple specifications, and then allows historical effects to differ across lags. Section 4.2 studies the reference price channel and establishes conditions for single block and constant price optimality.

4.1 Engagement channel

For the additive model in (2), define the marginal engagement effect of the kth remembered promotion by

$$\Delta_{a,k} = f_a(k) - f_a(k-1), \qquad 1 \le k \le a \le \mathscr{A}.$$

Assumption 4.1 (Diminishing additive engagement)**.** For every age a, the market share function has the form (2), with

$$\alpha_a, \eta_a \geq 0, \qquad \alpha_a + \eta_a + f_a(a) \leq 1.$$

Historical marginal effects are nonnegative and weakly decrease with exposure and age,

$$\begin{aligned} \Delta_{a,k} &\geq 0, && 1 \leq k \leq a \leq \mathscr{A}, \\ \Delta_{a,k+1} &\leq \Delta_{a,k}, && 1 \leq k < a \leq \mathscr{A}, \\ \Delta_{a+1,k} &\leq \Delta_{a,k}, && 1 \leq k \leq a < \mathscr{A}. \end{aligned}$$

Remembered promotions are substitutes in generating engagement because the marginal demand gain does not increase with prior exposure. The age condition requires this marginal gain to weaken as the customer relationship matures, holding the number of remembered promotions fixed. Nonnegative marginal effects and the upper bound ensure valid market shares throughout the engagement model.

4.1.1 Optimal arrangement and frequency

Additive separability gives $\Delta^H_{a,k} = \Delta^L_{a,k} = \Delta_{a,k}$, so $r_{a,k} = \Delta p\, \Delta_{a,k}$. Proposition 3.2 therefore yields

$$\psi_{b,k} = \begin{cases} p^H(\Delta_{b+1,k+1} - \Delta_{b+1,k}) + \Delta p(\Delta_{b+1,k} - \Delta_{b,k}), & 1 \leq k \leq b < \mathscr{A}, \\ -\Delta p\, \Delta_{\mathscr{A},k}, & 1 \leq k \leq b = \mathscr{A}. \end{cases}$$

The exposure and age conditions make the two interior terms nonpositive. Nonnegative engagement gives the terminal sign. Thus every window contribution F_b is concave. Assumption 4.1 provides behavioral conditions sufficient for this conclusion, while the general arrangement theorem requires only $\psi_{b,k} \leq 0$.

Corollary 4.1 (Optimal additive promotion policies)**.** *Under Assumption 4.1, every balanced cycle maximizes revenue at fixed $(\mathscr{T}, m)$. Its average revenue is a continuous, concave, piecewise linear function of promotion frequency. Maximizing over the candidate frequencies $\mathscr{B}$ in Proposition 3.7 gives a balanced cycle with period at most $\mathscr{A} + 1$ that is optimal over all feasible infinite price paths.*

The frequency objective is (1), with

$$g_0 = p^H \sum_{a=0}^{\mathscr{A}} \alpha_a, \qquad g_1 = \sum_{a=0}^{\mathscr{A}} (p^L \eta_a - \Delta p\, \alpha_a) + p^H \sum_{a=1}^{\mathscr{A}} a \Delta_{a,1}.$$

Current promotion effects η_a enter only g_1. They can therefore change optimal frequency without changing arrangement rankings at fixed $(\mathscr{T}, m)$. With nonlinear historical responses, rearranging promotions may also change total expected demand. The proof and frequency derivation appear in Section C.1.

The following examples illustrate Assumption 4.1. The first gives a family of historical responses that separate age from cumulative exposure. The next two examine its saturating and affine specializations.

Example 4.1 (Age and exposure separability)**.** Let

$$f_a(y) = c_a h(y), \qquad c_0 = 0, \qquad c_1 \geq \cdots \geq c_{\mathscr{A}} \geq 0, \qquad h(0) = 0,$$

where h is nondecreasing and discretely concave on $\{0, \dots, \mathscr{A}\}$. Then

$$\Delta_{a,k} = c_a\big(h(k) - h(k-1)\big)$$

is nonnegative and weakly decreases in both k and a. Assumption 4.1 holds whenever $\alpha_a, \eta_a \geq 0$ and $\alpha_a + \eta_a + c_a h(a) \leq 1$.

Example 4.2 (Saturating engagement)**.** In Example 4.1, take

$$f_a(y) = c_a(1 - q^y), \qquad 0 < q < 1.$$

Here c_a bounds the potential increase in market share due to historical engagement. A simple mechanism is that each promotion realizes a fraction $1 - q$ of the remaining potential lift. Early reminders may bring customers back, while repeated reminders increasingly reach customers who are already engaged.

After $k - 1$ remembered promotions, the remaining potential lift is $c_a q^{k-1}$, so the next promotion adds

$$\Delta_{a,k} = c_a(1 - q)q^{k-1}.$$

These gains weakly decrease with both exposure and age when c_a is nonincreasing. A declining c_a represents less scope for promotions to increase engagement as the customer relationship matures. Under the market share bounds in Example 4.1, Assumption 4.1 therefore holds and Corollary 4.1 applies.

Example 4.3 (Affine engagement)**.** The choice $h(y) = y$ in Example 4.1 gives the affine specification. Write $c_a = \beta_a^{\text{aff}}$, $\alpha_a = \alpha_a^{\text{aff}}$, and $\eta_a = \eta_a^{\text{aff}}$. Marginal engagement is then constant in exposure, so Assumption 4.1 reduces to nonnegative coefficients, valid market shares, and

$$\beta_1^{\text{aff}} \geq \cdots \geq \beta_{\mathscr{A}}^{\text{aff}} \geq 0.$$

Setting $\beta_{\mathscr{A}+1}^{\text{aff}} = 0$ gives $\psi_{b,k} = -\Delta p(\beta_b^{\text{aff}} - \beta_{b+1}^{\text{aff}}) \leq 0$.

Affine engagement also gives a simple interpretation of the arrangement result. Rearranging promotions at fixed $(\mathscr{T}, m)$ preserves total expected demand, so balanced cycles minimize demand served at the promotional price. Write g^{aff} for their average revenue as a function of frequency. Figure 1 illustrates this objective for three slope profiles. Its explicit formula, the demand identity, and an $O(\mathscr{A}^2 + 1)$ frequency search appear in Section C.1. Example C.1 shows how current promotion effects can change the optimal frequency.

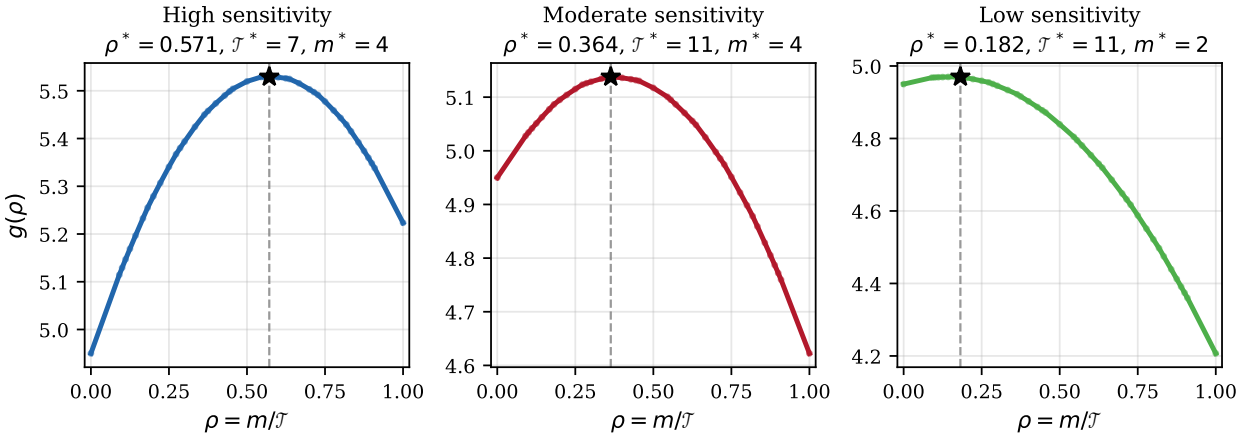


Figure 1: Promotion frequency in an affine engagement example. The balanced revenue $g^{\text{aff}}(\rho)$ is continuous and piecewise linear, and the stars mark the maximizing promotion frequency $\rho^\star$. The parameters are $\mathscr{A} = 10$, $p^L = 1$, $p^H = 1.8$, $\alpha_a^{\text{aff}} = 0.25$, and $\eta_a^{\text{aff}} = \beta_a^{\text{aff}}$. The high, moderate, and low sensitivity profiles are $\beta_a^{\text{aff}} = 0.1(0.85)^a$, $0.0625(0.88)^a$, and $0.0375(0.92)^a$, respectively.

4.1.2 Heterogeneous lag effects

The affine count model assigns the same marginal effect to every remembered promotion. We now allow engagement effects to differ across lags, so promotions one and two periods earlier can have different influences on current demand. This extension distinguishes histories with the same promotion count, so Assumption 4.1 no longer applies directly. The pairwise revenue representation nevertheless gives a separate condition for balanced optimality.

A related formulation appears in reference price research, where exponential smoothing assigns declining weights to past prices, as in Hardie et al. [16], Kopalle et al. [18], and Popescu and Wu [26]. Individual cohort coefficients in our model need not decline with lag. The following result establishes balanced optimality under a discrete convexity condition on the aggregate lag coefficients.

Consider the market share function

$$M_a^{\mathrm{loc}}(d_t, \dots, d_{t-a}) = \alpha_a^{\mathrm{loc}} + \beta_{a,0}^{\mathrm{loc}} d_t + \sum_{j=1}^{a} \beta_{a,j}^{\mathrm{loc}} d_{t-j}. \tag{3}$$

The coefficient $\beta_{a,0}^{\mathrm{loc}}$ measures the current promotion effect, while $\beta_{a,j}^{\mathrm{loc}}$ measures the effect of a promotion j periods earlier. Assume all coefficients are nonnegative and

$$\alpha_a^{\mathrm{loc}} + \sum_{j=0}^{a} \beta_{a,j}^{\mathrm{loc}} \le 1,$$

so market shares remain in $[0, 1]$.

The arrangement condition depends on historical effects aggregated across active cohorts. Define

$$\bar{\beta}_j^{\mathrm{loc}} = \sum_{\{a\,:\,a \ge j\}} \beta_{a,j}^{\mathrm{loc}}, \qquad j = 1, \dots, \mathscr{A},$$

and set

$$\bar{\beta}_{\mathscr{A}+1}^{\mathrm{loc}} = \bar{\beta}_{\mathscr{A}+2}^{\mathrm{loc}} = 0.$$

Proposition 4.2 (Balanced optimality with heterogeneous lag effects)**.** *If*

$$\bar{\beta}_b^{\mathrm{loc}} - 2\bar{\beta}_{b+1}^{\mathrm{loc}} + \bar{\beta}_{b+2}^{\mathrm{loc}} \ge 0, \qquad b = 1, \dots, \mathscr{A}, \tag{4}$$

then every balanced cycle maximizes revenue among cycles with the same length and number of promotions.

The result follows by expressing the revenue effects of specific lags through promotion combination counts of order two. Write $Q_b(C) = Q_{b,2}(C)$ and let

$$P_j(C) = \sum_{t=1}^{\mathscr{T}} d_t(C) d_{t-j}(C)$$

count promotional periods preceded by another promotion j periods earlier. At fixed $(\mathscr{T}, m)$, the arrangement component of revenue is

$$-\Delta p \sum_{j=1}^{\mathscr{A}} \bar{\beta}_j^{\mathrm{loc}} P_j(C).$$

Using

$$P_b(C) = Q_b(C) - 2Q_{b-1}(C) + Q_{b-2}(C), \qquad Q_0(C) = Q_{-1}(C) = 0,$$

the coefficient on $Q_b(C)$ becomes

$$-\Delta p\left(\bar{\beta}^{\mathrm{loc}}_b - 2\bar{\beta}^{\mathrm{loc}}_{b+1} + \bar{\beta}^{\mathrm{loc}}_{b+2}\right).$$

The resulting window contribution is $F_b(n) = -\Delta p(\bar{\beta}^{\mathrm{loc}}_b - 2\bar{\beta}^{\mathrm{loc}}_{b+1} + \bar{\beta}^{\mathrm{loc}}_{b+2})\binom{n}{2}$. Its second difference is nonpositive under the stated condition, so the window concavity argument of Theorem 3.5 gives balanced optimality. Details appear in Section C.1.3.

This condition recovers the affine case of Assumption 4.1 when historical coefficients do not depend on lag. If $\beta^{\mathrm{loc}}_{a,j} = \beta^{\mathrm{aff}}_a$ for $j \geq 1$, then $\bar{\beta}^{\mathrm{loc}}_j = \sum_{a=j}^{\mathscr{A}} \beta^{\mathrm{aff}}_a$, and

$$\bar{\beta}^{\mathrm{loc}}_b - 2\bar{\beta}^{\mathrm{loc}}_{b+1} + \bar{\beta}^{\mathrm{loc}}_{b+2} = \beta^{\mathrm{aff}}_b - \beta^{\mathrm{aff}}_{b+1}.$$

Thus declining historical slopes are the count model counterpart of the aggregate lag convexity condition.

The finite memory cutoff affects the convexity condition. For a geometric profile $\bar{\beta}^{\mathrm{loc}}_b = cr^b$, with $c > 0$, $0 < r < 1$, and $\mathscr{A} \geq 2$, the interior second differences are nonnegative. The condition at $b = \mathscr{A} - 1$ additionally requires

$$\bar{\beta}^{\mathrm{loc}}_{\mathscr{A}-1} \geq 2\bar{\beta}^{\mathrm{loc}}_{\mathscr{A}},$$

which holds when $r \leq 1/2$.

4.2 Reference price channel

The reference price channel has a different historical effect. A past discount lowers the price against which customers evaluate the current offer and can reduce demand at a given current price. We study a linear, loss neutral model and derive its implications for promotion arrangement and frequency.

For $a \geq 1$, let

$$r_{a,t} = \sum_{j=1}^{a} w_{a,j} p_{t-j}, \qquad w_{a,j} \geq 0, \qquad \sum_{j=1}^{a} w_{a,j} = 1,$$

and consider

$$M_{a,t} = u_a - v_a p_t + \gamma_a(r_{a,t} - p_t), \qquad v_a, \gamma_a \geq 0. \tag{5}$$

The market share function is loss neutral because the same coefficient γ_a applies to gains and losses relative to the reference price. At age zero, omit the reference price term and set $\gamma_0 = 0$. Parameters are chosen so that market shares lie in $[0, 1]$ for every feasible price history. Unequal reference weights distinguish histories with the same promotion count, giving a lag dependent affine model as in Section 4.1.2.

