Setting the Optimal Value of Loyalty Points

So Yeon Chun Dan A. Iancu Nikolaos Trichakis*

Abstract

A loyalty program introduces a new currency—the points—through which customers transact with a firm. We study the problem of optimally setting the monetary value of points, i.e., pricing in this new currency, in a multi-period setting. We first show that point pricing is different from cash pricing primarily due to the way points are accounted for, as liabilities on the firm's balance sheet. This introduces subtle channels through which the firm's decisions affect its financial performance, and exacerbates the importance of certain managerial considerations such as taxation or earnings smoothing incentives.

We characterize the optimal cash and point pricing policies, and find that they mimic "base-stock, list price" policies in inventory management. In particular, point prices/values are always set so that the total value of points reaches a "base-stock" target, and cash prices are charged so as to maximize the firm's cash flows under the optimal loyalty point values. Under a profit-maximizing policy, the total value of loyalty points is set independently of the firm's realized financial performance. In contrast, we find that under the aforementioned managerial considerations, the optimal value of points becomes state-dependent, and is increasing (decreasing) under strong (weak) operating performance. In this sense, our work shows that loyalty points can act as a hedging tool against uncertainty in future performance, providing a new rationale for their existence, even in the absence of competition.

1 Introduction

Originally designed as marketing tools for rewarding customers, loyalty programs have witnessed a dramatic expansion in size and scope during recent years, with memberships in the U.S. tripling between 2000 and 2014 (reaching 3.3 billion, or 10.3 on average per individual Berry 2015), and covering a wide array of industries, including retail (39%), travel and hospitality (27%), and financial services (17%).

*Chun is with the McDonough School of Business at Georgetown University, (soyeon.chun@georgetown.edu), Iancu is with the Graduate School of Business at Stanford University, (daniancu@stanford.edu), and Trichakis is with the Harvard Business School, (ntrichakis@hbs.edu).

In a typical "point-based" loyalty program (LP), members earn points for their purchases of products or services from an issuing firm, and are able to redeem accumulated points for awards, such as additional products, services or even cash. The existence of an LP thus effectively introduces a new *currency*—the loyalty points—through which customers can transact with the firm (Economist 2005*b*).[1] The monetary value of points is controlled by the firm, by explicitly setting an exchange rate or by posting "prices"—i.e., point requirements—in this new currency. For instance, airlines such as United or Delta dynamically change the mileage requirements for award tickets, charging anywhere from 10,000 to 25,000 miles for a one-way economy flight. Similarly, credit cards issuers operating LPs such as Chase or American Express routinely adjust the point requirements for the products and services offered for redemption on their reward platforms.

Understanding how these point valuation/pricing decisions should be taken optimally is the main focus of our paper. To that end, we first identify the critical ways in which setting prices in this new currency differs from—and impacts—traditional cash pricing, which will allow us to formulate a concrete set of research questions, and to propose a model that can address them.

To start, setting prices in both points and cash carries nontrivial implications on the firm's sales and revenues. For instance, reducing the point requirements could reduce cash sales—due to excessive redemptions—but could also increase them, as more customers purchase the products to earn the more valuable points. These *cannibalization* and *loyalty* effects, which have both been observed in practice (see Kopalle et al. 2012 and Lewis 2004, respectively), suggest that pricing in cash and points becomes akin to multi-product pricing (Maglaras and Meissner 2006), with both substitutability and complementarity effects possible.

More importantly, however, the differences between pricing in points and cash stem from the way LPs are accounted for: loyalty points constitute a liability on the issuing firm's balance sheet, as they are a promise for future service. Specifically, the unique financial accounting rules pertaining to LPs introduce new channels through which the firm's cash and point prices can affect its financial performance, and potentially exacerbates certain distortions in managerial decision-making.

To elaborate on these ideas, it is thus important to first understand the accounting standards for LPs. Under rules recently set by the International Financial Reporting Standard (IFRS),[2] a

[1] A significant monetary value is estimated to circulate as LP points; for example, frequent flier miles were already considered the world's most valuable currency in 2005, with an estimated 14 trillion miles outstanding worth more than \$700B (Economist 2005*a*), and the value of new rewards issued in the US alone every year exceeding \$48B (Gordon and Hlavinka 2015).

[2] The first LP-related rules, issued under "IFRIC 13 Customer Loyalty Programmes" in 2007 were a required standard in Europe, Canada, and Australia, and were also adopted by several U.S. firms. The new standard "IFRIC 15 Revenue from Contracts with Customers" will be required in the U.S. starting in 2017.

firm is required to treat any points issued in connection with a cash sale as a separate component of the sale. As such, a part of the sale's revenue is *deferred* and treated as a liability instead, which *decreases* the firm's profits. However, when points are redeemed or expire, the firm *recognizes* a corresponding amount from its deferred revenue liabilities as revenue, which *increases* its profits. IFRS guidelines stipulate that the deferred and recognized revenue amounts should reflect the points' value, i.e, the monetary value of the rewards for which the points can be redeemed. To that end, IFRS guidelines require the total value of a firm's LP-related deferred revenue to be calculated as the product of three terms: the total number of outstanding points, the value of a point, and the probability that the point will be redeemed (also referred to as the *redemption rate*).

In view of these rules, it can be seen how pricing in points impacts profitability due to the deferral process: increasing (decreasing) the value of a point translates into more (less) deferred revenue, which directly hurts (improves) profitability. For firms with billions of outstanding points, even a small change in the point value can thus have a *first-order* impact on profits. To illustrate this, consider the case of Delta Airlines, which was carrying \$4.2B deferred revenue liabilities in 2014: a 1% decrease in the value of points (implemented, for instance, by increasing the point requirements or by making fewer seats available for redemption, on average) would have increased its \$659M profits that year by \$42M, a boost of 6.4%. Such examples are not just hypothetical, but actually occur in practice: in 2008, Alaska Airlines shortened the expiration date of its points from three years to two, thus reducing their total value and the associated deferred revenue. This enabled the airline to claim an additional \$42.3M in revenue, and reduce its consolidated net losses for the year by a staggering 24%—see Item 7 of the 10K report, available at Alaska Airlines 2008.

The Alaska Airlines example also serves to highlight how reducing the deferred revenue associated with the loyalty program liabilities can improve the firm's operating performance in otherwise poor quarters. This suggests that earnings smoothing incentives can become particularly pertinent when considering point pricing/valuation decisions. Similarly, since the deferral process influences the revenue taxable year of inclusion, taxation can become another important managerial consideration influencing loyalty point valuations[3] (see Section 5 for a more detailed discussion).

In summary then, the subtle loyalty/cannibalization effects, the important accounting rules

[3]To further illustrate the significance of taxation when a firm operates an LP, we note that the U.S. Department of the Treasury included the relevant article (§451) of the U.S. Tax Law that determines taxable year of inclusion in its 2014-2015 Priority Guidance Plan. While this action was solely an indication that the law may undergo future changes (without specifying whether/when/what the changes might be), it prompted an immediate response from multiple trade organizations, including Airlines for America, the American Hotel & Lodging Association, the U.S. Travel association, etc., who wrote in an open letter to the Treasury Secretary Jacob Lew that *"Any change in accounting rules could result in billions of dollars in lost revenue to states and localities, as well as significant harm to small-business franchise owners."* (AHLA 2014)

governing loyalty programs, and the aforementioned managerial considerations all underscore the difficulty and unique challenges when setting point values/prices. Although this problem has been recognized by practitioners from a wide range of industries (see, e.g., Oracle 2008, SAS 2012, Ernst&Young 2014, and PwC 2015), to the best of our knowledge, little or no academic work has been done exploring optimal policies in this setting.

The present paper attempts to fill this gap by posing the following related research questions: How should a firm's manager adjust the value of a loyalty point, taking into account the important considerations highlighted above? When should point value increase or decrease, and how should cash prices be adjusted? How would the manager's policies be impacted by considerations such as taxation, earnings smoothing incentives, or risk aversion?

We address these questions by developing a dynamic model that integrates the practical details outlined so far. In particular, we study a firm that sells a single type of product, and awards customers who purchase in cash with points that can be exchanged for products. Reporting of cash flows and profits is subject to IFRS specifications. The firm's manager dynamically sets cash prices and point requirements over a discrete and finite time horizon, so as to maximize expected discounted rewards tied to profits. We consider both linear reward functions, so as to capture profit maximization, as well as concave reward functions, so as to capture the effects of tax considerations, earnings smoothing incentives, or risk aversion. We also extend our base model to allow for more complex firm operations, or rewards tied to both profits and cash flows.

Our findings and contributions. Our paper is the first to study how to dynamically adjust the monetary value of loyalty points. We formulate the problem as a Dynamic Program (DP) that captures the aforementioned important LP considerations, and that becomes tractable under mild assumptions.