With the average reference price, $w_{a,j} = 1/a$, we have $r_{a,t} = p^H - \Delta p\, y_{a,t}/a$. The market share function then has the additive form (2), with $\alpha_a = u_a - v_a p^H$, $\eta_a = (v_a + \gamma_a)\Delta p$, and

$$f_a(y) = \beta^{\mathrm{aff}}_a y, \qquad \beta^{\mathrm{aff}}_a = -\frac{\gamma_a \Delta p}{a}, \qquad a = 1, \ldots, \mathscr{A}.$$

A strictly negative historical slope violates the engagement assumption. Suppose γ_a/a weakly decreases with age, as it does when reference sensitivity is constant across ages. Then

$$\beta^{\mathrm{aff}}_1 \leq \cdots \leq \beta^{\mathrm{aff}}_{\mathscr{A}} \leq \beta^{\mathrm{aff}}_{\mathscr{A}+1} = 0,$$

and hence

$$\psi_{b,k} = -\Delta p\big(\beta^{\mathrm{aff}}_b - \beta^{\mathrm{aff}}_{b+1}\big) \geq 0, \qquad 1 \leq k \leq b \leq \mathscr{A}.$$

Every window contribution F_b is therefore convex.

Proposition 4.3 (Reference price benchmark)**.** *Suppose reference prices equal average past prices and γ_a/a weakly decreases with age. A single promotion block maximizes revenue at any fixed cycle length and number of promotions. When frequency is unrestricted, the better constant price is optimal over all feasible infinite price paths.*

The result follows from Theorem 3.5 and Proposition 3.8. Grouping promotions reduces the adverse effect of earlier discounts on demand during regular price periods. When frequency can also be chosen, a constant price weakly dominates every promotion cycle. The proof appears in Section C.2 of the supplementary appendix.

This constant price benchmark connects our arrangement analysis to classical reference price findings. Kopalle et al. [18] derive conditions for constant or cyclic pricing under asymmetric responses to gains and losses. Popescu and Wu [26] establish convergence to a constant price for loss neutral and loss averse demand under their assumptions, with continuously chosen prices and discounted revenue. Proposition 4.3 establishes constant price optimality under the stated reference price conditions, two given prices, and long run average revenue.

These arrangement and constant price conclusions also hold for reference prices smoothed exponentially within each cohort from a fixed initial reference, with common reference sensitivity across ages, by the same window convexity argument. The proof appears in Section C.2 of the supplementary appendix.

For general nonnegative reference weights summing to one, the better constant price remains optimal even without the age monotonicity condition. Equation (C.3) expresses average cycle revenue as the frequency weighted average of the two constant price revenues minus a nonnegative penalty for price variation. This frequency conclusion does not require convexity of every window contribution and does not by itself establish single block optimality at fixed frequency.

5 Multiplicative Market Share

We next study a multiplicative market share function as another special case of the general model. Multiplicative demand formulations have a long history in promotion research, including the analysis of dynamic promotion effects by van Heerde et al. [33]. Cohen et al. [9] use a product of current and historical price effects in promotion planning. These formulations provide a foundation for studying how proportional effects on demand accumulate across periods.

A practical setting is a subscription or repeat purchase service in which customers may churn before the maximum age $\mathscr{A}$. Customers arrive engaged. In each period, an engaged customer renews and purchases at price p with probability $s(p)$, and otherwise churns permanently. This gives a Markov chain with two states, engaged and churned, with churned absorbing. Promotions increase retention, so we write $s(p^L) = s_1$ and $s(p^H) = s_2$, where $0 < s_2 < s_1 < 1$.

A customer arriving at time $t-a$ purchases at time t only if every renewal through age a succeeds. The resulting market share is

$$M_a^{\mathrm{mul}}(p_t, y_{a,t}) = \prod_{j=0}^{a} s(p_{t-j}) = s(p_t) s_1^{y_{a,t}} s_2^{a-y_{a,t}}.$$

The historical factors describe survival up to the current period, and $s(p_t)$ is the probability of the current renewal. Repeated promotions reinforce demand by preserving a larger customer base for subsequent renewals.

Define

$$\theta = \frac{s_1}{s_2}, \qquad \kappa = \frac{p^H}{p^L}.$$

Replacing a remembered regular price with a promotional price multiplies market share by $\theta > 1$. The ratio $\kappa > 1$ measures the regular price relative to the promotional price. These two ratios determine the curvature of the window revenue contributions.

5.1 Multiplicative window contributions

For multiplicative market share, the marginal effect of the kth remembered promotion is

$$\Delta^H_{a,k} = s_2^{a+1}(\theta - 1)\theta^{k-1}, \qquad \Delta^L_{a,k} = \theta\Delta^H_{a,k}, \qquad 1 \le k \le a \le \mathscr{A}.$$

Substituting these effects into Proposition 3.2 gives the window contributions and their second differences. Cohort arrival mass is one, and retention is already incorporated in M_a^{mul}.

Corollary 5.1 (Multiplicative window contributions)**.** *For $b = 1, \dots, \mathscr{A}$, define*

$$S_b^{\mathrm{mul}}(\kappa) = \begin{cases} s_1(1 - s_2) - \kappa s_2(1 - s_1), & 1 \le b < \mathscr{A}, \\ s_1 - \kappa s_2, & b = \mathscr{A}. \end{cases}$$

The window revenue contributions are

$$F_b(n) = p^L s_2^b S_b^{\mathrm{mul}}(\kappa)\frac{\theta^n - 1 - n(\theta - 1)}{\theta - 1}, \qquad n = 0, \dots, b + 1,$$

and their second differences satisfy

$$\psi_{b,k} = F_b(k + 1) - 2F_b(k) + F_b(k - 1) = p^L s_2^b(\theta - 1)\theta^{k-1} S_b^{\mathrm{mul}}(\kappa), \qquad k = 1, \dots, b.$$

The factor multiplying $S_b^{\mathrm{mul}}(\kappa)$ in $\psi_{b,k}$ is positive. Thus, at each depth b, all second differences have the same sign across promotion counts k. Thus F_b is concave when $S_b^{\mathrm{mul}}(\kappa) \le 0$ and convex when $S_b^{\mathrm{mul}}(\kappa) \ge 0$. These are the shape conditions in Theorem 3.5. The derivation appears in Section D.1 of the supplementary appendix.

5.2 Optimal arrangement regimes

The price and retention ratios determine whether all window contributions are convex, all are concave, or their second differences have mixed signs. These regimes provide sufficient conditions for single block and balanced optimality at a prescribed cycle length and number of promotions.

If $\mathscr{A} = 0$, there are no terms that depend on arrangement, and all arrangements with the same $(\mathscr{T}, m)$ earn the same revenue. For $\mathscr{A} \ge 2$, define

$$\tau^{\mathrm{mul}} = \frac{s_1(1 - s_2)}{s_2(1 - s_1)}.$$

The threshold $\theta = s_1/s_2$ compares the retention gain from a promotion with its price reduction. The inequality $\kappa \le \theta$ is equivalent to $p^L s_1 \ge p^H s_2$. A promotion then weakly increases expected current revenue from an engaged customer while improving retention. The interior threshold also reflects churn reduction through the identity

$$\tau^{\mathrm{mul}} = \theta\frac{1 - s_2}{1 - s_1} = \frac{s_1/(1 - s_1)}{s_2/(1 - s_2)}.$$

It is the ratio of retention odds under the two prices, and $\tau^{\mathrm{mul}} > \theta$. At the terminal age, no further renewal remains, so the threshold reduces to the immediate revenue comparison θ.

Proposition 5.2 (Multiplicative arrangement regimes). *Fix a cycle length $\mathscr{T}$ and number of promotions m, and suppose $\mathscr{A} \geq 1$.*

(i) If $\kappa \leq \theta$, every F_b is discretely convex, and a single block cycle is optimal.

(ii) Every F_b is discretely concave if and only if $\kappa \geq \tau^{\mathrm{mul}}$ when $\mathscr{A} \geq 2$, or $\kappa \geq \theta$ when $\mathscr{A} = 1$. Under these conditions, every balanced cycle is optimal.

(iii) For the remaining parameter values, $\mathscr{A} \geq 2$ and $\theta < \kappa < \tau^{\mathrm{mul}}$. The interior contributions F_b, $b < \mathscr{A}$, are strictly convex, while $F_{\mathscr{A}}$ is strictly concave.

When $\kappa \leq \theta$, the retention gain covers the immediate price reduction, and convex window contributions favor concentrating promotions. For $\mathscr{A} \geq 2$, when $\kappa \geq \tau^{\mathrm{mul}}$, concave window contributions favor spreading promotions because concentrating discounts on a customer base already enlarged by earlier promotions becomes less attractive. These are the two cases of Theorem 3.5. In the region with mixed second-difference signs $\theta < \kappa < \tau^{\mathrm{mul}}$, a promotion reduces immediate expected revenue but retention benefits make the interior window contributions convex, while the terminal window contribution is concave. The resulting tradeoff requires evaluation of the full revenue objective. The proof appears in Section D.1 of the supplementary appendix.

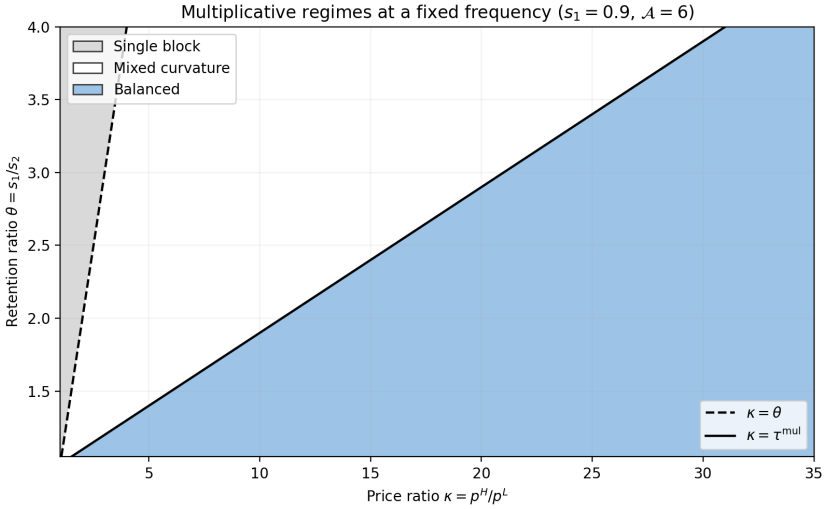


Figure 2: Sufficient arrangement regimes from Proposition 5.2 for prescribed $(\mathscr{T}, m)$, with $s_1 = 0.9$ and $\mathscr{A} = 6$. The dashed line is $\kappa = \theta$, and the solid line is $\kappa = \tau^{\mathrm{mul}}$. In the unshaded region, the interior contributions are convex and the terminal contribution is concave, so neither uniform shape condition applies.

For frequency optimization, the conclusions of Propositions 3.8 and 3.7 apply directly. Figure 2 illustrates these regimes for $s_1 = 0.9$ and $\mathscr{A} = 6$. In the region with mixed second-difference signs, the uniform concavity or convexity conditions do not apply, and the full revenue objective determines the optimal cycle. Section 6 bounds the losses of simple policies in this region, and Section 6.2 gives general and multiplicative memory truncation guarantees.

6 Approximation under Mixed Curvature

When the window contributions have mixed curvature, the structural results in Section 3 cannot identify an optimal policy from their curvature conditions. Exact optimization remains possible on the finite state

graph, but its size grows exponentially with memory. This motivates two computational approximations. The first builds on the preceding structural analysis. Balanced and single block policies can be optimized over finite candidate sets, and we bound their losses when the corresponding curvature conditions fail. The second truncates memory to reduce the state space, with a controlled revenue loss when distant prices have little effect on demand.

6.1 Balanced and single block policies

Recall the candidate frequency set from Proposition 3.7,

$$\mathscr{B} = \{0, 1\} \cup \left\{ \frac{j}{b+1} \,:\, b = 1, \dots, \mathscr{A},\ j = 1, \dots, b \right\}.$$

Let C_B be a balanced cycle at a frequency maximizing g over $\mathscr{B}$. Even without concavity, g is affine between consecutive points of $\mathscr{B}$, so C_B is optimal among balanced cycles. For the single block comparison, let H, L denote the regular and promotional prices, and let C_S maximize revenue over

$$\mathscr{S} = \{H, L\} \cup \{H^h L^m \,:\, 1 \le h, m \le \mathscr{A}\}.$$

Here $H^h L^m$ consists of h regular price periods followed by m promotional periods. Constant prices are included as boundary cases. The proposition shows that this finite search finds the best single block policy over all block lengths.

Recall that C_B is globally optimal when every $\psi_{b,k} \le 0$, while a constant price, and hence C_S, is globally optimal when every $\psi_{b,k} \ge 0$. These conclusions follow from Propositions 3.7 and 3.8. The following bounds measure the loss from violations of the respective sign conditions.

Proposition 6.1 (Revenue guarantees for simple policies)**.** *Under Assumption 2.1, the following guarantees hold.*

(i) *The best balanced policy C_B can be chosen with period at most $\mathscr{A} + 1$, and*

$$0 \le \bar{R}^* - \bar{R}(C_B) \le \sum_{b=1}^{\mathscr{A}} \sum_{k=1}^{b} \max\{\psi_{b,k}, 0\} \frac{k(b+1-k)}{b+1} =: E_B. \tag{6}$$

(ii) *The best single block policy C_S is found by evaluating at most $\mathscr{A}^2 + 2$ cycles. For $\mathscr{A} \ge 1$, its period is at most $2\mathscr{A}$, and*

$$0 \le \bar{R}^* - \bar{R}(C_S) \le \sum_{b=1}^{\mathscr{A}} \sum_{k=1}^{b} \max\{-\psi_{b,k}, 0\} \frac{k(b+1-k)}{b+1} =: E_S.$$

For $\mathscr{A} = 0$, both candidates reduce to optimal constant prices. Choosing C_{app} as the better of C_B and C_S gives

$$0 \le \bar{R}^* - \bar{R}(C_{\text{app}}) \le \min\{E_B, E_S\}.$$

Only coefficients with the opposite sign to the corresponding optimality condition enter each bound. Positive second differences enter E_B, while the magnitudes of negative second differences enter E_S. Each bound vanishes when its sign condition holds. The balanced search takes $O(\mathscr{A}^3)$ arithmetic operations. The additional single block search can improve the attained revenue, as the first example below illustrates. Both searches include constant prices. The proof appears in Section E.1 of the supplementary appendix.

The following two examples illustrate how these approximations can provide small revenue loss guarantees despite mixed curvature.

Weak reinforcement. For $\mathscr{A} = 3$, $p^L = 1$, and $p^H = 2$, consider

$$M_a(p, y) = 0.10 + 0.05 \cdot \mathbf{1}\{p = p^L\} + 0.04y + \epsilon\binom{y}{2}, \qquad 0 < \epsilon \le 0.01.$$

The quadratic term represents reinforcement among remembered promotions, and all market shares lie in $[0, 1]$. Interior second differences equal $2\epsilon > 0$, while terminal ones are negative, so the concavity condition fails for every $\epsilon > 0$. Nevertheless, $E_B = 11\epsilon/3$ tends to zero with reinforcement strength. At $\epsilon = 0.001$, $C_B = HL$ earns 0.901, while $C_S = HHLL$ earns the unrestricted optimum 0.9015. The balanced policy loses 0.0555%, with a guaranteed loss below 0.407%. Evaluating both candidates recovers the optimum.

Weak promotion fatigue. For $\mathscr{A} = 2$ and the same prices, let

$$M_a(p, y) = 0.20 + (0.20 + \epsilon/3)\mathbf{1}\{p = p^L\} - \epsilon\binom{y}{2}, \qquad 0 < \epsilon \le 0.01.$$

Here demand falls after repeated past promotions, and all shares remain in $[0, 1]$. The window second differences are -2ϵ, 0, and ϵ, giving mixed curvature and $E_S = \epsilon$. At $\epsilon = 0.001$, both constant prices earn 1.2, while $C_B = C_S = HL$ earns the optimal revenue 1.2005. The single block guarantee bounds its loss below 0.084%, while the actual loss is zero. Exact optimality certificates for both examples appear in Section E.1 of the supplementary appendix.

6.2 Memory truncation for general market share

Choose $0 \le J < \mathscr{A}$, retain the current price and the last $\min\{a, J\}$ past prices at age a, and replace demand by

$$M_{a,J}\left(p_t, \sum_{j=1}^{\min\{a,J\}} d_{t-j}\right), \qquad M_{a,J}(p, y) \in [0, 1].$$

All age cohorts remain in the objective, and $M_{a,J} = M_a$ for $a \le J$. For each affected age, define the largest error over compatible omitted histories by

$$\Delta^{\mathrm{tr}}_{a,J} := \max_{\substack{p \in \{p^L, p^H\} \\ 0 \le y \le J,\ 0 \le z \le a-J}} |M_a(p, y + z) - M_{a,J}(p, y)|, \qquad \varepsilon^{\mathrm{agg}}_J := p^H \sum_{a=J+1}^{\mathscr{A}} \Delta^{\mathrm{tr}}_{a,J},$$

where y, z are integers. Write $\bar{R}_J$ for average revenue under the truncated model, let $C^\star_J$ maximize it, and set $\bar{R}^\star_J = \bar{R}_J(C^\star_J)$.

Proposition 6.2 (General memory truncation guarantee)**.** *For every cycle C,*

$$|\bar{R}(C) - \bar{R}_J(C)| \le \varepsilon^{\mathrm{agg}}_J, \qquad \bar{R}(C^\star_J) \ge \bar{R}^* - 2\varepsilon^{\mathrm{agg}}_J.$$

If all market share errors have the same sign across ages, prices, and histories, the factor two in the policy bound can be replaced by one. These comparisons also hold over all feasible infinite price paths.

The truncated problem has only 2^J past-price histories and can be optimized by the maximum mean cycle method. Its policy is then evaluated under the original demand model. A short memory is justified when the aggregate error above is small relative to policy revenue. Finite memory alone does not imply a small truncation error.

One explicit choice is the midpoint of the smallest and largest values of $M_a(p, y+z)$ over $0 \le z \le a-J$. It is a valid market share and minimizes the largest error for each (a, p, y). Its error is half that range. Alternatively, the maximum gives an upper approximation and permits the one-sided bound. Both choices require no parametric demand form.

For example, suppose $M_a(p, y) = q^a m_a(p, y)$, where $0 < q < 1$ and $0 \le m_a \le 1$. Choosing $M_{a,J}(p, y) = q^a m_a(p, y)$ on the retained domain gives $\Delta^{\mathrm{tr}}_{a,J} \le q^a$, and hence

$$2\varepsilon_J^{\mathrm{agg}} \le \frac{2p^H q^{J+1}}{1-q}.$$

Thus fading activity with age can justify short memory even with nonlinear market share. Dividing revenues and bounds by $\mathscr{A}+1$ gives the corresponding normalized guarantees. The proof and a posteriori policy certificates appear in Section E.2 of the supplementary appendix.

Multiplicative demand. For the multiplicative market share in Section 5, truncation has a one-sided bound. For $0 \le J < \mathscr{A}$, retaining only the last $\min\{a, J\}$ past prices gives

$$M^{\mathrm{mul}}_{a,J}(p, y) = s(p) s_1^{y} s_2^{\min\{a,J\}-y}.$$

Omitting factors below one overstates survival, with age-specific error $\Delta^{\mathrm{tr}}_{a,J} = s_1^{J+1}(1 - s_2^{a-J})$ for $a > J$.

Proposition 6.3 (Multiplicative memory truncation)**.** *Define $\widehat{R} = \bar{R}/(\mathscr{A}+1)$ and $\widehat{R}_J = \bar{R}_J/(\mathscr{A}+1)$. If $C_J^\star$ optimizes the truncated model, then every cycle satisfies*

$$0 \le \widehat{R}_J(C) - \widehat{R}(C) \le \hat{\varepsilon}_J, \qquad \widehat{R}(C_J^\star) \ge \widehat{R}^\star - \hat{\varepsilon}_J,$$

where

$$\hat{\varepsilon}_J = \frac{p^H s_1^{J+1}}{\mathscr{A}+1} \sum_{a=J+1}^{\mathscr{A}} (1 - s_2^{a-J}) \le p^H s_1^{J+1}.$$

Thus J can be chosen to meet a revenue tolerance using the geometric bound. The factor one follows from Proposition 6.2. The age-specific error calculation is given in Section E.2 of the supplementary appendix.

7 Numerical Study

We examine how to implement the promotion planning procedure, how discarding price order affects policy revenue, and how the two computational approximations perform relative to their revenue loss bounds. The experiments use specified demand functions rather than an empirical calibration.

7.1 From estimated demand to a pricing cycle

Consider a monthly access service with regular price $p^H = 9.99$, promotional price $p^L = 5.99$, and five active ages $a = 0, \dots, 4$. The baseline market share is

$$M_a(p_t, y_{a,t}) = \alpha_a + \beta_a(d_t + y_{a,t}), \qquad \begin{aligned} \alpha &= (0.08, 0.15, 0.20, 0.23, 0.25), \\ \beta &= (0, 0.0675, 0.054, 0.042, 0.030). \end{aligned}$$

We simulate 400 independent groups over 24 measured months, with 60 users at each active age per group and month. Each group has a promotion probability drawn uniformly from $[0.2, 0.8]$. Prices are assigned

independently across months conditional on that probability, with four initial prices supplying the first measured histories. Conditional purchases are binomial with 60 trials and the specified market share. Using 300 groups, we estimate the age-specific intercepts and slopes by weighted least squares on observed cohort purchase proportions, with weights equal to cohort size and full-history memory fixed at four.

All estimated window contributions are concave. Theorem 3.5 and Proposition 3.7 therefore reduce optimization to balanced cycles at the candidate frequencies and select $\rho = 1/2$, implemented by HL. Under the true model, this cycle earns monthly revenue 9.6438, compared with 9.0909 under constant regular pricing and 9.1348 under constant promotional pricing. These values use unit mass arrivals and aggregate demand across five active ages. The improvement over constant regular pricing is 6.08%. Resampling the 300 estimation groups in 500 bootstrap replications selects HL in 99.2% of replications.

Figure 3 compares implementable price calendars and their revenue gains. Alternation and a single promotion block both use five promotions in ten periods, but alternation earns 2.32% more. The two constant price calendars show the revenue consequences of removing price variation altogether.



Figure 3: Ten-month price calendars and revenue gains under the true baseline demand. H and L denote regular and promotional prices. The balanced and single block calendars have the same promotion frequency, isolating the value of arrangement. The other calendars use constant prices.

7.2 Policy robustness under reference-price demand

To test sensitivity to price order, retain the baseline prices, ages, and coefficients, and let $e \in \{0, 1\}$ indicate historical engagement. Define

$$q^e_{a,t} = \alpha_a + \beta_a(d_t + e y_{a,t}), \qquad r_{a,t}(\lambda) = \frac{\sum_{j=1}^{a} \lambda^{j-1} p_{t-j}}{\sum_{j=1}^{a} \lambda^{j-1}}, \qquad z_{a,t} = \frac{r_{a,t}(\lambda) - p_t}{p^H - p^L}.$$

Following the gain and loss mechanisms of Kopalle et al. [18] and Popescu and Wu [26], set

$$M^{\text{true}}_{a,t} = \sigma\big(\text{logit}(q^e_{a,t}) + \gamma\{(z_{a,t})_+ - \ell(-z_{a,t})_+\}\big), \qquad \sigma(v) = \frac{1}{1 + \exp(-v)}, \tag{7}$$

with no reference term at age zero. For each e, use one $\gamma = 0$ control and all combinations of $\lambda \in \{0.2, 0.6, 1\}$, $\gamma \in \{0.5, 1.5\}$, and $\ell \in \{0.5, 2\}$. These 26 cases comprise ten count-based controls and 16 order-dependent cases. Smaller λ gives greater weight to recent prices.

The population count model averages true probabilities within each age, current-price, and past-count cell under the baseline randomized histories. For each case, we also generate 100 data sets using the baseline

sampling design and estimate the 30 cells from 300 groups by $\widehat{M}_a(p,y) = (Y_{a,p,y} + 1/2)/(N_{a,p,y} + 1)$, where Y and N are purchases and trials. No shape restriction is imposed. We optimize each count model and evaluate its policy under true demand, reporting

$$\mathrm{Regret}(\widehat{C}) = 100\frac{\bar{R}^*_{\mathrm{true}} - \bar{R}_{\mathrm{true}}(\widehat{C})}{\bar{R}^*_{\mathrm{true}}}.$$

Independent calculations on the equivalent graph with 16 past-price states and edge revenues verify all optima within 10^{-10} revenue units.

Mean estimated-policy regret across the order-dependent cases is 0.108% (Table A.2 in the supplementary appendix). The population count model selects a true optimum in 15 of 16 cases. The exception has $e = 0$, $\lambda = 0.2$, $\gamma = 1.5$, and $\ell = 0.5$. It selects HHL instead of HL, losing 1.045%. With estimation, mean regret is 1.195% (Monte Carlo standard error 0.043 percentage points, P95 2.413%). Thus strong gain sensitivity and emphasis on the latest price can make order consequential, despite good performance in most tested cases.

7.3 Computational approximation guarantees

We test both approximations in Section 6 on 1,296 specified demand functions, including 1,044 with mixed curvature, without estimation noise. Set $\mathscr{A} \in \{4, 6, 8\}$ and scale revenues and bounds by $1/(\mathscr{A} + 1)$. Losses and bounds are percentages of the unrestricted full-model optimum, used only as a numerical benchmark. A loss of 1% is the reference tolerance.

For general demand, take $p^L = 1$, $p^H = 2$, and

$$M^d_a(y) = q^a\{\alpha + \eta d + f^d_a(y)\}, \qquad \eta = \alpha - \frac{\sum_{a=0}^{\mathscr{A}} q^a f^1_a(a)}{\sum_{a=0}^{\mathscr{A}} q^a} + u, \qquad d \in \{0, 1\}.$$

The constant prices tie at $u = 0$. Cross $q \in \{0.4, 0.7, 0.95, 1\}$, controlling fading with age, and $u \in \{-0.04, 0, 0.04\}$, changing constant-price profitability, with two response grids. The first uses $\alpha = 0.30$ and

$$f^0_a(y) = f^1_a(y) = \beta y + \epsilon\binom{y}{2}, \qquad \beta = \frac{\pm 0.08}{\mathscr{A}}, \quad \epsilon = \frac{\pm\nu}{\binom{\mathscr{A}}{2}}, \quad \nu \in \{0.002, 0.02, 0.10, 0.18\},$$

with independent signs, giving 576 cases with varying response and curvature. The second uses $\alpha = 0.25$, $f^0_a = 0$, and

$$f^1_a(y) = \frac{\nu}{\mathscr{A}}\cos\left(\frac{2\pi(a+\phi)}{P}\right)y, \qquad \nu \in \{0.04, 0.12, 0.30\}, \quad P \in \{2, 3, 4\}, \quad \phi \in \{0, 1\},$$

giving 648 cases with effects changing sign across ages. Both general grids have shares in $[0, 0.824]$. The remaining 72 cases use multiplicative demand with $s_1 \in \{0.4, 0.7, 0.9, 0.98\}$, $s_2/s_1 \in \{0.5, 0.8\}$, $p^L = 1$, and $\kappa = \theta + t(\tau^{\mathrm{mul}} - \theta)$, $t \in \{0.25, 0.5, 0.75\}$. These lie in the mixed regime, with price ratios from approximately 1.292 to 38.75, including stress cases with large price gaps.

Balanced and single block policies. Table 2 compares the finite searches in Proposition 6.1, grouping general cases by $\omega = E_B/(E_B + E_S)$. The better policy loses at most 0.124%, versus 2.395% for balanced policies alone. Its curvature bound is below 1% in 47.0% of cases, but reaches 32.222% when the two sign budgets are comparable. Thus a large bound can coexist with high policy revenue.

Table 2: Maximum loss and curvature bound (%). The first three rows partition the general cases. Better selects between balanced and single block policies. Bound is $100 \min\{E_B, E_S\}/\bar{R}^*$. Column maxima may occur in different cases.

Setting	Cases	Balanced	Single block	Better	Bound
$E_B/(E_B+E_S) \le 0.2$	201	0.071	0.048	0.000	1.317
$0.2 < E_B/(E_B+E_S) < 0.8$	822	0.715	0.214	0.124	32.222
$E_B/(E_B+E_S) \ge 0.8$	201	0.000	0.000	0.000	1.333
Multiplicative, mixed	72	2.395	0.000	0.000	13.353

Memory truncation. We optimize every $J = 0, \dots, \mathscr{A}-1$, giving 7,776 truncated problems, and evaluate each selected cycle under full demand. General demand uses the compatible-range midpoint and bound $2\varepsilon_J^{\text{agg}}$. Multiplicative demand omits survival factors and uses the one-sided bound in Proposition 6.3. All 9,072 full and truncated optima have cycle and potential certificates within 10^{-9}. Independent linear programs also verify the full-model values.

At $J = 2$, four past-price states replace 16, 64, or 256. Table A.3 in the supplementary appendix gives the detailed comparison. General-demand loss reaches only 0.133% at $q = 0.4$, but 6.667% at $q \ge 0.95$. At $J = 3$, the bound is below 0.702% in every $q = 0.4$ case. Persistent memory can require the full history to meet the 1% bound benchmark. Multiplicative truncation loses up to 19.184% at $J = 2$, and none of its bounds is below 1%, even though all $s_1 = 0.4$ cases have zero loss to numerical precision. Factor omission inflates older-cohort demand, making the uniform bound conservative.