We derive the optimal policies, and demonstrate that the firm's manager should adjust the value of a loyalty point so that the value of all outstanding points (i.e., the firm's LP-related deferred revenue) hits a particular target. More precisely, we show that the total value of loyalty points can be operationally thought of as "inventory," with the optimal policy mimicking the well-known "base-stock, list-price" policy in operations management, whereby the cash prices are set so as to maximize the cash flows under the optimal inventory "base-stock." Such policies are well understood in practice, and are easy to calibrate, communicate, and implement. This connection also enables us to better highlight—in a somewhat more familiar setting—the tradeoffs in determining the value of points. On the one hand, maintaining excessively valuable points (i.e., holding "too much inventory") results in increased costs due to redemption, cannibalization, and time-value loss of money—corresponding to inventory overage/holding costs. On the other hand, maintaining less valuable points (i.e., holding "too little inventory") results in opportunity costs in the form of lost sales due to a weak loyalty effect—corresponding to underage/backlogging costs.

We find that under a profit-maximizing policy, the base-stock of total point value is independent of prior operating performance. In contrast, we find that under managerial considerations such as taxation, earnings smoothing incentives, or risk aversion, the base-stock becomes state-dependent, increasing (decreasing) under strong (weak) operating performance. This exactly shows how the deferred revenues associated with the loyalty program can act as a buffer against uncertainty in operating performance: when facing strong performance, the manager adjusts the value of points so as to defer a larger portion of the revenue for future access; when facing mediocre performance, the manager recognizes more of the deferred revenue, thus boosting the current profits. In this sense, our work provides a new rationale for loyalty programs: while traditionally viewed as means of softening competition (Kim et al. 2004), we show that such programs can be beneficial even in the absence of competition, due to their hedging capability.

Using our models, we also derive a number of insightful comparative statics. We show that higher discounting and higher redemption servicing costs lead to a lower value of points. Under profit-maximizing policies, we find that the points' value is set independently of the variability in future cash flows; in contrast, under considerations such as taxation, earnings smoothing incentives, or risk aversion, the points' value *increases* with variability.

Finally, we show that our main findings are robust under several extensions to our base model, e.g., when the firm runs more complex operations (with multiple decisions that are updated more frequently), or when the manager's rewards are tied to both cash flows and profits.

1.1 Literature Review

Our paper contributes to the large literature on multi-period dynamic pricing and revenue management (see, e.g., Talluri and van Ryzin 2005 and Bitran and Caldentey 2003 for reviews and additional references). By integrating pricing in the two currencies—cash and point—we introduce several unique considerations that are typically absent in this literature, as we discussed above. Our paper also draws several connections with the broader revenue management literature on coordinating pricing and inventory decisions (see, e.g., Elmaghraby and Keskinocak 2003, Chen and Simchi-Levi 2004, and references therein). While our focus is different, we find that the optimal policies for pricing in points and cash resemble policies for replenishing and pricing inventory.

Our paper is also related to the growing body of literature integrating broadly-defined concepts from revenue management and customer relationship management (for a general review, see Tang and Teck 2004, Coughlan and Shulman 2010, Tang 2010, and references therein). Aflaki and Popescu (2014) propose a dynamic model where retention depends on customer satisfaction, and characterize optimal policies for maximizing the expected lifetime value of a customer. Afeche et al. (2015) study profit-maximizing policies for an inbound call center with abandonment by controlling customer acquisition, retention and service quality via promotions, priorities, and staffing. Similarly, Ovchinnikov et al. (2014) study the effect of limited capacity on a firm's optimal acquisition and retention policies. Specific to loyalty programs and closer to our work, Kim et al. (2004) study the interaction between LPs and capacity decisions in a competitive environment, showing that accumulated reward points could be used to reduce excess capacities in a period of low demand. Sun and Zhang (2014) study the problem of optimally setting the expiration date of points, and show that this can be used as a price-segmentation mechanism, and Baghaie et al. (2015) design optimal policies for setting reward levels in an LP using social media. Chung et al. (2015) present a dynamic model in which customers choose whether to purchase using cash or points, and investigate the impact of reimbursement terms for redemptions on the firm's pricing and inventory decisions. Chun and Ovchinnikov (2015) consider the recent change in the airline industry from a "mileage-based" to a "spending-based" design, and study the impact of strategic customer behavior on the firm's optimal pricing and premium-status LP qualification requirements. Lu and Su (2015) also study the same two LP designs for a firm setting capacity limits for loyalty awards in a classical Littlewood two-type model; they find that LPs allow firms to extract high valuations from low type customers, and that the switch to a "spending-based" design could be profitable. In contrast to these, our paper is the first to provide a holistic, dynamic model for optimally pricing with loyalty points and cash, integrating relevant practical considerations such as financial accounting.

For this reason, our paper is also related to the extensive literature in (financial) accounting that studies income smoothing, a form of earnings management. We refer the reader to Dechow et al. (1995), Healy and Wahlen (1999), Leuz et al. (2003) for reviews of this topic. Our work is closer to *real* earnings management, which is the practice of altering earnings by changing operational decisions, as opposed to pure accounting or reporting manipulation. We contribute to this literature by showing how decisions related to an LP can serve as a tool for earnings management.

Dealing with operating decisions under financial accounting considerations also relates our paper to the growing literature on the interface of finance and operations (see, e.g., Xu and Birge 2008, Caldentey and Haugh 2009, Babich 2010, Li et al. 2013, Kouvelis et al. 2013, Dong and Tomlin

2012, and references therein) We differ from this literature in both focus, as well as questions: we examine loyalty programs, and their specific associated accounting considerations.

Our modeling assumptions are motivated by several empirical papers documenting the positive impact of LPs on sales (revenues), for firms in financial services (Verhoef 2003), retail (Lewis 2004, Liu 2007), as well as travel and hospitality (Lederman 2007). Of particular relevance to our work, Taylor and Neslin (2005) and Smith and Sparks (2009) study the "loyalty effect," empirically illustrating that LPs can increase sales through two separate mechanisms ("points pressure effect," whereby customers purchase more in an effort to earn a reward, as well as "rewarded behavior effect," whereby customers purchase more after receiving a reward), and can also increase (the rate of) redemptions. Dorotic et al. (2014) and Kopalle et al. (2012) provide evidence that higher sales may lead to higher redemptions, and raise the issue of potential sales cannibalization, highlighting that setting the right point requirements involves complex trade-offs. Our model is aligned with these empirical results and is flexible, capturing the relevant dynamics of both the "loyalty" and "cannibalization" effects. Furthermore, our results show that sales cannibalization can, in fact, be the outcome of optimal behavior under some conditions.

2 Model

Consider a firm run by a manager over a discrete time-frame of $T+1$ periods, indexed by $t \in \{1, \ldots, T+1\}$. A period in our model corresponds to a fiscal period, e.g., a financial quarter or year. We make the exact sequencing and timing precise below, once we introduce all events.

The firm is selling a single type of product to its customers, operating as a monopoly. The product can be produced and delivered at zero cost, and is perishable, so that the firm does not carry any unused inventory across successive periods. The firm also runs a loyalty program (LP), whereby all customers who purchase products using cash are automatically awarded points. We use w_t to denote the balance of outstanding points at the beginning of period t. Points never expire, and can be redeemed to acquire more units of the same product, with any such redemption causing the firm to incur a per-unit servicing cost of c.

The firm's customers can acquire products by purchasing in cash or by redeeming points. During period t, we denote by p_t the unit cash price charged by the firm, and by q_t the number of points required in exchange for one product, i.e., the *point requirement* or *point price*. Equivalently, since any point requirement induces a monetary value of $\theta_t = \frac{p_t}{q_t}$ for each point, we can also consider the firm's decisions as the *cash price* p_t and the *point value* θ_t.

During period t, the firm's customers buy s_t products in cash, and acquire r_t products by redeeming points; both the *cash sales* s_t and *redemptions* (or *point sales*) r_t are random, and depend on the cash price and point value, and the number of outstanding points. We make no assumptions concerning the monotonicity of these dependencies, and only require that the randomness is independent across time.

In connection with the cash sales, the firm awards points to its customers at a given rate of λ points for every dollar spent, resulting in a total of $\lambda\, p_t\, s_t$ new points issued during period t. In contrast, redemptions result in a total of $q_t r_t$ points deducted from customer accounts, so that the balance of outstanding points at the end of period t (and beginning of period $t+1$) becomes:

$$w_{t+1} = w_t + \lambda\, p_t\, s_t - q_t\, r_t. \tag{1}$$

Revenues, costs, and profits. In period t, the firm generates *sales revenue* of $p_t\, s_t$. Adjusting for the deferred components associated with the newly issued and redeemed points, the firm's *revenues* at the end of period t are

$$\text{revenues} = (\text{sales revenue } p_t\, s_t) - (\text{newly deferred revenue}) + (\text{newly recognized revenue}).$$

If we let L_t denote the total value of the firm's deferred revenue in the beginning of period t, we can rewrite the equation above as

$$\text{revenues} = p_t\, s_t + L_t - L_{t+1}, \tag{2}$$

since the difference $L_{t+1} - L_t$ between the firm's total deferred revenues in periods $t+1$ and t is precisely equal to the newly deferred revenue net of the newly recognized revenue in period t.