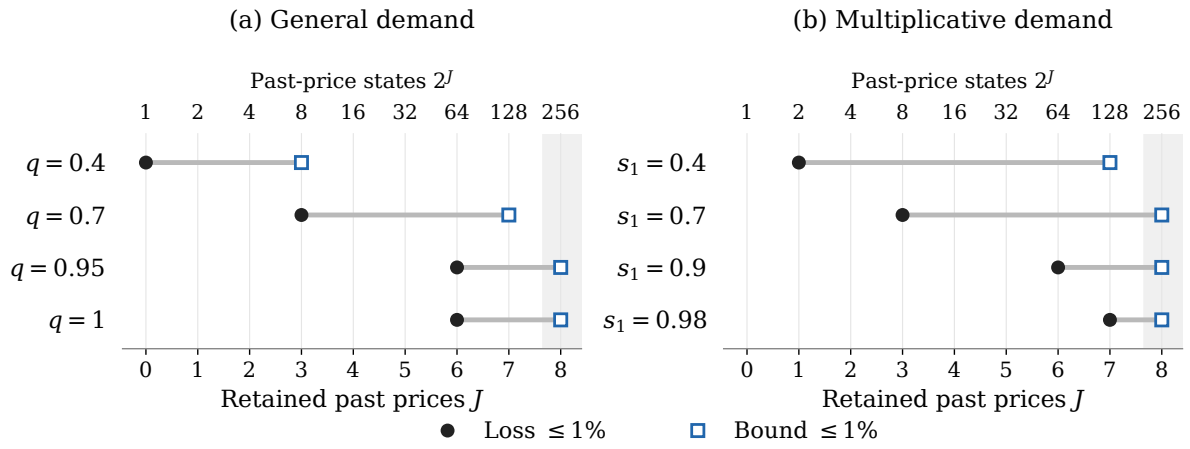


Figure 4: Minimum retained memory at $\mathscr{A} = 8$ needed to keep loss or its bound at or below 1%. Filled circles and open squares show the respective cutoffs, each computed using the maximum over the remaining parameters. Panel (a) pools the general response grids, and panel (b) uses multiplicative factor omission. Upper axes report the corresponding 2^J past-price states. At $J = 8$, full memory gives zero loss and bound.

These grids support simple-policy searches even when their curvature certificate is loose. Figure 4 shows that memory truncation benefits from fading exposure, while persistent memory requires larger cutoffs.

8 Conclusion

We study the joint choice of promotion frequency and arrangement when demand depends on current and remembered prices. A telescoping representation of marginal promotion effects yields a revenue decomposition into moving window contributions, whose curvature determines the preferred arrangement. Concave contributions favor balanced cycles, and frequency optimization reduces to a finite search with a tight bound on the optimal period. Convex contributions favor a single promotion block at fixed frequency and a constant price when frequency is unrestricted. The additively separable and multiplicative cases give economic interpretations of these conditions through engagement, reference prices, and retention. An affine extension allows historical effects to vary across lags.

Beyond the uniform curvature regimes, finite searches over balanced and single block policies and memory truncation provide revenue loss bounds. The simple-policy bounds measure violations of the corresponding curvature conditions, while the truncation bounds depend on demand errors caused by omitted histories. The numerical experiments illustrate implementation, examine sensitivity to price order, and distinguish policy performance from bound tightness. On the tested grids, the better balanced or single block policy loses at most 0.124%, although its bound can be much looser. Truncation performs well with fading demand, while persistent histories and survival-factor omission can produce substantial losses and conservative bounds. These findings support choosing approximations using their error budgets and evaluating candidate policies under full demand. Future work can sharpen these bounds, assess the methods using empirical demand, and extend the analysis to continuously chosen prices and multiple products.

AI Disclaimer. Use of Generative Artificial Intelligence. The authors used Codex GPT-6.1 Sol Medium to assist with language editing, notation checks, proof organization, exploration of mathematical arguments, and computational verification. The system was instructed to work within the stated model assumptions, preserve notation consistency, explain derivations, and check proposed results against computational examples. The authors reviewed, verified, and edited the resulting material and take responsibility for the manuscript's mathematical claims, computations, and presentation.

References

[1] Agrawal S, Tang W (2024) Dynamic pricing and learning with long-term reference effects. Working paper, arXiv:2402.12562.

[2] Altman E, Gaujal B, Hordijk A (2000a) Balanced sequences and optimal routing. *J. ACM* 47(4):752–775.

[3] Altman E, Gaujal B, Hordijk A (2000b) Multimodularity, convexity, and optimization properties. *Math. Oper. Res.* 25(2):324–347.

[4] Amazon (2025) Double the days, millions of deals. Prime Day event announcement, press release (June 16). Accessed October 9, 2026, Amazon Press Center.

[5] Anderson ET, Simester DI (2004) Long-run effects of promotion depth on new versus established customers: Three field studies. *Marketing Sci.* 23(1):4–20.

[6] Baardman L, Cohen MC, Panchamgam K, Perakis G, Segev D (2019) Scheduling promotion vehicles to boost profits. *Management Sci.* 65(1):50–70. doi:10.1287/mnsc.2017.2926.

[7] Baek J, Ma W, Mitrofanov D (2025) Personalized promotions in practice: Dynamic allocation and reference effects. Working paper, arXiv:2512.23781.

[8] Cinemark (n.d.) Discount Tuesdays movie deals. Accessed October 9, 2026, https://www.cinemark.com/discounts/discount-tuesdays.

[9] Cohen MC, Leung NHZ, Panchamgam K, Perakis G, Smith A (2017) The impact of linear optimization on promotion planning. *Oper. Res.* 65(2):446–468.

[10] Cohen MC, Kalas JJ, Perakis G (2021) Promotion optimization for multiple items in supermarkets.

Management Sci. 67(4):2340–2364. doi:10.1287/mnsc.2020.3641.

[11] Cohen-Hillel T, Panchamgam K, Perakis G (2023) High-low promotion policies for peak-end demand models. *Management Sci.* 69(4):2016–2050.

[12] Fader PS, Hardie BGS, Lee KL (2005) "Counting your customers" the easy way: An alternative to the Pareto/NBD model. *Marketing Sci.* 24(2):275–284. doi:10.1287/mksc.1040.0098.

[13] Fibich G, Gavious A, Lowengart O (2003) Explicit solutions of optimization models and differential games with nonsmooth (asymmetric) reference-price effects. *Oper. Res.* 51(5):721–734. doi:10.1287/opre.51.5.721.16758.

[14] Greenleaf EA (1995) The impact of reference price effects on the profitability of price promotions. *Marketing Sci.* 14(1):82–104. doi:10.1287/mksc.14.1.82.

[15] Hajek B (1985) Extremal splittings of point processes. *Math. Oper. Res.* 10(4):543–556.

[16] Hardie BGS, Johnson EJ, Fader PS (1993) Modeling loss aversion and reference dependence effects on brand choice. *Marketing Sci.* 12(4):378–394. doi:10.1287/mksc.12.4.378.

[17] Karp RM (1978) A characterization of the minimum cycle mean in a digraph. *Discrete Math.* 23(3):309–311.

[18] Kopalle PK, Rao AG, Assunção JL (1996) Asymmetric reference price effects and dynamic pricing policies. *Marketing Sci.* 15(1):60–85. doi:10.1287/mksc.15.1.60.

[19] Krishna A (1991) Effect of dealing patterns on consumer perceptions of deal frequency and willingness to pay. *J. Marketing Res.* 28(4):441–451. doi:10.1177/002224379102800406.

[20] Krishna A (1994) The impact of dealing patterns on purchase behavior. *Marketing Sci.* 13(4):351–373. doi:10.1287/mksc.13.4.351.

[21] Lal R, Rao R (1997) Supermarket competition: The case of every day low pricing. *Marketing Sci.* 16(1):60–80. doi:10.1287/mksc.16.1.60.

[22] Mazumdar T, Raj SP, Sinha I (2005) Reference price research: Review and propositions. *J. Marketing* 69(4):84–102. doi:10.1509/jmkg.2005.69.4.84.

[23] Mela CF, Gupta S, Lehmann DR (1997) The long-term impact of promotion and advertising on consumer brand choice. *J. Marketing Res.* 34(2):248–261.

[24] Nasiry J, Popescu I (2011) Dynamic pricing with loss-averse consumers and peak-end anchoring. *Oper. Res.* 59(6):1361–1368.

[25] Neslin SA, van Heerde HJ (2009) Promotion dynamics. *Found. Trends Marketing* 3(4):177–268. doi:10.1561/1700000010.

[26] Popescu I, Wu Y (2007) Dynamic pricing strategies with reference effects. *Oper. Res.* 55(3):413–429.

[27] Raju JS, Srinivasan V, Lal R (1990) The effects of brand loyalty on competitive price promotional strategies. *Management Sci.* 36(3):276–304. doi:10.1287/mnsc.36.3.276.

[28] Reinartz WJ, Kumar V (2000) On the profitability of long-life customers in a noncontractual setting: An empirical investigation and implications for marketing. *J. Marketing* 64(4):17–35. doi:10.1509/jmkg.64.4.17.18077.

[29] Schmittlein DC, Morrison DG, Colombo R (1987) Counting your customers: Who are they and what will they do next? *Management Sci.* 33(1):1–24. doi:10.1287/mnsc.33.1.1.

[30] Silva-Risso JM, Bucklin RE, Morrison DG (1999) A decision support system for planning manufacturers' sales promotion calendars. *Marketing Sci.* 18(3):274–300.

[31] Singer G, Holston R (2017) Driving smarter decisions on pricing, promotion. *Analytics Magazine* (January 2). Accessed October 10, 2026, INFORMS Analytics Magazine.

[32] Tsujikawa R, Nakano R, Hoshino T (2025) Coupon redemption, churn, and customer lifetime value for a newly-launched noncontractual product: Evidence from a large-scale field experiment in supermarkets. Keio-IES Discussion Paper DP2025-025, Keio University, working paper.

[33] van Heerde HJ, Leeflang PSH, Wittink DR (2000) The estimation of pre- and postpromotion dips with store-level scanner data. *J. Marketing Res.* 37(3):383–395. doi:10.1509/jmkr.37.3.383.18782.

[34] Winer RS (1986) A reference price model of brand choice for frequently purchased products. *J. Consumer Res.* 13(2):250–256. doi:10.1086/209064.

[35] Zhang DJ, Dai H, Dong L, Qi F, Zhang N, Liu X, Liu Z, Yang J (2020) The long-term and spillover effects of price promotions on retailing platforms: Evidence from a large randomized experiment on Alibaba. *Management Sci.* 66(6):2589–2609.

Supplementary Appendix

A Proofs and supplementary results for Section 2

Proof of Proposition 2.1. Fix any infinite price path. At time t, the payoff relevant history is the vector

$$s_t = (p_{t-\mathscr{A}}, \dots, p_t) \in \{p^L, p^H\}^{\mathscr{A}+1}.$$

Choosing the next price $p_{t+1} \in \{p^L, p^H\}$ shifts this vector from

$$(p_{t-\mathscr{A}}, \dots, p_t) \quad \text{to} \quad (p_{t-\mathscr{A}+1}, \dots, p_t, p_{t+1}),$$

which is exactly an edge in $\mathscr{G}_{\mathscr{A}}$. Hence every feasible infinite price path induces an infinite walk in $\mathscr{G}_{\mathscr{A}}$. Conversely, every infinite walk in $\mathscr{G}_{\mathscr{A}}$ induces a feasible price path by reading the last coordinate of each visited state.

The node weight $w(s)$ was defined so that $w(s_t)$ is exactly the one period revenue at time t. Thus the seller's long run average revenue is the long run average node weight of the corresponding walk. Let $\lambda^\star$ be the maximum mean weight of a directed cycle in $\mathscr{G}_{\mathscr{A}}$. Any finite prefix of a walk in a finite graph can be decomposed into a residual simple path plus directed cycles. The residual path has length at most the number of nodes, while each extracted cycle has mean weight at most $\lambda^\star$. Therefore the average weight of a prefix of length N is at most $\lambda^\star + O(1/N)$. Taking $N \to \infty$ shows that no infinite price path can beat $\lambda^\star$.

For the reverse inequality, take a directed cycle attaining $\lambda^\star$ and post its periodic price sequence from time zero. After at most $\mathscr{A}$ periods, the fixed initial history has left every memory window and the induced states follow that cycle. The initial discrepancy affects only finitely many bounded rewards, so its contribution to average revenue vanishes. The posted prices are periodic from time zero and their long run average revenue is $\lambda^\star$. Hence the optimal long run value equals the maximum mean cycle value, and an optimal policy can be chosen periodic. □

A.1 Revenue reversal symmetry

Reversal preserves revenue because it permutes the promotion counts in windows of each length.

Proposition A.1 (Revenue reversal symmetry)**.** *Under Assumption 2.1, any price cycle $C = (p_1, \dots, p_{\mathscr{T}})$ and its reversal $C_{\text{rev}} = (p_{\mathscr{T}}, \dots, p_1)$ have the same average revenue.*

Proof. Write $\tilde{d}_t = d_{\mathscr{T}+1-t}$ for the reversed promotion indicators. For each depth b,

$$\widetilde{N}_{b,t} = \sum_{i=0}^{b} d_{\mathscr{T}+1-t+i} = N_{b,\mathscr{T}+1-t+b}.$$

Cyclic summation therefore gives $\sum_t F_b(\widetilde{N}_{b,t}) = \sum_t F_b(N_{b,t})$, including when windows exceed the cycle length. Reversal also preserves $\mathscr{T}$ and m. Proposition 3.2, whose proof does not use reversal symmetry, now gives equal cycle revenues. Dividing by $\mathscr{T}$ proves the claim. □

Thus a search over cycles can retain one representative from each class related by rotation or reversal.

B Proofs and supplementary results for Section 3

B.1 Market share representation and revenue decomposition

Proof of Lemma 3.1. For $y = 0, \dots, a$, the sum of the first differences telescopes to give

$$\alpha_a^j + \sum_{k=1}^{y} \Delta_{a,k}^j = M_a^j(y).$$

Evaluating at zero and taking consecutive differences proves uniqueness, separately for each age and current price. □

Proof of Proposition 3.2. For $a \geq 1$, write the historical response at the regular price as

$$M_a^H(y) - \alpha_a^H = \Delta_{a,1}^H y + \sum_{k=1}^{a-1} (\Delta_{a,k+1}^H - \Delta_{a,k}^H)(y-k)_+.$$

Cyclic summation gives $\sum_t y_{a,t} = am$ and $\sum_t (y_{a,t} - k)_+ = V_{a-1,k}$. Since d_t is binary,

$$d_t \mathbf{1}\{y_{a,t} \geq k\} = (y_{a,t} + d_t - k)_+ - (y_{a,t} - k)_+,$$

whose cycle sum is $V_{a,k} - V_{a-1,k}$. Expanding revenue with these identities yields

$$\begin{aligned} R(C) = \mathscr{T} g_0 + m g_1 + p^H \sum_{a=1}^{\mathscr{A}} \sum_{k=1}^{a-1} (\Delta_{a,k+1}^H - \Delta_{a,k}^H) V_{a-1,k} \\ - \sum_{a=1}^{\mathscr{A}} \sum_{k=1}^{a} r_{a,k} (V_{a,k} - V_{a-1,k}). \end{aligned}$$

The constant and linear terms are precisely the stated g_0 and g_1, with the historical term omitted at age zero. Collecting the coefficient of $V_{b,k}$ gives $p^H(\Delta_{b+1,k+1}^H - \Delta_{b+1,k}^H) + r_{b+1,k} - r_{b,k}$ for $b < \mathscr{A}$. At $b = \mathscr{A}$, no older cohort contributes, leaving $-r_{\mathscr{A},k}$. Substituting $V_{b,k} = \sum_t (N_{b,t} - k)_+$ proves the decomposition. The second difference of $(n-k)_+$ at an integer ℓ equals one if $\ell = k$ and zero otherwise, so the second difference of F_b at k is $\psi_{b,k}$. If $\mathscr{A} = 0$, all window sums are empty and the same formula holds. □

B.2 Balanced cycles and arrangement comparisons

Proof of Lemma 3.3. First, d_t is binary. Since $0 \leq m/\mathscr{T} \leq 1$, the quantity $tm/\mathscr{T}$ increases by at most one when t increases by one, so

$$d_t = \left\lfloor \frac{tm}{\mathscr{T}} \right\rfloor - \left\lfloor \frac{(t-1)m}{\mathscr{T}} \right\rfloor \in \{0, 1\}.$$

Second, the sequence is $\mathscr{T}$ periodic because m is an integer,

$$d_{t+\mathscr{T}} = \left\lfloor \frac{tm}{\mathscr{T}} + m \right\rfloor - \left\lfloor \frac{(t-1)m}{\mathscr{T}} + m \right\rfloor = d_t.$$

Third, one period contains exactly m promotions, since the sum telescopes,

$$\sum_{t=1}^{\mathscr{T}} d_t = \left\lfloor \frac{\mathscr{T} m}{\mathscr{T}} \right\rfloor - \lfloor 0 \rfloor = m.$$

It remains to prove the balance property. For any starting time u and block length k,

$$\sum_{t=u+1}^{u+k} d_t = \left\lfloor \frac{(u+k)m}{\mathscr{T}} \right\rfloor - \left\lfloor \frac{um}{\mathscr{T}} \right\rfloor.$$

Let $x = um/\mathscr{T}$, $y = km/\mathscr{T}$, and $\{x\} := x - \lfloor x \rfloor$. Since $0 \le \{x\} < 1$,

$$\lfloor x + y \rfloor - \lfloor x \rfloor = \lfloor \{x\} + y \rfloor \in \{\lfloor y \rfloor, \lceil y \rceil\}.$$

When y is an integer, the difference equals y. Thus

$$\left\lfloor \frac{(u+k)m}{\mathscr{T}} \right\rfloor - \left\lfloor \frac{um}{\mathscr{T}} \right\rfloor \in \left\{ \left\lfloor \frac{km}{\mathscr{T}} \right\rfloor, \left\lceil \frac{km}{\mathscr{T}} \right\rceil \right\}.$$

This holds for every k, and therefore in particular for all payoff relevant block lengths $k = 1, \dots, \mathscr{A} + 1$. Hence the constructed cycle is balanced. □

Balanced cycle uniqueness when $\mathscr{T} \le \mathscr{A} + 1$. We prove uniqueness by induction on $\mathscr{T}$. The cases $m = 0$, $m = 1$, and $m = \mathscr{T}$ each have only one rotation class, so assume $1 < m < \mathscr{T}$. When $\mathscr{T} \le \mathscr{A} + 1$, balance applies to all block lengths through $\mathscr{T}$. Write the promotion positions as $u_1 < \cdots < u_m$, extend them by $u_{i+m} = u_i + \mathscr{T}$, and define the cyclic gaps $g_i = u_{i+1} - u_i$.

For $1 \le q < m$, put $D_{i,q} := u_{i+q} - u_i$. The interval from $u_i + 1$ to $u_{i+q} - 1$ has length $D_{i,q} - 1$ and contains $q - 1$ promotions. Including both endpoints gives length $D_{i,q} + 1$ and $q + 1$ promotions. Balance, with the empty interval interpreted as having count zero, gives

$$\left\lfloor \frac{(D_{i,q} - 1)m}{\mathscr{T}} \right\rfloor \le q - 1, \qquad q + 1 \le \left\lceil \frac{(D_{i,q} + 1)m}{\mathscr{T}} \right\rceil.$$

These inequalities imply $q\mathscr{T}/m - 1 < D_{i,q} < q\mathscr{T}/m + 1$. Since $D_{i,q}$ is an integer,

$$\sum_{j=0}^{q-1} g_{i+j} = D_{i,q} \in \left\{ \left\lfloor \frac{q\mathscr{T}}{m} \right\rfloor, \left\lceil \frac{q\mathscr{T}}{m} \right\rceil \right\}.$$

The same assertion holds for $q = m$, because the gaps sum to $\mathscr{T}$.

Write $\mathscr{T} = km + r$, where $k = \lfloor \mathscr{T}/m \rfloor$ and $0 \le r < m$. The case $q = 1$ shows that $g_i = k + e_i$, with $e_i \in \{0, 1\}$ and $\sum_i e_i = r$. Subtracting qk from the consecutive gap bounds gives

$$\sum_{j=0}^{q-1} e_{i+j} \in \left\{ \left\lfloor \frac{qr}{m} \right\rfloor, \left\lceil \frac{qr}{m} \right\rceil \right\}, \qquad q = 1, \dots, m.$$

Thus $(e_1, \dots, e_m)$ is a fully balanced circular binary word of length $m < \mathscr{T}$. By induction its rotation class is unique. The gap sequence is therefore unique up to a cyclic shift, and choosing its first promotion position changes the original cycle only by rotation. The construction in Lemma 3.3 supplies this unique rotation class.

For $\mathscr{T} > \mathscr{A} + 1$, balanced cycles need not be unique up to rotation. With $\mathscr{T} = 8$, $m = 2$, and $\mathscr{A} = 2$, both

$$(1, 0, 0, 1, 0, 0, 0, 0) \qquad \text{and} \qquad (1, 0, 0, 0, 1, 0, 0, 0)$$

are balanced. Every window of length at most three contains either zero or one promotion, as required by the balance condition. Their successive promotion gaps are $(3, 5)$ and $(4, 4)$, respectively, so the two cycles are not cyclic rotations of one another.

Proof of Lemma 3.4. Fix b and a discretely convex function ϕ on $\{0, \dots, b+1\}$. The total exposure is fixed.

$$\sum_{t=1}^{\mathscr{T}} N_{b,t} = (b+1)m.$$

Thus $\sum_t \phi(N_{b,t})$ is minimized when the window counts differ by at most one. If $x_i \geq x_j + 2$, replacing (x_i, x_j) by $(x_i - 1, x_j + 1)$ weakly decreases the sum. A balanced cycle realizes this equalized vector simultaneously at every depth. For a strictly convex function, the decrease is strict whenever the numbers of promotions in two windows differ by at least two. Applying the same comparison to the negative of a strictly concave function gives a strict increase in its total contribution.

For the strict comparison, suppose $\mathscr{T} \leq \mathscr{A} + 1$ and the cycle is nonbalanced. Since the average number of promotions is fixed at each window length, nonbalance implies that two windows have numbers of promotions differing by at least two. Removing complete cycles from both windows subtracts the same multiple of m and preserves this difference. The residual length cannot be zero, one, or $\mathscr{T} - 1$, because the numbers of promotions in windows of these lengths differ by at most one. Thus such a pair exists at some depth $1 \leq b \leq \mathscr{T} - 2$. At that depth, strict convexity or strict concavity makes the comparison with a balanced cycle strict.

For the maximum, identify positions with the vertices of the cyclic group $\mathbb{Z}_{\mathscr{T}}$. We first establish the required polarization inequality. Let σ be any reflection of the cycle and let H be one of the two closed semicircles bounded by its reflection axis. From each nonfixed orbit $\{i, \sigma(i)\}$, denote by i the representative in H. The polarization toward H of the promotion set $P = \{i : d_i = 1\}$ keeps the number of promotions in every reflected pair but places the promotion in H whenever exactly one member of that pair is promotional. Equivalently,

$$d_i^H - d_{\sigma(i)}^H = \left| d_i - d_{\sigma(i)} \right|, \qquad i \in H.$$

Write the window length as $b + 1 = q\mathscr{T} + s$, where $0 \leq s < \mathscr{T}$. Every window then contains qm promotions from its q complete turns around the cycle, plus the promotions in a residual cyclic interval I of length s. The complete turns are arrangement independent, so define the discrete convex function

$$\varphi(z) := \phi(qm + z).$$

For a residual interval I and its reflection $\sigma(I)$, let their numbers of promotions before polarization be (x, y) and afterward be (x', y'). For every nonfixed reflected pair, set

$$a_i := d_i - d_{\sigma(i)}, \qquad c_i := \mathbf{1}\{i \in I\} - \mathbf{1}\{\sigma(i) \in I\}.$$

All nonzero c_i's have the same sign. Indeed, rotate coordinates so that $\sigma(z) = -z \pmod{\mathscr{T}}$ and the interior of H is $(0, \mathscr{T}/2)$. Opposite signs would give $u, v \in H$ with $u, -v \in I$ and $-u, v \notin I$. These four vertices alternate between membership and nonmembership around the circle, which is impossible for a cyclic interval. Boundary vertices fixed by σ contribute equally to the two counts. Therefore

$$x - y = \sum_{i \in H} c_i a_i, \qquad x' - y' = \sum_{i \in H} c_i |a_i|.$$

Polarization also preserves $x + y$, because it preserves the number of promotions in each reflected pair. The common sign property now gives

$$|x' - y'| = \sum_{i \in H} |c_i|\,|a_i| \geq \left| \sum_{i \in H} c_i a_i \right| = |x - y|.$$

Thus (x', y') majorizes (x, y). Discrete convexity of φ implies

$$\varphi(x') + \varphi(y') \geq \varphi(x) + \varphi(y).$$

Pairing every residual interval with its reflection shows that one polarization weakly increases $\sum_t \phi(N_{b,t})$. The argument uses the same polarization for every depth and convex function, so all convex window sums increase simultaneously.

It remains to show that these polarizations lead to a single promotion block. Embed the vertices on a circle and choose a reference point c on the circle whose arc distances to the vertices are all distinct. For each reflection σ, let $H_{\sigma,c}$ contain the vertices i satisfying $\mathrm{dist}(i, c) < \mathrm{dist}(\sigma(i), c)$, together with any vertices fixed by σ. This is one of the two semicircles bounded by the reflection axis. Every nontrivial polarization toward $H_{\sigma,c}$ strictly decreases the potential

$$\Psi(P) := \sum_{i \in P} \mathrm{dist}(i, c),$$

because each moved promotion is replaced by its strictly closer reflected partner. Since there are finitely many promotion sets, repeated centered polarizations terminate. At termination, P must consist of the m vertices closest to c. Otherwise, some regular price vertex $u \notin P$ would be closer to c than a promotional vertex $v \in P$, and the reflection $\sigma(i) = u + v - i \pmod{}T$, which exchanges u and v, would produce another nontrivial centered polarization. The m closest vertices to c are consecutive. Hence the terminal cycle has a single promotion block, and because no polarization decreases any $\sum_t \phi(N_{b,t})$, a single block cycle maximizes every convex window sum. Applying the argument to the negative of a concave function gives the concave comparisons. □

Numerical illustration. The numerical illustration shows how the market share function changes the preferred arrangement at a common promotion frequency. Set $\mathscr{T} = 12, m = 4, \mathscr{A} = 5, p^L = 1$, and $p^H = 1.8$. We enumerate all $\binom{12}{4} = 495$ labeled cycles, each with frequency $1/3$, under five market share functions. Each specification is a special case of $M_a(p, y)$ expressed through

$$x = y + \mathbf{1}\{p = p^L\},$$

which counts remembered past promotions together with the current promotion. Table A.1 reports the comparison.

Table A.1: Arrangement effects at fixed $\mathscr{T} = 12, m = 4, \mathscr{A} = 5, p^L = 1$, and $p^H = 1.8$. The market share formulas use $x = y + \mathbf{1}\{p = p^L\}$. Arrangements are reported up to cyclic rotation as promotion indicator cycles $d = (d_1, \ldots, d_{12})$, where $d_t = 1$ denotes a promotion and $d_t = 0$ denotes the regular price. The revenue range is (best − worst)/best.

Market share function	Best cycle d	Worst cycle d	Range
$M_a(p, y) = 0.04 + 0.2(0.82)^a x$	(0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1)	(0, 0, 0, 0, 0, 0, 0, 0, 1, 1, 1, 1)	7.88%
$M_a(p, y) = 0.04 + 0.2(0.84)^a(1 - 0.15^x)$	(0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1)	(0, 0, 0, 0, 0, 0, 0, 0, 1, 1, 1, 1)	27.54%
$M_a(p, y) = 0.04 + 0.008(0.85)^a x^3$	(0, 0, 0, 0, 0, 0, 0, 0, 1, 1, 1, 1)	(0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1)	31.70%
$M_a(p, y) = 0.04 + 0.48(0.85)^a \mathbf{1}\{x \geq 2\}$	(0, 0, 0, 0, 1, 1, 0, 0, 0, 0, 1, 1)	(0, 0, 0, 1, 0, 0, 1, 0, 1, 0, 0, 1)	28.66%
$M_a(p, y) = 0.95^x 0.15^{a+1-x}$	(0, 0, 0, 0, 0, 0, 0, 0, 1, 1, 1, 1)	(0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1)	37.64%

For these parameter values, balanced arrangements attain the highest revenue under the affine and concave specifications, while single block arrangements do so under the cubic and multiplicative specifications. The threshold market share function favors two separated promotion blocks. These comparisons incorporate both the demand generated by each cycle and the transaction prices.