In accordance with the IFRS rules concerning the calculation of LP-related deferred revenue, the total value of the firm's deferred revenue in period t is equal to the product of three terms: the total number of points w_t, the value of a point θ_t, and the redemption rate g_t. That is,

$$L_t = w_t\, \theta_t\, g_t. \tag{3}$$

Similarly to the sales, the redemption rate g_t depends on the cash and point prices, and the number of outstanding points—we make these dependencies explicit in the next section. We shall also refer to L_t as the *value of the LP* in period t. It is worth noting that by equations (2) and (3), the firm's

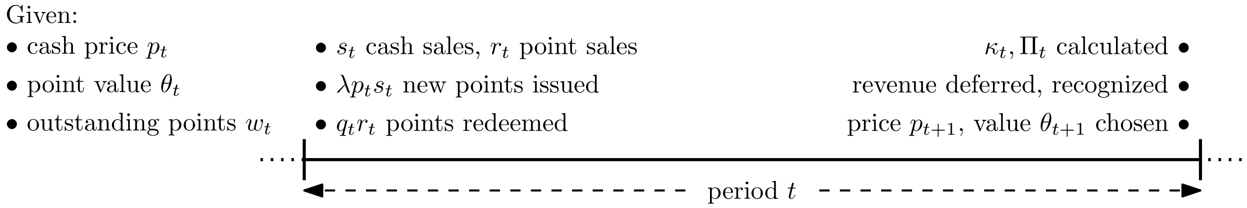

Figure 1: Timeline of events during period t.

revenues at the end of period t implicitly depend on p_{t+1} and θ_{t+1}. Consequently, this means that all these values are essentially decided at the end of period t (instead of the beginning of period $t+1$), jointly with the revenue deferral. The exact timeline of events is depicted in Figure 1.

The firm incurs *redemption servicing costs* of $c\,r_t$. Let $\kappa_t \overset{\text{def}}{=} p_t\,s_t - c\,r_t$ denote the firm's *(operating) cash flow* at the end of period t. Accordingly, the firm's (pre-tax) *profit* at the end of period t is given by

$$\Pi_t \overset{\text{def}}{=} \underbrace{p_t\,s_t + L_t - L_{t+1}}_{\text{revenues}} - \underbrace{c\,r_t}_{\text{costs}} = \kappa_t + L_t - L_{t+1}. \tag{4}$$

The manager's decision problem. The manager obtains a reward $f_t(\Pi_t)$ that is tied to the firm's profits, where f_t is a concave, increasing function.[4] The manager's problem is to select a policy for setting the cash price and point value, $\{p_t, \theta_t\}_{t=1}^{T}$, so as to maximize his cumulative, discounted expected rewards over the given time-frame, i.e., $\sum_{t=1}^{T+1} \alpha^t\,\mathbb{E}[f_t(\Pi_t)]$, where $\alpha \in (0,1]$ is a discount factor. When f_t is linear, we recover the objective of maximizing the firm's profits. Studying a strictly concave reward function f_t allows us to capture important practical managerial considerations such as taxes, income smoothing, or risk aversion, which become particularly pertinent in the context of LP management (see Section 5 for a more detailed discussion).

The model of the firm's operations and decision-making we introduced here includes only the necessary ingredients to capture the key drivers underpinning managerial decisions on loyalty point values/prices (in particular, financial accounting considerations). This will allow us to derive optimal policies and structural insights in a general-purpose and industry-independent setting. In Section 6, we show that our insights are broadly applicable, by considering more detailed operational models (e.g., selling multiple products, adjusting prices more frequently, carrying inventory, taking additional operating decisions such as inventory replenishments), as well as more general compensation schemes (depending on cash flows and profits). We discuss other modeling choices

[4]We use the terms increasing (decreasing) in their non-strict sense, i.e., to denote non-decreasing (non-increasing).

and limitations in Section 7.1, where we outline fruitful directions for future research.

3 Dynamic Programming Formulation

The manager's decision problem can be formulated as a dynamic program (DP) (Bertsekas 2001). A sufficient state is given by the number of outstanding points, the cash price, and the point value, i.e., the triple (w_t, p_t, θ_t), since s_t, r_t and g_t depend only on it. With J_t denoting the manager's value function at the beginning of period t, the Bellman recursion can now be written as:

$$
\begin{aligned}
J_t(w_t, p_t, \theta_t) &= \mathbb{E}\Big[\max_{\substack{p_{t+1}\geq 0\\ \theta_{t+1}\geq 0}} \Big(f_t(\Pi_t) + \alpha\, J_{t+1}(w_{t+1}, p_{t+1}, \theta_{t+1})\Big)\Big] && (5)\\
\Pi_t &= \kappa_t + L_t - L_{t+1} && \forall\, t \in \{1, \ldots, T\}\\
\kappa_t &= p_t\, s_t - c\, r_t, && \forall\, t \in \{1, \ldots, T+1\}\\
L_t &= w_t\, \theta_t\, g_t, && \forall\, t \in \{1, \ldots, T+1\}\\
w_{t+1} &= w_t + \lambda\, p_t\, s_t - q_t\, r_t, && \forall\, t \in \{1, \ldots, T\},
\end{aligned}
$$

where $J_{T+1}(w_{T+1}, p_{T+1}, \theta_{T+1}) = \mathbb{E}\big[f_{T+1}\big(\kappa_{T+1} + L_{T+1}\big)\big]$, i.e., all deferred revenue is recognized at the end of the terminal period.

Note that the order of the maximization and expectation operators in (5) reflects the fact that the decisions p_{t+1} and θ_{t+1} are taken at the end of period t, after observing the realized cash and point sales, s_t and r_t, respectively (see our discussion in Section 2). As stated, the problem is not readily amenable for analysis, due to the high-dimensional state, and the non-linear dynamic evolution. Fortunately, it turns out that the following mild assumptions enable tractability.

Assumption 1: "No Money Illusion." Outstanding points affect customers' purchasing and redemption behavior only through their monetary value. In particular, s_t, r_t, and g_t depend on p_t and $w_t \cdot \theta_t$. In other words, sales and redemptions do not depend separately on the total outstanding points w_t or the monetary value of a single point θ_t. To exemplify this assumption, suppose that an airline is issuing miles, with each mile having a value of \$0.01. If the airline were to exchange every 10 miles with 1 point, with each point having a value of \$0.10, then, under "No Money Illusion," sales and redemption propensities would not be affected.

This is akin to standard assumptions made in finance and economics. We align our study with this vast literature, and assume that rational decision makers should not exhibit "money illusion," i.e., that their purchasing decisions should be in terms of the real value of money, instead of the nominal one (Fisher 1928).

Assumption 2: Redemption Increasing in Value. The redemption rate is increasing in the value of outstanding points, i.e., $g_t(p_t, w_t\,\theta_t)$ is increasing in $w_t\,\theta_t$, for any fixed p_t. In the absence of integrality or point budget considerations, this assumption becomes natural, as there is no reason for customers to stockpile points.

In view of Assumptions 1 and 2, note that (3) implies that $L_t = w_t\,\theta_t\,g_t(p_t, w_t\,\theta_t)$ is a strictly increasing function in $w_t\,\theta_t$, for any fixed p_t. Given this one-to-one relation, we can equivalently take all the quantities of interest to depend only on the cash price and the value of outstanding points, i.e.,

$$s_t,\ r_t,\ g_t,\ \text{and } \kappa_t \text{ are functions only of } (p_t, L_t).$$

Several important remarks are in order. First, note that our assumptions do not imply any monotonicity of s_t in L_t. In particular, increasing the LP value L_t can lead to either an increase in cash sales s_t (the loyalty effect) or a decrease (the cannibalization effect). This is in line with empirical findings in the literature, see, e.g., Dorotic et al. (2014) and Kopalle et al. (2012).

Second, note that the point value (or point price) can be exactly inferred from the value of outstanding points, given the cash price and the number of points. More precisely, we have $\theta_t = \frac{\phi_t(p_t,L_t)}{w_t}$ (and $q_t = \frac{p_t w_t}{\phi_t(p_t,L_t)}$, respectively), where $\phi_t(p_t, L)$ is a strictly increasing function of L for any fixed p_t, $\phi_t(p_t, 0) = 0$, and $\lim_{L\to\infty} \phi_t(p_t, L) = \infty$.

Finally, for tractability purposes we assume concavity of sales revenues net of costs for any realization of the noise.

Assumption 3: κ_t is concave in (p_t, L_t). This parallels classic requirements in the literature, see, e.g., Petruzzi and Dada 1999, Federgruen and Heching 1999, and Talluri and van Ryzin 2005 for similar assumptions. For illustration purposes, Appendix B includes a discussion and examples of sales and redemption functions satisfying this requirement. We note that this assumption is not required for some of our results, particularly those in Section 4.