B.3 Promotion frequency for balanced arrangements

Proof of Proposition 3.6. Fix a balanced cycle C^{bal} with $\mathscr{T}$ periods, m promotions, and frequency $\rho = m/\mathscr{T}$. For each $b = 1, \dots, \mathscr{A}$, cyclic summation gives

$$\frac{1}{\mathscr{T}} \sum_{t=1}^{\mathscr{T}} N_{b,t}(C^{\mathrm{bal}}) = (b+1)\rho.$$

This identity also holds when the window is longer than the cycle, since each cyclic position is counted $b+1$ times, with multiplicity. Write $n_b = \lfloor (b+1)\rho \rfloor$ and $\xi_b = (b+1)\rho - n_b$. When $\xi_b > 0$, balance restricts the window counts to n_b and $n_b + 1$, and their mean implies that the respective proportions are $1 - \xi_b$ and ξ_b. Thus

$$\frac{1}{\mathscr{T}} \sum_{t=1}^{\mathscr{T}} F_b\big(N_{b,t}(C^{\mathrm{bal}})\big) = (1-\xi_b)F_b(n_b) + \xi_b F_b(n_b+1) = \widetilde{F}_b\big((b+1)\rho\big).$$

When $\xi_b = 0$, every window contains n_b promotions, so the same average equals $F_b(n_b) = \widetilde{F}_b((b+1)\rho)$. In particular, at $\rho = 1$ this value is $F_b(b+1)$, without evaluating F_b outside its domain. Dividing the revenue decomposition in Proposition 3.2 by $\mathscr{T}$ now gives

$$\bar{R}(C^{\mathrm{bal}}) = g_0 + g_1\rho + \sum_{b=1}^{\mathscr{A}} \widetilde{F}_b\big((b+1)\rho\big) = g(\rho).$$

No concavity assumption is needed for this representation. □

Proof of Proposition 3.7. For each b, discrete concavity means that the increments $F_b(k+1) - F_b(k)$, $k = 0, \dots, b$, are nonincreasing. These increments are the slopes of $\widetilde{F}_b$ on consecutive unit intervals. Hence $\widetilde{F}_b$ is continuous, piecewise linear, and concave. Composing it with $(b+1)\rho$, summing over b, and adding $g_0 + g_1\rho$ preserves these properties, so g is continuous, piecewise linear, and concave on $[0,1]$.

The slope of $\widetilde{F}_b((b+1)\rho)$ can change only at $\rho = j/(b+1)$, $j = 1, \dots, b$. Thus g is affine between consecutive distinct points of $\mathscr{B}$. It attains its maximum on the compact interval $[0,1]$, and an affine function on each such interval attains its maximum at an endpoint. Therefore some maximizer $\rho^\star$ belongs to $\mathscr{B}$.

Write $\rho^\star = m^\star/\mathscr{T}^\star$ in lowest terms, using $\mathscr{T}^\star = 1$ at either endpoint. If $\rho^\star = j/(b+1)$ is interior, its reduced denominator satisfies $\mathscr{T}^\star \le b+1 \le \mathscr{A}+1$. Lemma 3.3 constructs a balanced cycle $C^\star$ of that length and frequency. At the endpoints, use the corresponding constant price cycle of length one. Proposition 3.6 gives $\bar{R}(C^\star) = g(\rho^\star)$.

Finally, let C be any cycle with $\mathscr{T}$ periods and m promotions. By Lemma 3.3, there is a balanced cycle C^{bal} with the same $(\mathscr{T}, m)$. Theorem 3.5 and Proposition 3.6 imply

$$\bar{R}(C) \le \bar{R}(C^{\mathrm{bal}}) = g(m/\mathscr{T}) \le g(\rho^\star) = \bar{R}(C^\star).$$

Thus $C^\star$ maximizes average revenue over all cycles. Proposition 2.1 then establishes optimality over all feasible infinite price paths. When $\mathscr{A} = 0$, the window sum is empty, g is affine, and $\mathscr{B} = \{0,1\}$. The same argument selects a constant price cycle of length one. □

Sharpness of the period bound in Proposition 3.7. Fix $\mathscr{A} \geq 1$, set $L = \mathscr{A} + 1$, $p^L = 1$, $p^H = 2$, and let

$$M_a(p, y) = \alpha + \beta y, \qquad \beta = \frac{1}{4\mathscr{A}}, \qquad \alpha = \beta\left(\mathscr{A} - \frac{1}{2}\right).$$

These shares lie in $[\alpha, \alpha + \mathscr{A}\beta] \subset (0, 1)$. Proposition 3.2 gives $F_b = 0$ for $b < \mathscr{A}$, $F_{\mathscr{A}}(n) = -\beta\binom{n}{2}$, and $g_1 = \beta L/2$. On $j/L < \rho < (j+1)/L$, the slope of g is $\beta L(1/2 - j)$. It is positive for $j = 0$ and negative thereafter, so $1/L$ is the unique maximizing frequency. Every cycle is dominated by a balanced cycle at its frequency. Hence every optimal periodic policy has frequency $1/L$ and period divisible by L, while a balanced cycle of period L attains the optimum. □

Algorithm 1 Balanced frequency search under concave window contributions

```
Form ℬ = {0, 1} ∪ {j/(b + 1) : 1 ≤ b ≤ 𝒜, 1 ≤ j ≤ b}.
Set g* = −∞ and ρ* = 0.
for each ρ ∈ ℬ do
    Set v = g_0 + g_1 ρ.
    for b = 1, ..., 𝒜 do
        Set n = ⌊(b + 1)ρ⌋ and ξ = (b + 1)ρ − n.
        if ξ = 0 then
            Add F_b(n) to v.
        else
            Add (1 − ξ)F_b(n) + ξF_b(n + 1) to v.
        end if
    end for
    if v > g* then
        Set g* = v and ρ* = ρ.
    end if
end for
Write ρ* = m*/𝒯* in lowest terms, using 𝒯* = 1 at the endpoints.
Construct the corresponding balanced cycle by Lemma 3.3.
```

Given the window contribution tables, each frequency evaluation takes $O(\mathscr{A})$ arithmetic operations. There are $O(\mathscr{A}^2)$ candidates, giving $O(\mathscr{A}^3)$ operations and $O(\mathscr{A}^2)$ space for $\mathscr{A} \geq 1$. At $\mathscr{A} = 0$, compare the two constant prices. These are unit cost arithmetic bounds.

B.4 Promotion frequency under convex window contributions

Proof of Proposition 3.8. Discrete convexity bounds each window contribution by its endpoint chord, giving

$$F_b(n) \leq \left(1 - \frac{n}{b+1}\right) F_b(0) + \frac{n}{b+1} F_b(b+1), \qquad n = 0, \ldots, b+1.$$

Since $F_b(0) = 0$ and $\mathscr{T}^{-1} \sum_t N_{b,t} = (b+1)\rho$, Proposition 3.2 gives

$$\bar{R}(C) \leq g_0 + \rho \left[g_1 + \sum_{b=1}^{\mathscr{A}} F_b(b+1) \right] = g_0 + \rho D.$$

To identify the bracket, evaluate the revenue decomposition at the two constant cycles. Their revenues are g_0 and $g_0 + g_1 + \sum_b F_b(b+1)$. Their difference is D. The upper bound is maximized at $\rho = 0$ if $D \leq 0$, and at $\rho = 1$ if $D \geq 0$. The corresponding constant cycle attains the bound. Both are optimal when $D = 0$. Proposition 2.1 extends optimality to all feasible infinite price paths. □

B.5 Heterogeneous active lifetimes

Suppose type i has proportion $\lambda_i > 0$ in each arriving cohort, maximum active age $\mathscr{A}_i$, and count-based market share $M_{i,a}(p, y)$, with $\sum_{i=1}^{K} \lambda_i = 1$. All types face the same price path and remember prices since arrival. Write $\mathscr{A}_{\max} = \max_i \mathscr{A}_i$.

Proposition B.1 (Revenue decomposition with heterogeneous customers)**.** *Let $g_s^{(i)}$, $F_b^{(i)}$, and $\psi_{b,k}^{(i)}$ be the quantities in Proposition 3.2 for type i, extending the last two by zero for $b > \mathscr{A}_i$. Then*

$$R^{\text{het}}(C) = \mathscr{T}\, g_0^{\text{het}} + m g_1^{\text{het}} + \sum_{b=1}^{\mathscr{A}_{\max}} \sum_{t=1}^{\mathscr{T}} F_b^{\text{het}}(N_{b,t}(C)),$$

where

$$g_s^{\text{het}} = \sum_i \lambda_i g_s^{(i)}, \quad s = 0, 1, \qquad F_b^{\text{het}} = \sum_i \lambda_i F_b^{(i)}, \qquad \psi_{b,k}^{\text{het}} = \sum_i \lambda_i \psi_{b,k}^{(i)}.$$

The arrangement and frequency results apply with $\mathscr{A}_{\max}$ replacing $\mathscr{A}$, whenever the aggregate window contributions satisfy the corresponding concavity or convexity conditions. Individual types need not satisfy these conditions.

Proof. Aggregate revenue is $\sum_i \lambda_i R_i(C)$. Apply Proposition 3.2 separately to each type. All type-specific decompositions use the same window counts $N_{b,t}(C)$, so collecting the terms gives the stated weighted sums. Taking second differences gives the formula for $\psi_{b,k}^{\text{het}}$. The window comparisons and interpolation arguments then apply to the aggregate contributions. Aggregate rewards have finite history length $\mathscr{A}_{\max}$, so Proposition 2.1 also applies. □

C Proofs and supplementary results for Section 4

C.1 Engagement channel

C.1.1 Additive policy result

Proof of Corollary 4.1. The historical marginal effects are the same at both current prices, so $r_{a,k} = \Delta p\, \Delta_{a,k}$. Substitution into Proposition 3.2 yields the displayed $\psi_{b,k}$ in Section 4.1. Under Assumption 4.1, decreasing marginal effects in exposure and age make both interior terms nonpositive, while nonnegative marginal effects give the terminal sign. Therefore every F_b is concave. Theorem 3.5 gives balanced optimality. Proposition 3.7 gives the frequency and short period conclusions, including optimality over all feasible infinite price paths.

The additive frequency coefficients are

$$g_0 = p^H \sum_{a=0}^{\mathscr{A}} \alpha_a, \qquad g_1 = \sum_{a=0}^{\mathscr{A}} (p^L \eta_a - \Delta p\, \alpha_a) + p^H \sum_{a=1}^{\mathscr{A}} a \Delta_{a,1}.$$

The window contributions are $F_b(n) = \sum_{k=1}^{b} (n-k)_+ \psi_{b,k}$, and their linear interpolations give the frequency objective by Proposition 3.6. □

C.1.2 Affine specialization

For affine engagement, the balanced frequency objective simplifies to

$$g^{\text{aff}}(\rho) = g_0^{\text{aff}} + g_1^{\text{aff}} \rho - \Delta p \sum_{b=1}^{\mathscr{A}} (\beta_b^{\text{aff}} - \beta_{b+1}^{\text{aff}}) q_b^{\text{aff}}(\rho), \tag{C.1}$$

where $q_b^{\mathrm{aff}}(\rho) = q_{b,2}(\rho)$, $\beta_{\mathscr{A}+1}^{\mathrm{aff}} = 0$, and

$$g_0^{\mathrm{aff}} = p^H \sum_{a=0}^{\mathscr{A}} \alpha_a^{\mathrm{aff}},$$

$$g_1^{\mathrm{aff}} = -\Delta p \sum_{a=0}^{\mathscr{A}} \alpha_a^{\mathrm{aff}} + \sum_{a=0}^{\mathscr{A}} (p^L \eta_a^{\mathrm{aff}} + p^H a \beta_a^{\mathrm{aff}}).$$

Affine revenue representation. For affine engagement, $\Delta_{a,k} = \beta_a^{\mathrm{aff}}$. Proposition 3.2 gives

$$\psi_{b,k} = -\Delta p(\beta_b^{\mathrm{aff}} - \beta_{b+1}^{\mathrm{aff}}), \qquad F_b(n) = -\Delta p(\beta_b^{\mathrm{aff}} - \beta_{b+1}^{\mathrm{aff}}) \binom{n}{2},$$

with $\beta_{\mathscr{A}+1}^{\mathrm{aff}} = 0$. The frequency terms are

$$K_{\mathscr{T},m}^{\mathrm{aff}} = (\mathscr{T} p^H - \Delta p\, m) \sum_a \alpha_a^{\mathrm{aff}} + m \sum_a (p^L \eta_a^{\mathrm{aff}} + p^H a \beta_a^{\mathrm{aff}}) = \mathscr{T} g_0^{\mathrm{aff}} + m g_1^{\mathrm{aff}}.$$

Hence the cycle revenue is

$$R^{\mathrm{aff}}(C) = \mathscr{T} g_0^{\mathrm{aff}} + m g_1^{\mathrm{aff}} - \Delta p \sum_{b=1}^{\mathscr{A}} (\beta_b^{\mathrm{aff}} - \beta_{b+1}^{\mathrm{aff}}) Q_{b,2}(C). \tag{C.2}$$

For a balanced cycle, dividing by $\mathscr{T}$ and applying Proposition 3.6 gives (C.1). □

Summing $M_a^{\mathrm{aff}} = \alpha_a^{\mathrm{aff}} + \eta_a^{\mathrm{aff}} d_t + \beta_a^{\mathrm{aff}} y_{a,t}$ over the cycle and using $\sum_t y_{a,t} = am$ gives

$$\sum_{t=1}^{\mathscr{T}} \sum_{a=0}^{\mathscr{A}} M_a^{\mathrm{aff}}\big(p_t(C), y_{a,t}(C)\big) = \mathscr{T} \sum_{a=0}^{\mathscr{A}} \alpha_a^{\mathrm{aff}} + m \sum_{a=0}^{\mathscr{A}} \big(\eta_a^{\mathrm{aff}} + a\beta_a^{\mathrm{aff}}\big).$$

Thus rearranging promotions at fixed $(\mathscr{T}, m)$ preserves total expected demand. Increasing current promotion effects by u_a adds $p^L m \sum_a u_a$ to cycle revenue without affecting any window contribution. This latter conclusion also holds for nonlinear additive engagement.

For heterogeneous affine customers, the preceding representation uses the weighted historical slope $\beta_b^{\mathrm{het}} = \sum_{\{i:\, \mathscr{A}_i \geq b\}} \lambda_i \beta_{i,b}^{\mathrm{aff}}$, with $\beta_{\mathscr{A}_{\max}+1}^{\mathrm{het}} = 0$. A nonincreasing, nonnegative aggregate slope sequence therefore gives the same concavity and computation results, by Proposition B.1.

Corollary C.1 (Computational complexity)**.** *For affine engagement satisfying Assumption 4.1, g_0^{aff} and g_1^{aff} can be computed in $O(\mathscr{A}+1)$ arithmetic operations. For $\mathscr{A} \geq 1$, a subsequent slope sweep computes an optimal affine cycle in $O(\mathscr{A}^2 \log(\mathscr{A}+1))$ operations and $O(\mathscr{A}^2)$ auxiliary space. Generating breakpoints in increasing Farey order reduces the sweep to $O(\mathscr{A}^2)$ operations. These bounds count arithmetic operations and comparisons at unit cost. For $\mathscr{A} = 0$, preprocessing and comparison of the two constant price policies take $O(\mathscr{A}+1)$ operations.*

Proof of Corollary C.1. A single pass over ages reads the slopes and computes g_0^{aff} and g_1^{aff} using $O(\mathscr{A}+1)$ operations including initialization. If $\mathscr{A} = 0$, the objective is affine and comparing its endpoints completes the search. Assume henceforth that $\mathscr{A} \geq 1$.