Returning to DP formulation (5), we can now collapse the state into a single-dimensional variable $y_t = \kappa_t + L_t$, which yields a tractable DP model, as formalized in the next result.

Theorem 1. *The manager's optimal value function can be written as*

$$J_t(w_t, p_t, \theta_t) = \mathbb{E}[V_t(y_t)],$$

where $y_t \stackrel{\text{def}}{=} p_t\,s_t(p_t, L_t) - c\,r_t(p_t, L_t) + L_t$, *for any* $t \in \{1, \dots, T+1\}$, *and the function* V_t *satisfies* *the following one-dimensional Bellman recursion*

$$V_t(y) = \max_{\substack{p_{t+1}\geq 0 \\ L_{t+1}\geq 0}} \Big[f_t\big(y - L_{t+1}\big) + \alpha\, \mathbb{E}\big[V_{t+1}(y_{t+1})\big]\Big], \tag{6}$$

where $V_{T+1}(y) = f_{T+1}(y)$. *Furthermore, the following structural properties hold:*

(i) if $\tilde{p}_{t+1}(y_t)$ *and* $\tilde{L}_{t+1}(y_t)$ *denote the optimal actions for the maximization problem in* (6)*, then the optimal pricing policies in* (5) *can be obtained as*

$$p^\star_{t+1} = \tilde{p}_{t+1}(y_t), \qquad \theta^\star_{t+1} = \frac{\phi_{t+1}(\tilde{p}_{t+1}(y_t), \tilde{L}_{t+1}(y_t))}{w_{t+1}}.$$

(ii) the function V_t *is concave and increasing.*

Interestingly, according to Theorem 1, one can equivalently think of the manager's decisions to be the cash price p_{t+1}, and the value of the LP L_{t+1} for the next period, instead of the point value θ_{t+1}. In particular, once p_{t+1} and L_{t+1} are optimally determined, the corresponding optimal θ_{t+1} can be readily derived through ϕ_{t+1} and w_{t+1}, as previously discussed. Although this requires tracking an additional (state) variable, it does not complicate the solution of the DP in (6), which retains its one-dimensional, concave structure, and can thus be efficiently solved (Bertsekas 2001).

Our new state variable y_t, given by the sum of the cash flow and the LP value just prior to the decision point, can be interpreted as the firm's *reward potential* at the end of period t. In this sense, by deciding on L_{t+1}, the manager splits the reward potential into (i) current realized profits, $y_t - L_{t+1}$, and (ii) total value of the LP for the next period, L_{t+1}. While the former quantity generates immediate rewards, the latter is "invested" in the future, impacting outcomes in a complex fashion (note that future reward potentials are affected both explicitly, as well as implicitly, through the modified cash and point sales due to the loyalty effect).

To the latter point, an alternative interpretation of L_{t+1} is actually as a loan "financed" by the firm's customers, through the promised LP rewards, and carrying a zero excess interest rate beyond the firm's internal cost of capital. Unlike regular debt, however, the firm here has the choice of when, if, and how much of the debt to repay or "write off"—by adjusting the value of LP rewards—and thus faces no associated bankruptcy risk. This is in line with the informal view that LPs act as "zero-interest loans" from consumers, held by some practitioners.

4 Linear Rewards: Profit Maximization

We first analyze the case where the rewards function f_t is linear, taken without loss to be $f_t(\Pi) = \Pi$, $\forall t \in \{1, \ldots, T+1\}$. This corresponds to the case where the manager's objective is to simply maximize (expected discounted) profits.

For a linear reward function, the optimal policy characterized in Theorem 1 simplifies considerably, as discussed in the next result.[5]

Lemma 1. *The profit-maximizing cash price p_t and point value θ_t are given by:*

$$(p_t^\star, L_t^\star) \in \arg\max_{p \geq 0, L \geq 0} \big\{ \alpha \cdot \mathbb{E}[\, \kappa_t(p, L)\,] - (1-\alpha) \cdot L \big\} \tag{7a}$$

$$\theta_t^\star = \frac{\phi_t(p_t^\star, L_t^\star)}{w_t}. \tag{7b}$$

Furthermore, for any $t \in \{1, \ldots, T\}$ and any reward potential y,

$$V_t(y) = y - L_{t+1}^\star + \sum_{\tau=t+1}^{T+1} \alpha^{\tau - t} \big[\, \mathbb{E}[\kappa_\tau(p_\tau^\star, L_\tau^\star)] + (L_\tau^\star - L_{\tau+1}^\star)\,\big], \tag{8}$$

where $L_{T+2}^\star \stackrel{\text{def}}{=} 0$.

This lemma affords several interesting implications.

• **Tradeoffs.** The objective in (7a) reveals the tradeoffs that the managerial decisions, the LP value and the cash price, control: 1) The LP value trades off potential future benefits of an increased loyalty effect with the (time-value) loss incurred by deferring more revenue. In other words, by increasing L, the next-period cash flow κ_t may rise, but the deferred revenue incurs a time-value loss. 2) The cash price trades off sales revenues with costs related to servicing redemptions.

• **Base-stock, list price policy.** From a structural standpoint, Lemma 1 suggests an intriguing connection between the optimal policy for setting the LP value and well-known policies in operations management for replenishing and pricing inventory. In the latter context, such policies, known as "base-stock, list price," are governed by two (time-dependent) values $(\tilde{S}_t, \tilde{p}_t)$; when inventory level falls below the *base-stock* $\tilde{S}_t$, an order is placed to raise the inventory to this target, and the *list price* $\tilde{p}_t$ is charged. Otherwise, no order is placed, and the price charged is lower than $\tilde{p}_t$, and decreasing in the initial inventory level (see, e.g., Petruzzi and Dada 1999, Federgruen and Heching 1999, Elmaghraby and Keskinocak 2003, Chen and Simchi-Levi 2004, and references therein). In

[5] The concavity of κ_t is no longer required for this result.

our setting, the LP value—or, equivalently, the pool of deferred revenue—plays the role of inventory, and the values $L_t^\star$ and $p_t^\star$ can be interpreted as the "base-stock, list price" targets, respectively. From a practical standpoint, base-stock policies have multiple advantages: they are widely used for inventory management, and familiar to managers; they are simple, and easy to communicate; they rely on few parameters, and can thus be easily calibrated and are likely robust.

Further exploring the similarities between LP and inventory management, note that in our context maintaining an excessively valuable LP or large pool of deferred revenue (i.e., holding "too much inventory") results in increased costs due to redemption and time-value loss—which are the corresponding overage/holding costs. On the other hand, maintaining a less valuable LP or small pool of deferred revenue (i.e., holding "too little inventory") results in a weak loyalty effect and missed sales opportunity costs—which are the corresponding underage/backlogging costs.

However, an important distinction between our model and classical inventory management lies in the manager's ability to make *downwards* adjustments in the "inventory" of deferred revenue—as such, the desired base-stock $L_t^\star$ is always attainable here, and consequently the list price $p_t^\star$ is always charged. In the next section, we further explore (dis)similarities between LP and inventory management, highlighting how tax considerations, income smoothing, and/or risk aversion in the former setting produce a similar effect to convex ordering costs in the latter. As we will show, the resulting optimal policies for managing the LP preserve their base-stock structure, albeit the base-stock targets become state-dependent.

- **LP value independent of current operating performance.** Lemma 1 shows that the optimal LP value is chosen independently of the current cash flows and the reward potential y. Note that the optimal LP value is also independent of the variability in future sales or redemptions. As our discussion in Section 5 will highlight, these insights critically rely on the linear/profit-maximizing objective, and no longer hold when the reward function is concave.

- **Zero-excess interest rate loan.** According to (7a), a profit-maximizing manager faces a series of equivalent one-period LP design problems. In every such problem, the LP is set up from scratch, by investing an amount L, which drives the expected cash flow obtained during the period, $\mathbb{E}[\,\kappa_t(p, L)\,]$. At the end of the period, the LP is dissolved, and the original investment of L is fully recovered. Properly discounting the payoffs, the value extracted in every such design problem equals $-L + \alpha(\mathbb{E}[\,\kappa_t(p, L)\,] + L)$, which is precisely the objective in (7a). This intuition shows how L acts as a loan "financed" by the firm's customers through the promised LP rewards, carrying a zero excess interest rate beyond the firm's internal time discounting.

- **Firm/LP valuation.** Finally, equation (8) implies that the value function V_t decomposes into the sum of the firm's future discounted profits, yielding a valuation measure of the firm. Individual profit terms have two components, corresponding to the expected cash flow and adjustments in the deferred revenue, respectively. This formula could also be used to elicit the value added to a firm by running an LP. In particular, the manager of a firm without an LP would charge prices p_k that maximize sales revenues, and project the profits as $y + \sum_{k=t+1}^{T+1} \alpha^{k-t}\mathbb{E}[\kappa_k(p_k, 0)]$. Thus, the difference between $V_t(y)$ and this value yields a fair quantification of the LP-added value.

Next, we present several properties pertaining to the optimal cash price and LP value.

Corollary 1 (No cannibalization). *When the reward function is linear,*
(i) the LP value is never inflated so as to cannibalize cash flow, i.e., $\frac{\partial \mathbb{E}[\kappa_t]}{\partial L}|_{(p_t^\star, L_t^\star)} \geq 0$;
(ii) if expected redemptions are increasing in the cash price (LP value), then the optimal cash price (LP value) is set so as to never cannibalize sales revenues, i.e., $\frac{\partial \mathbb{E}[r_t]}{\partial p_t} \geq 0 \Rightarrow \frac{\partial \mathbb{E}[p_t s_t]}{\partial p_t}|_{(p_t^\star, L_t^\star)} \geq 0$ $\left(\frac{\partial \mathbb{E}[r_t]}{\partial L_t} \geq 0 \Rightarrow \frac{\partial \mathbb{E}[p_t s_t]}{\partial L_t}|_{(p_t^\star, L_t^\star)} \geq 0\right)$.

The fact that cash flow and sales revenue are increasing at optimality in L_t suggests that a profit-maximizing manager would never grow the LP to the extent that it would cannibalize the regular revenue source. This reinforces the intuition that the optimal L_t involves a strict trade-off between the two competing objectives, of maximizing time-value of money and leveraging loyalty to increase sales income. Thus, cannibalization can never happen here, as it would lead to both effects negatively impacting the manager's objective. Interestingly, this intuition breaks down when the manager's objective is concave, as we discuss in Section 5.

Furthermore, since sales revenues are increasing in p_t at optimality, this suggests that the manager should charge lower prices than the ones that would maximize sales revenues (i.e., without considering redemption costs). This is driven by the fact that higher prices also induce more redemptions, and thus larger servicing costs. As a side note, the requirements in the corollary are mild, as it is natural to expect increasing redemptions when cash prices increase (as more customers substitute the "expensive" cash purchase with points) or when the LP value increases.

4.1 Comparative Analysis

We next present several comparative results describing how the profit-maximizing decisions depend on the manager's discount factor α, the redemption cost c, and the presence of the LP.

Lemma 2a. *The value of the loyalty program is increasing with the discount factor* α*.*

This result is driven by the decreased opportunity cost of the (inaccessible) deferred revenue associated with a higher α. In view of our interpretation of deferred revenue as inventory, a lower time-value of money (i.e., higher α) corresponds to lower holding costs, leading to higher "inventory" levels. Viewed under a financial lens, the result suggests that firms facing lower cost of capital (i.e., higher α) will tend to operate under higher leverage, by increasing the LP-related liabilities.

Lemma 2b. *If the expected redemptions are increasing in the cash price and the LP value (i.e., $\frac{\partial \mathbb{E}[r_t]}{\partial p_t} \geq 0, \frac{\partial \mathbb{E}[r_t]}{\partial L_t} \geq 0$), then the optimal cash price and the LP value are decreasing in the per-unit redemption cost, i.e., $\frac{\partial p_t^\star}{\partial c} \leq 0$, and $\frac{\partial L_t^\star}{\partial c} \leq 0$. If, additionally, the redemption rate g_t is also independent of the cash price p_t, then, ceteris paribus, the point value is also decreasing in the per-unit redemption cost, i.e., $\frac{\partial \theta_t^\star}{\partial c} \leq 0$.*

Facing increased redemption servicing costs, the manager devalues the LP, and at the same time charges lower cash prices. While the devaluation seems to be an intuitive response to increased redemption costs—decreasing the points' value averts redemptions—lowering cash prices appears counterintuitive at first: why would a firm decrease prices under increased costs? This is because customers would prefer cash purchases under lower prices, which would reduce costly redemptions. More broadly, this suggests that by making redemption procedures more efficient, firms would not only benefit from cost savings, but also from their ability to command higher cash prices.

Lemma 2c. *If the expected cash flow $\mathbb{E}[\kappa_t(p_t, L_t)]$ is supermodular (submodular) in (p_t, L_t), then the optimal price charged by a firm running an LP is larger (smaller) than the price charged by a firm without an LP, i.e., $p_t^\star(L_t^\star) \geq (\leq) p_t^\star(0)$.*

This result elicits a condition that could help explain whether managers may charge lower of higher cash prices if they operate an LP, depending on whether the two have complementary effects on the cash flow, i.e., on whether $\mathbb{E}[\kappa_t]$ is supermodular. For instance, in contexts where the loyalty effect does not decrease the customers' willingness to pay, $\mathbb{E}[\kappa_t]$ is likely to be supermodular. Several recent empirical papers confirm this to be the case in the travel and hospitality industry (see, e.g., Mathies and Gudergan 2012, McCaughey and Behrens 2011 and Brunger 2013), so that here one might expect higher cash prices under more valuable loyalty programs. In contrast, when LPs attract a larger population of customers that are also more price-sensitive, $\mathbb{E}[\kappa_t]$ is likely to be submodular, so the presence of (more valuable) loyalty programs would warrant lower cash prices.

5 Concave Rewards

We now analyze the case where the reward function f_t is a (non-linear) concave function. Concavity is routinely employed in the literature to reflect the way managerial decision making is affected by a multitude of important economic factors and behavioral incentives, which we review next, without compromising analytical tractability.

(a) *Tax management.* Post-tax profit can be expressed as a concave function of pre-tax profit (Smith and Stulz 1985). While the effect of taxation is often ignored in the operations literature—based on the assumption that maximizing pre-tax and post-tax profit is equivalent—such a simplification can be problematic in our setting. In fact, the value of the loyalty points has a significant and subtle effect on taxation that needs to be accounted for: according to U.S. Income Tax Law, the taxable year of inclusion of LP-related deferred revenue could depend on when the revenue is in fact recognized, e.g., due to redemption (Ernst&Young 2014). Thus, taxable income at time t is influenced by the newly deferred/recognized revenue, making the post-tax profit a concave function of the profit Π_t.

(b) *Income smoothing.* It is well established empirically that managers of large firms are averse to fluctuations in income, and thus employ practices that result in their smoothing (see, e.g., DeFond and Park 1997, Healy and Wahlen 1999, Kasznik 1999, Leuz et al. 2003, and references therein). The vast empirical evidence is also complemented by several theoretical models rationalizing such behavior (see, e.g., Trueman and Titman 1988, Fudenberg and Tirole 1995, Beyer 2009). A concave reward function adequately captures such incentives: for low profits, the marginal reward is high, whereas for high profits, it is low (Lambert 1984). Changes in the loyalty point value influence profits, and thus facilitate income smoothing: for instance, devaluing the LP can lead to increased recognized revenue, helping to boost current profit. It is very important to note at this point that, unlike earnings manipulation, the process we describe here involves real business decisions, which could also affect future operational performance. This practice is also referred to as real earnings management in the literature (see, e.g. Graham et al. 2005, Cohen et al. 2008). For instance, devaluation of the LP in order to boost current earnings might compromise next period's cash sales, as customers would be offered less valuable reward points.

(c) *Risk aversion.* Managers are often averse to risks (see, e.g., Pratt 1964, Smith and Stulz 1985, etc). Concave utility functions have been widely used to capture such effects, for instance through the hyperbolic absolute risk aversion (HARA) function.

With this motivation, we first characterize the manager's optimal policies, and then investigate the impact of uncertainty, degree of concavity in the reward function, and time. To derive analytical results, we also make the following simplifying assumption concerning the randomness, effective throughout the remainder of this section.

Assumption 4: Additive Noise. The cash flow is affected by additive noise, $\kappa_t(p_t, L_t) =$

$\bar{\kappa}_t(p_t, L_t) + \sigma\varepsilon_t$, where $\bar{\kappa}_t(p_t, L_t) \stackrel{\text{def}}{=} \mathbb{E}[\kappa_t(p_t, L_t)]$, ε_t are independent across time, with zero mean and unit variance, and $\sigma \geq 0$. While many of our results continue to hold under more general noise models, we adopt this parameterization to streamline the analysis. We note that it has been used in several papers in the literature (see, e.g., Federgruen and Heching 1999, Chen and Simchi-Levi 2004), and it provides a simple and intuitive way of quantifying variability, through a single parameter—the standard deviation, σ—with larger σ corresponding to increased variability/uncertainty.

Recall that the manager's optimal value function satisfies the Bellman recursion in equation (6), reproduced below for convenience:

$$V_{t-1}(y) = \max_{\substack{p_t \geq 0 \\ L_t \geq 0}} \Big[f_{t-1}\big(y - L_t\big) + \alpha\, \mathbb{E}\big[V_t(y_t)\big] \Big],$$

where $y_t = \kappa_t(p_t, L_t) + L_t$ is the firm's reward potential at the end of period t. This problem is generally not decomposable across time, unlike the case discussed in Section 4. As a consequence, the optimal policy no longer admits a closed-form characterization, but is rather obtained by solving the (one-dimensional) DP problem. The next result derives several structural properties of the optimal actions.