At each interior breakpoint $j/(b+1)$, $j = 1, \ldots, b$, the slope of q_b^{aff} increases by $b+1$. Define the corresponding slope decrement of g^{aff} by

$$\zeta_b := \Delta p\,(b+1)\big(\beta_b^{\mathrm{aff}} - \beta_{b+1}^{\mathrm{aff}}\big).$$

Form the $O(\mathscr{A}^2)$ events $(j/(b+1), \zeta_b)$, sort them by frequency, and add decrements at coincident frequencies. Initialize the objective value and the best value at g_0^{aff}, attained at $\rho = 0$, and the slope at g_1^{aff}. Between successive event frequencies, advance the value by the current slope times the frequency increment, record the best value, and then subtract the combined decrement. Finally advance to $\rho = 1$ and compare its value with the current best. Each event is processed once, so sorting dominates the $O(\mathscr{A}^2 \log(\mathscr{A}+1))$ operation count and the event list uses $O(\mathscr{A}^2)$ space.

To avoid sorting, observe that an interior frequency u/q in lowest terms is an event for depth b exactly when q divides $b+1$. Its total slope decrement therefore depends only on the reduced denominator. Precompute

$$\delta_q := \sum_{k=1}^{\lfloor (\mathscr{A}+1)/q \rfloor} \zeta_{kq-1}, \qquad q = 2, \dots, \mathscr{A}+1.$$

The number of summands across these computations is $O(\mathscr{A} \log(\mathscr{A}+1))$. Generate the Farey sequence of order $\mathscr{A}+1$ starting from $0/1$ and $1/(\mathscr{A}+1)$, and stop at $1/1$. If $u/v < w/z$ are consecutive fractions, their successor is

$$\frac{hw-u}{hz-v}, \qquad h = \left\lfloor \frac{\mathscr{A}+1+v}{z} \right\rfloor.$$

This recurrence takes a constant number of operations per fraction. At each interior fraction with denominator q, apply the same sweep update using δ_q. There are $O(\mathscr{A}^2)$ fractions, so preprocessing the decrements and sweeping them takes $O(\mathscr{A}^2)$ operations. Constructing the maximizing balanced cycle takes $O(\mathscr{A}+1)$ further operations by Lemma 3.3. □

Example C.1 (Finding an optimal affine arrangement)**.** Let $\mathscr{A} = 4$, $p^L = 1$, and $p^H = 1.5$. Set $\alpha_a^{\mathrm{aff}} = 0.12$ for every age and $(\beta_0^{\mathrm{aff}}, \dots, \beta_4^{\mathrm{aff}}) = (0, 0.1, 0.04, 0, 0)$, initially with $\eta_a^{\mathrm{aff}} = \beta_a^{\mathrm{aff}}$. Then $\sum_a \alpha_a^{\mathrm{aff}} = 0.6$. These values give $g_0^{\mathrm{aff}} = 0.9$, $g_1^{\mathrm{aff}} = 0.11$, and nonzero window contributions only at depths one and two, so it suffices to evaluate candidate frequencies $\{0, 1/3, 1/2, 2/3, 1\}$. Thus

$$g^{\mathrm{aff}}(\rho) = 0.9 + 0.11\rho - 0.5\big[0.06 q_1^{\mathrm{aff}}(\rho) + 0.04 q_2^{\mathrm{aff}}(\rho)\big].$$

The slope is positive before $\rho = 1/2$ and negative after it, so the shortest optimal balanced cycle has $(\mathscr{T}^\star, m^\star) = (2, 1)$. The optimal arrangement alternates promotional and regular prices.

Now increase only η_0^{aff} from zero to 0.04. The historical coefficients and all window contributions are unchanged, while g_1^{aff} rises to 0.15. The maximizing frequency becomes $2/3$, implemented by the balanced cycle HLL, with average revenue 0.97. This comparison isolates the effect of the current promotion on frequency while retaining the same arrangement rule.

C.1.3 Heterogeneous lag effects

Proof of Proposition 4.2. Write $Q_b = Q_{b,2}$. Aggregate the historical coefficients across active cohorts as

$$\bar{\beta}_j^{\mathrm{loc}} = \sum_{\{a\,:\,a \geq j\}} \beta_{a,j}^{\mathrm{loc}}, \qquad j = 1, \dots, \mathscr{A}.$$

Define the lag counts

$$P_j(C) = \sum_{t=1}^{\mathscr{T}} d_t d_{t-j}.$$

Since $p_t = p^H - \Delta p\, d_t$ and $\sum_t d_{t-j} = m$, the terms involving the baseline and current promotion response are fixed by $(\mathscr{T}, m)$, as is the regular price part of each historical term. The remaining revenue term is

$$-\Delta p \sum_{j=1}^{\mathscr{A}} \bar{\beta}_j^{\mathrm{loc}} P_j(C).$$

The promotion combination counts satisfy

$$Q_b(C) - Q_{b-1}(C) = \sum_{j=1}^{b} P_j(C), \qquad P_b(C) = Q_b(C) - 2Q_{b-1}(C) + Q_{b-2}(C),$$

where $Q_0(C) = Q_{-1}(C) = 0$. With $\bar{\beta}^{\mathrm{loc}}_{\mathscr{A}+1} = \bar{\beta}^{\mathrm{loc}}_{\mathscr{A}+2} = 0$, substitution gives

$$-\Delta p \sum_{j=1}^{\mathscr{A}} \bar{\beta}^{\mathrm{loc}}_j P_j(C) = -\Delta p \sum_{b=1}^{\mathscr{A}} \left(\bar{\beta}^{\mathrm{loc}}_b - 2\bar{\beta}^{\mathrm{loc}}_{b+1} + \bar{\beta}^{\mathrm{loc}}_{b+2}\right) Q_b(C).$$

Equivalently, the arrangement contribution is $\sum_{b,t} F_b(N_{b,t})$, with

$$F_b(n) = -\Delta p \left(\bar{\beta}^{\mathrm{loc}}_b - 2\bar{\beta}^{\mathrm{loc}}_{b+1} + \bar{\beta}^{\mathrm{loc}}_{b+2}\right) \binom{n}{2}.$$

The identities above follow by counting pairs of positions in a window using $Q_b = \sum_{j=1}^{b}(b+1-j)P_j$. They remain valid with cyclic multiplicities when the window exceeds the cycle length. Under (4), every F_b is concave, so the window comparison in Theorem 3.5 proves balanced optimality. For $\mathscr{A} = 0$, the arrangement sum is empty. □

For the finite geometric profile $\bar{\beta}^{\mathrm{loc}}_b = cr^b$, the second difference is $cr^b(1-r)^2$ for $b \le \mathscr{A}-2$, $cr^{\mathscr{A}-1}(1-2r)$ for $b = \mathscr{A}-1$, and $cr^{\mathscr{A}}$ for $b = \mathscr{A}$. Thus the cutoff condition is $r \le 1/2$ when $\mathscr{A} \ge 2$. With $\mathscr{A} = 1$, the sole second difference is positive.

C.2 Reference price channel

C.2.1 Average reference prices

Proof of Proposition 4.3. If $\mathscr{A} = 0$, demand depends only on the current price, so the better constant price is optimal. Suppose $\mathscr{A} \ge 1$. Under $w_{a,j} = 1/a$, the reference price is $p^H - \Delta p\, y/a$, and the market share function becomes

$$M_a(p, y) = u_a - v_a p + \gamma_a \left(p^H - \frac{\Delta p}{a} y - p\right), \qquad a \ge 1.$$

It therefore has the additive form (2) with a linear historical response and coefficients

$$\alpha^{\mathrm{aff}}_a = u_a - v_a p^H, \qquad \eta^{\mathrm{aff}}_a = (v_a + \gamma_a)\Delta p, \qquad \beta^{\mathrm{aff}}_a = -\frac{\gamma_a \Delta p}{a}.$$

At age zero, take $\alpha^{\mathrm{aff}}_0 = u_0 - v_0 p^H$, $\eta^{\mathrm{aff}}_0 = v_0 \Delta p$, and $\beta^{\mathrm{aff}}_0 = 0$. The assumed market share bounds ensure feasibility even though historical slopes are nonpositive.

For every age $a \ge 1$, the marginal effects at both current prices are $\Delta^H_{a,k} = \Delta^L_{a,k} = \beta^{\mathrm{aff}}_a$. Hence $r_{a,k} = \Delta p\, \beta^{\mathrm{aff}}_a$, and Proposition 3.2 gives

$$\psi_{b,k} = \begin{cases} (\Delta p)^2 \left(\dfrac{\gamma_b}{b} - \dfrac{\gamma_{b+1}}{b+1}\right), & 1 \le k \le b < \mathscr{A}, \\ (\Delta p)^2 \dfrac{\gamma_{\mathscr{A}}}{\mathscr{A}}, & 1 \le k \le b = \mathscr{A}. \end{cases}$$

Both expressions are nonnegative because γ_a/a is weakly decreasing and $\gamma_a \geq 0$. Equivalently, setting $\beta^{\text{aff}}_{\mathscr{A}+1} = 0$ yields

$$F_b(n) = -\Delta p\big(\beta^{\text{aff}}_b - \beta^{\text{aff}}_{b+1}\big)\binom{n}{2}, \qquad n = 0, \dots, b+1,$$

whose second differences are nonnegative. Thus every F_b is discretely convex. Theorem 3.5(ii) gives single block optimality at each fixed $(\mathscr{T}, m)$, and Proposition 3.8 gives optimality of the better constant price over all feasible infinite price paths. □

C.2.2 Exponentially smoothed reference prices

Suppose $\mathscr{A} \geq 1$ and each cohort starts from a fixed initial reference price $r_{0,t} = \bar{r}$. At age $a \geq 1$, its reference price follows

$$r_{a,t} = \lambda r_{a-1,t-1} + (1-\lambda)p_{t-1}, \qquad 0 < \lambda < 1.$$

Retain the linear, loss neutral market share function in (5), with common sensitivity $\gamma_a = \gamma \geq 0$ at ages $1, \dots, \mathscr{A}$, and omit the reference effect at age zero. Assume market shares lie in $[0,1]$ for every feasible price history. Then the arrangement and constant price conclusions of Proposition 4.3 continue to hold.

Proof for exponential smoothing. Expanding the recursion gives

$$r_{a,t} = \lambda^a \bar{r} + (1-\lambda)\sum_{j=1}^{a} \lambda^{j-1} p_{t-j}.$$

Substituting $p_t = p^H - \Delta p\, d_t$ yields the lag-dependent affine form in Section 4.1.2. For $a \geq 1$, its coefficients are

$$\alpha^{\text{loc}}_a = u_a - v_a p^H + \gamma\lambda^a(\bar{r} - p^H), \qquad \beta^{\text{loc}}_{a,0} = (v_a + \gamma)\Delta p,$$
$$\beta^{\text{loc}}_{a,j} = -\gamma\Delta p(1-\lambda)\lambda^{j-1}, \qquad j = 1, \dots, a.$$

At age zero, take $\alpha^{\text{loc}}_0 = u_0 - v_0 p^H$ and $\beta^{\text{loc}}_{0,0} = v_0\Delta p$. Aggregating the historical coefficients across active cohorts gives

$$\bar{\beta}^{\text{loc}}_j = -\gamma\Delta p(1-\lambda)(\mathscr{A} - j + 1)\lambda^{j-1}, \qquad j = 1, \dots, \mathscr{A}.$$

Set $\bar{\beta}^{\text{loc}}_{\mathscr{A}+1} = \bar{\beta}^{\text{loc}}_{\mathscr{A}+2} = 0$. The lag-count identities in Section 4.1.2 are algebraic and remain valid for these negative historical coefficients. They give the arrangement-dependent window contributions

$$F_b(n) = -\Delta p\left(\bar{\beta}^{\text{loc}}_b - 2\bar{\beta}^{\text{loc}}_{b+1} + \bar{\beta}^{\text{loc}}_{b+2}\right)\binom{n}{2}, \qquad n = 0, \dots, b+1.$$

Their second differences are independent of the promotion count and equal

$$-\Delta p\left(\bar{\beta}^{\text{loc}}_b - 2\bar{\beta}^{\text{loc}}_{b+1} + \bar{\beta}^{\text{loc}}_{b+2}\right) = \begin{cases} \gamma(\Delta p)^2(1-\lambda)\lambda^{b-1}\left[(\mathscr{A}-b)(1-\lambda)^2 + 1 - \lambda^2\right], & 1 \leq b < \mathscr{A}, \\ \gamma(\Delta p)^2(1-\lambda)\lambda^{\mathscr{A}-1}, & b = \mathscr{A}. \end{cases}$$

Both expressions are nonnegative, including at the finite lifetime cutoff. Thus every F_b is discretely convex. The convex-window comparison in Theorem 3.5, applied to this lag-dependent representation, establishes single block optimality at each fixed $(\mathscr{T}, m)$.

For unrestricted frequency, apply the endpoint chord bound to each F_b, as in Proposition 3.8. The remaining terms in average cycle revenue are affine in frequency, so every cycle C with frequency ρ satisfies

$$\bar{R}(C) \leq (1-\rho)\bar{R}(C^H) + \rho\bar{R}(C^L) \leq \max\{\bar{R}(C^H), \bar{R}(C^L)\},$$

where $C^H = (p^H)$ and $C^L = (p^L)$ are the two constant price cycles. Each cohort starts from the fixed reference $\bar{r}$ and remains active only through age $\mathscr{A}$, so current rewards depend on at most the current and previous $\mathscr{A}$ prices. The finite graph argument of Proposition 2.1 therefore extends this bound to all feasible infinite price paths. The initial reference can affect which constant price is better through the baseline demand terms, but does not affect window convexity. When $\mathscr{A} = 0$, history is absent and the better constant price is optimal directly. □

C.2.3 General reference weights

Constant price optimality also holds for any nonnegative reference weights satisfying $\sum_{j=1}^{a} w_{a,j} = 1$, without the age monotonicity condition used for the fixed-frequency arrangement result. Cyclic reindexing gives

$$\sum_{t=1}^{\mathscr{T}} p_t(p_{t-j} - p_t) = -\frac{1}{2}\sum_{t=1}^{\mathscr{T}}(p_t - p_{t-j})^2.$$

Consequently, for the linear, loss neutral model in (5),

$$\begin{aligned}\bar{R}(C) = \frac{1}{\mathscr{T}}\sum_{t=1}^{\mathscr{T}}\sum_{a=0}^{\mathscr{A}} p_t(u_a - v_a p_t) \\ - \frac{1}{2\mathscr{T}}\sum_{a=1}^{\mathscr{A}}\gamma_a\sum_{j=1}^{a} w_{a,j}\sum_{t=1}^{\mathscr{T}}(p_t - p_{t-j})^2.\end{aligned} \tag{C.3}$$

The first term is the frequency-weighted average of the two constant price revenues, and the subtracted term is nonnegative. The better constant price is therefore optimal over cycles and, by Proposition 2.1, over all feasible infinite price paths. This argument establishes constant price optimality for general weights. It does not assert convexity of every window contribution.