Theorem 2. *Under a concave reward function,*[6] *for any time $t \in \{2, \ldots, T+1\}$,*
(a) the optimal LP value increases in the reward potential, i.e., $L_t^\star(y)$ is increasing in y;
(b) the firm's profit increases in the reward potential, i.e., $\Pi_t^\star = y - L_t^\star(y)$ is increasing in y;
(c) the optimal cash price and point value satisfy

$$p_t^\star(y) \in \arg\max_{p \geq 0} \bar{\kappa}_t(p, L_t^\star(y)), \qquad \theta_t^\star(y) = \frac{\phi_t(p_t^\star(y), L_t^\star(y))}{w_t}.$$

Theorem 2 derives a novel insight: the LP can be used to smooth a firm's financial performance, acting as a revenue buffer against poor performance. To illustrate this, suppose that, ceteris paribus, the firm's operating cash flows κ_t increase (decrease), e.g., due to stronger (weaker) sales. Consequently, the reward potential y would increase (decrease), and according to Theorem 2(a), the manager would increase (decrease) the LP value. Effectively then, when facing strong operating performance, the manager would defer a larger part of the revenue, for future access; when fac-

[6] Part (a) of the theorem holds independent of Assumption 4, and part (b) holds provided that the condition $\bar{\kappa}(p, L) \geq \bar{\kappa}(p', L')$ implies that $\kappa(p, L)$ dominates $\kappa(p', L')$ in the sense of second-order stochastic dominance, for any pair of values (p, L) and (p', L'). This holds for the additive noise model of Assumption 4, but also for a multiplicative noise model, i.e., $\kappa_t(p_t, L_t) = \bar{\kappa}_t(p_t, L_t) \cdot \varepsilon_t$, or a combination thereof.

ing mediocre performance, the manager would boost current profits by recognizing some deferred revenue. This smoothing function provides a new rationale for the existence of an LP, even in the absence of competition for the firm. Part (b) of the result implies that the manager would nonetheless never engage in "excessive smoothing," and would ensure that an increased (decreased) reward potential always results in increased (decreased) profit for the firm.

Our findings also have interesting implications for the firm's customers, suggesting that they always "share the pain and the gain" with the firm. More precisely, increased operating performance always induces the manager to promise more value for loyal customers in the future, through an inflated LP, while at the same time recording larger immediate profits for the firm (and larger rewards for himself). On the flip-side, poor performance leads the manager to decrease the promise for future value to the customers, as well as the firm's (and his own) immediate benefit.

In comparison with Lemma 1, our discussion above also illustrates that when the rewards function is concave, setting the optimal LP value depends on current operating performance and variability in future sales or redemptions. We explore the latter dependence in more detail in our Comparative Analysis subsection below.

Finally, the structure of the manager's optimal actions allows us to reinforce the connection between LP decisions and inventory management, better highlighting the impact of concavity. In particular, Theorem 2 suggests that in choosing the optimal LP value (or the "inventory" of deferred revenue), the manager effectively follows a *state-dependent* base-stock policy, with target level $L^\star(y)$. Thus, concavity of the reward function is indeed equivalent in effect to convex ordering/production costs in classical inventory theory: just as the latter lead to a smoothing of orders/production across time, giving rise to state-dependent base-stock policies (Porteus 2002), the former induces the manager to smoothen profits, and to set state-dependent LP values.

The next result further outlines the differences of the manager's policy under linear and concave reward functions.

Corollary 2. *At optimality, the LP value can be inflated so as to cannibalize cash flow,*

$$\left.\frac{\partial \bar{\kappa}_t}{\partial L}\right|_{(p_t^\star(y), L_t^\star(y))} > -1, \quad \textit{for all } t = 2, \ldots, T+1 \textit{ and } y. \tag{9}$$

The manager might inflate the LP to the extent that it cannibalizes expected cash flow, in stark contrast to the optimal behavior exhibited in Section 4 (see Corollary 1). Since κ_t is concave, cannibalization will occur when $L_t^\star(y)$ is relatively high, which in turn translates to y being high in period $t-1$, based on our previous discussion. This suggests that managers faced with very strong performance in a period could promise such high value to their customers through the LP that marginal redemption costs could start outweighing marginal sales revenue in the future. Furthermore, and more interestingly, this suggests a potentially counterintuitive behavior: in anticipation of a weak period (e.g., due to weak sales revenues or increased redemption costs), the manager may in fact *exacerbate* the outcomes, by increasing the LP value, thus fueling more redemptions and possibly further lowering sales. It is worth emphasizing that such outcomes are possible here (and not in Section 4) due to the concavity of f_t, as the marginal impact of a future reward increase can outweigh both time-value loss and sales cannibalization.

Interpreted in a different way, inequality (9) also implies an upper bound on the value of the LP (since $\bar{\kappa}_t$ is concave, its derivative is decreasing, so that a lower bound on this derivative is equivalent to an upper bound on $L_t^\star$). This result adds a different perspective to the intuition from Theorem 2 (a-b), suggesting that managers would eventually cap the increase in LP value as they face increasingly strong operational performance.

5.1 Comparative Analysis

Paralleling our results in Lemma 2b-2c, it can be confirmed that the optimal cash price decreases in the per-unit redemption cost and increases (decreases) in the LP value under super(sub)-modularity of $\bar{\kappa}_t$, and the value of the LP is increasing in the manager's discount factor.

Our next results examine the impact of variabity, degree of concavity of f_t, and time. In order to distill the impact of these factors, it will be useful to consider instances where all problem primitives are stationary in our analysis, i.e., $s_t = s$, $r_t = r$, etc.

Impact of Variability

As our results so far have demonstrated, managers with concave rewards would utilize LPs and their associated "inventory" of deferred revenue as a means of protection against future fluctuations in operating performance. In this sense, *variability* in outcomes may play a critical role, unlike in Section 4 (recall that only *expected* cash flows actually mattered in the decision process of a profit-maximizing manager).

To highlight the dependency on variability, let $V_t(y, \sigma)$ be the value function defined as in (6) when the standard deviation of cash flows is σ, and similarly denote the optimal policies by $L_t^\star(y, \sigma)$ and $p_t^\star(y, \sigma)$.

Theorem 3. *Suppose that the model primitives are stationary. Let* $\rho(L) \stackrel{\text{def}}{=} \max_p \bar{\kappa}(p, L)$*. If* u' *and* ρ' *are convex, then for all* $t = 1, \ldots, T+1$ *and for all* y*, we have that* *(a) the value function $V_t(y, \sigma)$ is decreasing in σ,*

(b) the optimal LP value $L_t^\star(y, \sigma)$ is increasing in σ, and

(c) ceteris paribus, if $\bar{\kappa}$ is supermodular (submodular) in (p, L), the optimal cash price $p_t^\star(y, \sigma)$ is increasing (decreasing) in σ.

We first comment on the conditions in Theorem 3. Convexity of the derivative of the reward function is a reasonable assumption, and is satisfied by the vast majority of commonly used utility functions, including the entire broad family of HARA utilities. Convexity of ρ' is a technical condition, introduced for tractability purposes. Note that it is readily satisfied for the important class of quadratic revenue models, i.e., under linear price impact (see our discussion in Appendix B).

Part (a) suggests that a manager derives less rewards under increased variability—an intuitive consequence of Jensen's inequality. Surprisingly, part (b) argues that increased variability would also cause a manager to actually *increase* the value of the LP, and thus the future promised rewards to the customers. Such action may seem counterintuitive at first, particularly when recognizing that it is tantamount to an increased liability on the firm's balance sheet. What sheds light on this outcome is the interpretation of the LP as "safety stock" held in anticipation of future fluctuations in performance, with a larger stock being preferable under increased uncertainty.

In this sense, the impact of uncertainty on the manager's policy differs slightly from classical inventory management, where increased uncertainty could have subtle effects on the optimal inventory level, leading to either increased or decreased values, depending on the overage and underage costs. For the management of loyalty point value, however, the relationship between the virtual "inventory" of deferred revenue and variability is monotonic.

Part (c) confirms that increased uncertainty may lead to either decreased or increased prices, depending on their complementarity with the LP value in the expected cash flow $\bar{\kappa}$. For example, based on our discussion of Lemma 2c, this suggests that managers should charge higher (lower) prices in contexts where loyalty effects increase (decrease) customers' willingness to pay.

Impact of Degree of Concavity

In a similar fashion to our prior analysis, we parameterize the reward function in order to quantify the degree of concavity of f_t. In particular, we consider functions of the form

$$f_t(\Pi) = \begin{cases} \gamma \cdot (\Pi - \hat{\Pi}), & \Pi \le \hat{\Pi}, \\ \Pi - \hat{\Pi}, & \Pi > \hat{\Pi}, \end{cases} \quad \text{for all } t = 1, \ldots, T, \tag{10}$$

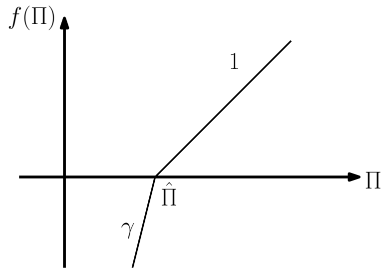

Figure 2: The reward function considered in Section 5.

where $\gamma \geq 1$ (see Figure 2). For $\gamma = 1$, we recover the case of linear rewards. As the value of γ increases, the effects of concavity of f_t become more pronounced. Note that piece-wise linear rewards (utilities) of this type have been studied in the literature (see, e.g., Ben-Tal and Teboulle 2007). Although other parameterizations are possible, our choice yields tractability, and remains suitable for capturing important effects such as taxation (e.g., with γ being a lower taxation rate for profits below a threshold $\hat{\Pi}$), earnings smoothing or risk aversion (e.g., with γ being the manager's aversion for shortfalls with respect to a pre-set benchmark/target $\hat{\Pi}$).

In keeping with the notation used above, let $L_t^\star(y, \gamma)$ be the optimal value of the LP when the reward function is of the form in (10).

Theorem 4. *Suppose that the model primitives are stationary. For all $t = 2, \ldots, T+1$, there is a threshold $\hat{y}_t$ such that*

(a) the optimal LP value $L_t^\star(y, \gamma)$ is increasing in γ if $y > \hat{y}_t$, and decreasing in γ if $y \leq \hat{y}_t$;

(b) $\hat{y}_t$ is decreasing in t;

(c) $\hat{y}_t$ is increasing in γ.

While intuition might suggest that increasing risk aversion would lead to maintaining a smaller pool of deferred revenue, i.e., a less valuable LP, Theorem 4 shows that this is not true when the firm's current reward potential y is higher than a threshold $\hat{y}_t$. In fact, in that case increasing risk aversion would lead to *inflating* the LP. To understand the effect, note that $\hat{y}_t$ can be thought of as an *adjusted* target that the manager himself sets. Reward potentials above this milestone are considered "gains," whereas below are considered "losses." Increased loss aversion would thus result in more deferred revenue and an inflated LP if the firm faces "gains," so as to hedge against future losses. On the contrary, if the firm faces "losses," increased loss aversion would result in less deferred revenue and a deflated LP, so as to mitigate the present losses. Finally, the threshold $\hat{y}_t$ is decreasing in t and increasing in γ, i.e., managers would set higher targets early in their planning horizon, and as they become more loss averse.

The non-monotonic relationship between the "inventory" of deferred revenue and loss aversion is in fact in line with findings in classical inventory management. In particular, Chen et al. (2007) documents a similar non-monotic relationship between the order-up-to-level and risk aversion in a multiperiod inventory control problem.

Impact of Time

We conclude by examining the dependence of the impact of time on the LP value, as well as on the manager's marginal valuation of income.

Lemma 3. *If the problem primitives are stationary, then under the optimal policy, for all* y,
(a) the optimal value of the LP is decreasing in time, i.e., $L_t^\star(y)$ *is decreasing in* t*;*
(b) the marginal value of reward potential is decreasing in time, i.e., $V_t'(y)$ *is decreasing in* t*.*

Part (a) suggests that managers would tend to prefer more valuable LPs earlier in the planning horizon, and as such would tend to inflate them early on. The intuition behind this is that, apart from boosting sales due to the loyalty effect, a valuable LP also results in a buffer of deferred revenues that could be used to hedge against future fluctuation in operating performance in subsequent periods. Both capacities, however, diminish as less time steps remain in the planning horizon.

Finally, in view of our interpretation of deferred revenue as virtual "inventory," part (b) parallels classical results in operations management, which maintain that the marginal value of a unit of inventory decreases over time (see, e.g. Talluri and van Ryzin 2005).

6 Extensions

We now explore several extensions of our framework. In particular, we consider firms whose operations are more complex than selling a single product, update their prices more frequently, or employ managers whose rewards are tied to both cash flows and profits. For each setting, we confirm that our main results and insights remain unchanged.

6.1 More Complex Operating Model

Our model so far focused on a firm selling a single product with perishable inventory, and endowed with two decisions, the cash price and point value. To generalize this setting, consider first a firm that is providing multiple products or services to its customers, *without* running an LP. At the beginning of period t, the firm's state is given by a vector $x_t \in \mathbb{R}^n$, and the firm's manager takes a set of constrained actions $a_t \in \mathcal{A}(x_t) \subseteq \mathbb{R}^m$ corresponding to operating decisions. The firm's operations during period t generate total sales of s_t products, a cash flow of κ_t (also equal to the firm's profit Π_t), and causing the firm's state to transition to x_{t+1}. All quantities s_t, x_{t+1}, and κ_t depend on the initial operating state x_t, on the firm's actions a_t, and on an exogenous random vector ε_t. The firm's manager obtains a reward f_t tied to the firm's profit during the period, and seeks an operating policy $\{a_t\}_{t=1}^T$ that maximizes his total discounted rewards, i.e., $\sum_{t=1}^{T+1} \alpha^t \mathbb{E}[f_t(\Pi_t)]$.

To introduce the LP, assume the firm now rewards customers with points for their cash purchases, and allows point redemptions for its products. Let w_t denote the outstanding points at the beginning of period t. As in our base model, points do not expire, and the firm's only LP-related decision is the point value during period t, denoted by $\theta_t \in \mathbb{R}^+$, which induces a set of corresponding point requirements q_t for the products (equal to $\frac{p_t}{\theta_t}$ when cash prices are p_t). During period t, the firm now observes cash sales of s_t and point sales of r_t, and correspondingly issues Λ_t new points and retracts $r_t^\top q_t$ points, so that $w_{t+1} = w_t + \Lambda_t - r_t^\top q_t$. As a result of the sales, the firm's operating state transitions to x_{t+1}, and the firm records a cash flow of κ_t and an operating profit of $\kappa_t + L_t - L_{t+1}$, where $L_t = w_t \, \theta_t \, g_t$ is the total deferred revenue associated with the LP, calculated under an estimated redemption rate g_t. All quantities $x_{t+1}, s_t, r_t, \Lambda_t$, and g_t now depend on the operational state x_t, on the operating decisions a_t, and on the outstanding points w_t and the monetary value θ_t; excepting g_t, all are also affected by exogenous randomness ε_t. As before, the firm's manager seeks a policy for setting the operating decisions and monetary point values $\{a_t, \theta_t\}_{t=1}^T$ that maximizes his cumulative, discounted rewards.

It is worth noting that our base model is a special case of this more general framework, with $x_t = \emptyset$, $a_t = p_t$, $\Lambda_t = \lambda p_t s_t$, and $\kappa_t = p_t s_t - c r_t$. This framework allows capturing more complex dynamics, such as a retailer/manufacturer deciding replenishment/production quantities and selling prices, or an airline/hotel adjusting booking limits to manage capacity.

As in our analysis in Section 3, the manager's value function at the beginning of period t, J_t, can be obtained as the solution to the following Bellman recursion:

$$J_t(x_t, w_t, a_t, \theta_t) = \mathbb{E}_{\varepsilon_t}\left[\max_{\substack{a_{t+1} \in \mathcal{A}(x_t) \\ \theta_{t+1} \geq 0}} \Big(f_t\big(\kappa_t + L_t - L_{t+1}\big) + \alpha \, J_{t+1}(x_{t+1}, w_{t+1}, a_{t+1}, \theta_{t+1}) \Big) \right], \tag{11}$$

where J_{T+1} corresponds to a suitable terminal reward, and $\kappa_t, L_t, L_{t+1}, x_{t+1}, w_t$ exhibit appropriate dependencies on x_t, w_t, a_t, θ_t, and the exogenous noise. The presence of the additional state variables related to the LP and the nonlinear dependency of w_{t+1} complicates recursion (11), even if the underlying recursion for a firm with no LP (i.e., with $w_t \equiv \theta_t \equiv 0, \forall\, t$) is tractable.

In this context, under our earlier assumptions that (1) customers do not suffer from money illusion, and (2) the redemption rate g_t is (weakly) increasing in the face value of points $w_t\theta_t$, we can run through the same argument as in Section 3 to conclude that s_t, r_t, κ_t and x_{t+1} only depend on $(x_t, a_t, L_t, \varepsilon_t)$, and given x_t, a_t, and w_t, there is a one-to-one mapping between the face value of a point θ_t and the total LP value L_t. More precisely, we have $x_{t+1} = X_t(x_t, a_t, L_t, \varepsilon_t)$, $\kappa_t = K_t(x_t, a_t, L_t, \varepsilon_t)$, and $\theta_t = \frac{\phi_t(x_t,a_t,L_t)}{w_t}$, for some functions X_t, K_t, ϕ_t, where $\phi_t(x, a, \cdot)$ is strictly increasing for any x, a.