D Proofs and supplementary results for Section 5

D.1 Multiplicative window contributions and arrangement regimes

Proof of Corollary 5.1. Since $M_a^H(y) = s_2^{a+1}\theta^y$ and $M_a^L(y) = \theta M_a^H(y)$, their successive differences are

$$\Delta_{a,k}^H = s_2^{a+1}(\theta - 1)\theta^{k-1}, \qquad \Delta_{a,k}^L = \theta\Delta_{a,k}^H.$$

The revenue comparison of marginal effects in Section 3.2 is therefore

$$r_{a,k} = p^H\Delta_{a,k}^H - p^L\Delta_{a,k}^L = p^L(\kappa - \theta)\Delta_{a,k}^H.$$

For $1 \le k \le b < \mathscr{A}$, Proposition 3.2 gives

$$\begin{aligned}\psi_{b,k} &= p^H(\Delta_{b+1,k+1}^H - \Delta_{b+1,k}^H) + r_{b+1,k} - r_{b,k} \\ &= p^L s_2^{b+1}(\theta - 1)\theta^{k-1}[\theta(1 - s_2) - \kappa(1 - s_1)] \\ &= p^L s_2^b(\theta - 1)\theta^{k-1}S_b^{\mathrm{mul}}(\kappa).\end{aligned}$$

At the terminal depth, the same proposition yields

$$\psi_{\mathscr{A},k} = -r_{\mathscr{A},k} = p^L s_2^{\mathscr{A}+1}(\theta - 1)\theta^{k-1}(\theta - \kappa) = p^L s_2^{\mathscr{A}}(\theta - 1)\theta^{k-1}S_{\mathscr{A}}^{\mathrm{mul}}(\kappa).$$

Finally, use $F_b(n) = \sum_{k=1}^{b}(n-k)_+\psi_{b,k}$ and the finite geometric identity

$$\sum_{k=1}^{b}(n-k)_+\theta^{k-1} = \frac{\theta^n - 1 - n(\theta-1)}{(\theta-1)^2}, \qquad n = 0, \dots, b+1.$$

Substitution gives the stated formula for F_b, including $F_b(0) = F_b(1) = 0$. □

Proof of Proposition 5.2. By Corollary 5.1, the sign of every $\psi_{b,k}$ at depth b equals the sign of $S_b^{\mathrm{mul}}(\kappa)$. At the terminal depth,

$$S_{\mathscr{A}}^{\mathrm{mul}}(\kappa) = s_2(\theta - \kappa).$$

For every interior depth,

$$S_b^{\mathrm{mul}}(\kappa) = s_2[\theta(1-s_2) - \kappa(1-s_1)], \qquad b < \mathscr{A}.$$

If $\kappa \le \theta$, the terminal factor is nonnegative, and each interior factor is positive because

$$\theta(1-s_2) - \kappa(1-s_1) \ge \theta(s_1 - s_2) > 0.$$

Thus every F_b is convex, and Theorem 3.5(ii) proves part (i).

When $\mathscr{A} \ge 2$, interior factors are nonpositive if and only if $\kappa \ge \tau^{\mathrm{mul}}$. Since $\tau^{\mathrm{mul}} > \theta$, the terminal factor is then negative as well. When $\mathscr{A} = 1$, only the terminal factor remains, and it is nonpositive if and only if $\kappa \ge \theta$. These are precisely the conditions for every F_b to be concave, so Theorem 3.5(i) proves part (ii).

The remaining region has $\mathscr{A} \ge 2$ and $\theta < \kappa < \tau^{\mathrm{mul}}$. All interior second differences are positive and all terminal second differences are negative, establishing part (iii). In the convex regime, Proposition 3.8 gives a constant price optimum. Moreover, $p^L s_1^{a+1} \ge p^H s_2^{a+1}$ at every age, so constant promotion is globally optimal. In the concave regime, Proposition 3.7 gives an optimal balanced cycle with period at most $\mathscr{A} + 1$. □

Scope of the structural conditions. The short-period conclusion requires concave window contributions. For the multiplicative model, let $\mathscr{A} = 2$, $p^L = 1$, $p^H = 5$, $s_1 = 3/5$, and $s_2 = 1/5$. Direct enumeration of all cycles of period at most three gives maximum revenue $473/375$, whereas

$$\bar{R}(LLHH) = \frac{158}{125} > \frac{473}{375}.$$

Thus an optimal period need not be at most $\mathscr{A} + 1$ outside the concave regime. This example also rules out domination by mixtures of shorter cycle revenues.

Single block optimality at fixed frequency likewise needs a shape condition. With $\mathscr{A} = 1$, $p^L = 1$, $p^H = 2$, $s_1 = 3/10$, and $s_2 = 1/5$,

$$\bar{R}(LHLH) = \frac{11}{25} > \frac{7}{16} = \bar{R}(LLHH).$$

These parameters give a concave window contribution, so the comparison agrees with balanced optimality.

E Proofs and supplementary results for Section 6

E.1 Balanced and single block guarantees

Proof of Proposition 6.1. Fix a cycle C with frequency ρ. For $l = b + 1$, write

$$H_{b,k}(C) := \frac{1}{\mathscr{T}} \sum_t (N_{b,t} - k)_+.$$

Convexity and the endpoint chord of the hinge give

$$(l\rho - k)_+ \le H_{b,k}(C) \le \rho(l-k).$$

The lower bound is attained by balanced cycles because the hinge is linear between consecutive integers. Set

$$e_{b,k}(\rho) := \rho(b+1-k) - ((b+1)\rho - k)_+ = \min\{(b+1-k)\rho, k(1-\rho)\}.$$

Its maximum over $[0,1]$ is $k(b+1-k)/(b+1)$. Multiplying the lower hinge bound by nonpositive coefficients and the upper bound by positive coefficients in the revenue decomposition gives

$$\bar{R}(C) \le g(\rho) + \sum_{b,k} \max\{\psi_{b,k}, 0\} e_{b,k}(\rho) \le \bar{R}(C_B) + E_B.$$

Here $g(\rho) \le \max_{\mathscr{B}} g$ follows from piecewise linearity without a concavity assumption. Every frequency in $\mathscr{B}$ has reduced denominator at most $\mathscr{A}+1$, giving the period bound.

We first establish that the single block search is finite. Put $r_H = \bar{R}(H)$ and $r_L = \bar{R}(L)$. For $\mathscr{A} \ge 1$, consider any $H^h L^m$ with $h, m \ge 1$, and let $h' = \min\{h, \mathscr{A}\}$, $m' = \min\{m, \mathscr{A}\}$. Beyond the first $\mathscr{A}$ periods of a constant-price run, all remembered prices equal the current price. Removing such periods leaves every remaining history unchanged, so

$$(h+m)\bar{R}(H^h L^m) = (h'+m')\bar{R}(H^{h'} L^{m'}) + (h-h')r_H + (m-m')r_L.$$

Thus its average revenue is a convex combination of the shortened cycle revenue and the two constant-price revenues. Maximizing over $\mathscr{S}$ therefore finds a best single block policy over all lengths, with at most $\mathscr{A}^2 + 2$ candidates and period at most $2\mathscr{A}$.

For its revenue guarantee, endpoint evaluation of the decomposition gives

$$(1-\rho)r_H + \rho r_L = g_0 + g_1\rho + \sum_{b,k} \psi_{b,k}\rho(b+1-k).$$

The upper hinge bound controls positive coefficients, while the lower bound controls negative coefficients. Therefore

$$\bar{R}(C) \le (1-\rho)r_H + \rho r_L + \sum_{b,k} \max\{-\psi_{b,k}, 0\} e_{b,k}(\rho) \le \max\{r_H, r_L\} + E_S \le \bar{R}(C_S) + E_S.$$

The better of C_B and C_S has loss bounded by each of E_B, E_S, proving the combined guarantee. Proposition 2.1 extends both cycle comparisons to all feasible infinite price paths. When $\mathscr{A} = 0$, both searches reduce to the constant prices, the sums are empty, and both candidates are optimal. □

Verification of the examples. In the reinforcement example, the shares lie in $[0.10, 0.30]$ for the stated parameter range. Interior curvatures are 2ϵ, and terminal curvatures are $-(0.04 + \epsilon(k-1))$, giving $E_B = 11\epsilon/3$. In the fatigue example, shares lie in $[0.19, 0.404]$, and the curvatures are $(-2\epsilon, 0, \epsilon)$, giving $E_S = \epsilon$. Direct evaluation of the candidate frequencies and single block calendars gives the reported policy revenues.

At $\epsilon = 0.001$, exact node potentials certify the unrestricted optima. Encode the previous $\mathscr{A}$ indicators as a binary integer s, oldest first, and let $w(s,d)$ be current revenue. For the reinforcement example, set

$$\lambda = \frac{1803}{2000}, \qquad h = \frac{1}{2000}(0, 523, 162, 886, 0, 605, 281, 1054).$$

For the fatigue example, set

$$\lambda = \frac{2401}{2000}, \qquad h = \frac{1}{2000}(2, 1, 2, 0).$$

In each case, every edge satisfies

$$\lambda + h(s) \geq w(s,d) + h((2s+d) \bmod 2^{\mathscr{A}}).$$

The stated optimal cycle attains λ, so summing the inequalities proves optimality over all cycles and infinite price paths.

E.2 General memory truncation

Proof of Proposition 6.2. For each affected age, the definition of $\Delta^{\mathrm{tr}}_{a,J}$ bounds the demand error for every retained and omitted history. Multiplication by the current price, summation over ages, and averaging give the uniform revenue bound. For a full-model maximizer $C^\star$,

$$\bar{R}(C^\star_J) \geq \bar{R}_J(C^\star_J) - \varepsilon^{\mathrm{agg}}_J \geq \bar{R}_J(C^\star) - \varepsilon^{\mathrm{agg}}_J \geq \bar{R}^* - 2\varepsilon^{\mathrm{agg}}_J.$$

If $\bar{R}_J \geq \bar{R}$ pointwise, then $\bar{R}(C^\star_J) \geq \bar{R}^\star_J - \varepsilon^{\mathrm{agg}}_J \geq \bar{R}^* - \varepsilon^{\mathrm{agg}}_J$. If $\bar{R}_J \leq \bar{R}$ pointwise, then $\bar{R}(C^\star_J) \geq \bar{R}^\star_J \geq \bar{R}^* - \varepsilon^{\mathrm{agg}}_J$. Both objectives have finite memory and bounded rewards, so maximizing cycles exist. The per-period error bound also applies directly to every infinite path, proving the stated scope. □

Midpoint and fading-activity bounds. For fixed (a,p,y), let L and U be the smallest and largest compatible full-model shares. Any approximating value v has error at least $\max\{|L-v|, |U-v|\} \geq (U-L)/2$. The midpoint attains this lower bound and lies in $[0,1]$. For $M_a = q^a m_a$, both full and truncated shares lie in $[0, q^a]$, so $\Delta^{\mathrm{tr}}_{a,J} \leq q^a$. Summing the geometric tail gives the displayed policy bound.

Corollary E.1 (A posteriori certification for candidate cycles)**.** *Under Proposition 6.2, set* $U_J = \bar{R}^\star_J + \varepsilon^{\mathrm{agg}}_J$. *For every feasible cycle* C,

$$\bar{R}^* \leq U_J, \qquad \bar{R}^* - \bar{R}(C) \leq U_J - \bar{R}(C).$$

If $\bar{R}(C) > 0$, *then* $\bar{R}(C)/\bar{R}^* \geq \bar{R}(C)/U_J$. *When the truncated model bounds full-model demand from above, one may instead use* $U_J = \bar{R}^\star_J$.

Proof. The upper bound follows from the uniform error comparison, or directly from the one-sided comparison. Subtract the candidate's true revenue for the additive certificate, and divide when it is positive for the ratio certificate. □

Multiplicative demand.

Proof of Proposition 6.3. For $a > J$, let y and z count promotions in the retained J and omitted $a-J$ periods. Then

$$M^{\mathrm{mul}}_a(p, y+z) = M^{\mathrm{mul}}_{a,J}(p,y)\, s_1^z s_2^{a-J-z}.$$

The omitted factor lies in $[s_2^{a-J}, 1]$, and $M^{\mathrm{mul}}_{a,J}(p,y) \leq s_1^{J+1}$. Therefore

$$0 \leq M^{\mathrm{mul}}_{a,J}(p,y) - M^{\mathrm{mul}}_a(p, y+z) \leq s_1^{J+1}(1 - s_2^{a-J}).$$

The upper bound is attained at $p = p^L$, $y = J$, and $z = 0$. Apply the one-sided part of Proposition 6.2, divide by $\mathscr{A}+1$, and bound the number of affected ages by $\mathscr{A}+1$. This proves both the normalized error and the geometric policy bound. □

F Supplementary numerical results for Section 7

F.1 Reference price robustness

Table A.2 supplements Section 7.2. Each positive-γ row pools three reference smoothing parameters and two loss-to-gain ratios, with 100 replications per case. Population and Constant average case-level regrets. The fitted-policy mean, P95, and maximum use the pooled replications. Population excludes estimation noise. The low average losses coexist with the failure case discussed in the main text, so this experiment does not establish robustness beyond the specified demand grid.

Table A.2: True-demand regret (%). Rows with $\gamma > 0$ pool six cases and 100 replications per case. Population excludes sampling noise. Mean, P95, and Max summarize fitted-policy regret. Constant uses the better true constant price.

Channel	γ	Cases	Population	Mean	P95	Max	Constant
Reference only	0.0	1	0.000	0.000	0.000	0.000	0.000
Reference only	0.5	6	0.000	0.000	0.000	0.000	0.000
Reference only	1.5	6	0.174	0.225	1.045	2.413	2.511
With engagement	0.0	1	0.000	0.049	0.201	0.201	5.279
With engagement	0.5	6	0.000	0.050	0.480	0.806	4.006
With engagement	1.5	6	0.000	0.042	0.480	0.480	8.300

F.2 Memory truncation at a common cutoff

Table A.3 supplements Section 7.3 and Figure 4. At $J = 2$, every approximation uses four past-price states. General rows pool the two general demand grids and maximum ages 4, 6, 8, while multiplicative rows pool the price and regular-retention settings at those ages. Loss and Bound are separate maxima relative to the full-model optimum and can come from different cases. The last column reports how often the relative bound is below 1%. Bounds above 100% are retained to show their conservatism. They do not imply that actual loss exceeds 100%.

Table A.3: Truncation at $J = 2$. Loss and Bound are separate maxima (%). The last column gives the percentage of cases with bound below 1%. General rows pool both response grids. Bounds are uncapped.

Setting	Cases	Loss at $J = 2$	Bound at $J = 2$	Bound $< 1\%$
General, $q = 0.4$	306	0.133	2.095	79.7%
General, $q = 0.7$	306	1.590	11.984	8.8%
General, $q = 0.95$	306	6.667	30.873	1.0%
General, $q = 1$	306	6.667	35.000	1.0%
Multiplicative, $s_1 = 0.4$	18	0.000	138.037	0.0%
Multiplicative, $s_1 = 0.7$	18	4.314	313.784	0.0%
Multiplicative, $s_1 = 0.9$	18	19.184	451.343	0.0%
Multiplicative, $s_1 = 0.98$	18	3.101	495.756	0.0%