We can now state several results paralleling our earlier findings, but in this more general setting; the formal proofs follow similar arguments, and are omitted.

Proposition 1. *Under the more general model of the firm,*

i.) the manager's optimal value function is given by $J_t(x_t, w_t, a_t, \theta_t) = \mathbb{E}_{\varepsilon_t}[V_t(x_{t+1}, y_t)]$, *where* $y_t \stackrel{\text{def}}{=} \kappa_t + L_t$ *denotes the firm's reward potential during period* t, *and* V_t *is given by:*

$$V_t(x_{t+1}, y_t) = \max_{\substack{a_{t+1}\in\mathcal{A}(x_{t+1})\\ L_{t+1}\geq 0}} \Big[f_t\big(y_t - L_{t+1}\big) + \alpha\, \mathbb{E}_{\varepsilon_{t+1}}\big[V_{t+1}(x_{t+2}, y_{t+1})\big]\Big]. \tag{12}$$

Furthermore, optimal policies in (11) can be readily obtained from the optimal policies in (12), by keeping track of the balance of outstanding points w_t.

ii.) if the reward functions are linear $f_t(\Pi) = \Pi, \forall t$, *then* $V_t(x, y) = y + H_{t+1}(x)$, *where* $H_t(x)$ *and the optimal LP values* $L_t^\star(x)$ *and operating decisions* $a_t^\star(x)$ *are given by:*

$$H_t(x) = \max_{\substack{L\geq 0\\ a\in\mathcal{A}(x)}} \Big\{ -(1-\alpha)\cdot L + \alpha\cdot \mathbb{E}_{\varepsilon_t}\Big[K_{t+1}(x, a, L, \varepsilon_t) + H_{t+1}\big(X_{t+1}(x, a, L, \varepsilon_t)\big)\Big]\Big\}$$

$$\theta_t^\star(x) = \frac{\phi_t\big(x, a_t^\star(x), L_t^\star(x)\big)}{w_t}.$$

iii.) if the reward functions are concave, then the optimal LP value increases in the reward potential at any given firm state, i.e., $L_t^\star(y, x)$ *is increasing in* y, *for any* x.

Proposition 1 confirms that our main insights are quite robust, and persist under this more general model of the firm. Part *i.*) parallels Theorem 1, and reinforces our interpretation of L_{t+1} as an "investment" decision that splits the firm's reward potential y_t between realized profit during the present period, $y_t - L_{t+1}$, and LP value invested in the firm's future cash flows and profits.

Part *ii.*) parallels Lemma 1, and illustrates that all its implications carry over. Specifically, when the manager maximizes the firm's discounted profits, the value of the LP during a subsequent period is chosen independently of the present financial performance. As before, L_{t+1} trades off the immediate time-value loss incurred by larger deferrals with the potential future benefits of a larger LP value, and resembles a zero interest rate loan financed by the firm's customers. Structurally, the manager's optimal policy retains a "base-stock, list-price" flavor, with the target "base-stock" of LP value chosen first, and the operating decisions a_t chosen so as to maximize the firm's discounted cash flows, subject to the optimally-chosen (contingent) LP valuations. It is important to note that, although independent of financial performance, the optimal LP value $L^\star_{t+1}$ does depend on the firm's operating state x_{t+1}.

Finally, part *iii.*) mirrors Theorem 2(a), and reveals that when the manager's reward is concave, the LP acts as a buffer against uncertainty and a tool for smoothing the firm's performance. More precisely, the future LP value is influenced by the firm's current financial performance (i.e., $L^\star_{t+1}$ depends on y_t), and the manager always sets the LP target value so as to increase (decrease) the value of points whenever performance is better (worse). As before, optimal policies ensure that the firm's customers "share the pain and the gain" with the firm and its manager.

We note that the results above hold under no additional structural requirements on the primitives. In particular, we did not require concavity of cash flows. Under suitable restrictions, additional results are possible. For instance, if the functions X_t and K_t (describing the firm's state evolution and cash flows, respectively) are jointly concave, then one can readily check that V_t remains jointly concave in (x_{t+1}, y_t). Thus, although accounting for the LP introduces an additional state variable, it preserves the structural properties of the manager's optimization problem. Under concavity of X_t and K_t, one can also check that $y - L^\star_t(y)$ is increasing in y (for any fixed operating state x), mirroring the results in Theorem 2(b).

6.2 Frequent Updating of Prices

The firm in our base model could adjust its cash price p_t only at the beginning/end of a period (a financial quarter). To capture more frequent updates, suppose each "macro-period" t in our model is split into several "micro-periods" (t, i), $i \in \{1, \dots, N\}$, and the firm can change the price $p_{t,i}$ in each micro-period. In this case, we can think of p_t as a *target* price, which the firm chooses at the end of period $t-1$; the firm's subsequent (micro) pricing decisions $p_{t,i}$ would then have to be consistent with this target, i.e., they have to be equal on average.[7] Provided that the expected cash flow achieved during period t—when maximizing over price $\boldsymbol{p}_t \stackrel{\text{def}}{=} (p_{t,1}, \dots, p_{t,N})$ that are *consistent* with the target p_t—remains jointly concave in (p_t, L_t), our results will carry through. For instance, this would

[7]Note that, if the micro-prices were *not* equal on average with p_t, the firm's implemented prices would consistently bear no resemblance to the ones used in calculating the firm's reported profits, raising serious issues about fraudulent accounting and operating practices.

be the case if the expected cash flows achieved in every micro-period (t, i) were concave in the firm's decisions in that period. To see this, note that $\mathbb{E}[\kappa_t(p_t, L_t)] \stackrel{\text{def}}{=} \max_{\boldsymbol{p}:\boldsymbol{e}^\top \boldsymbol{p}/N = p_t} \sum_{i=1}^N \mathbb{E}[\kappa_{t,i}(p_{t,i}, L_t)]$ remains jointly concave in (p_t, L_t) if $\mathbb{E}[\kappa_{t,i}]$ are jointly concave, so that all our results carry over.

A same line of arguments could be employed to address more frequent updates of the point value θ_t, as well.

6.3 Rewards Tied to Profits and Cash Flows

Although our main treatment considered rewards tied to the firm's profits, in practice cash flows could also be relevant. Both profits and cash flows are fundamental measures of firm performance, widely employed in debt covenants, in the prospectuses of firms seeking to go public, and by investors and creditors (see Dechow 1994). Furthermore, ample empirical evidence suggests that profits and cash flows critically drive managerial decisions, as they are used in compensation plans (see, for instance, Fox 1980, Healy 1985, and Ittner et al. 1997).

We now assume that the manager's reward is $f_t(x_t)$, where $x_t \stackrel{\text{def}}{=} \xi \cdot \Pi_t + (1-\xi) \cdot \kappa_t$ for some $\xi \in [0, 1]$, retaining all other assumptions in our model. Such a convex combination of profits and cash flows could correspond to the typical weights used in compensation plans (see, e.g., Delta Airlines (2014)). The following result extends our main findings to this more general setting. The proof is omitted.

Lemma 4. *When the manager's rewards depend on x_t,*

i.) The manager's optimal value function at the beginning of period t can be written as $\mathbb{E}[V_t(y_t)]$, where $y_t \stackrel{\text{def}}{=} p_t\, s_t(p_t, L_t) - c\, r_t(p_t, L_t) + \xi \cdot L_t$ is the reward potential, and V_t satisfies

$$V_t(y) = \max_{\substack{p_{t+1} \geq 0 \\ L_{t+1} \geq 0}} \Big[f_t\big(y - \xi \cdot L_{t+1}\big) + \alpha\, \mathbb{E}\big[V_{t+1}(y_{t+1})\big] \Big], \tag{13}$$

where $V_{T+1}(y) = f_{T+1}(y)$. Furthermore, V_t is concave, and the manager's optimal policy can be directly obtained from the optimal actions in the maximization problem above.

ii.) If the rewards are linear $f_t(x) = x$, the optimal cash price p_t and point value θ_t are given by:

$$(p_t^\star, L_t^\star) \in \arg\max_{p \geq 0, L \geq 0} \Big\{ \alpha\, \mathbb{E}[\, \kappa_t(p, L)\,] - (1-\alpha)\xi L \Big\}, \qquad \theta_t^\star = \frac{\phi_t(p_t^\star, L_t^\star)}{w_t}.$$

Furthermore, $V_t(y) = y - \xi\, L_{t+1}^\star + \sum_{\tau=t+1}^{T+1} \alpha^{\tau-t} \big[\, \mathbb{E}[\kappa_\tau(p_\tau^\star, L_\tau^\star)] + \xi(L_\tau^\star - L_{\tau+1}^\star) \,\big]$, where $L_{T+2}^\star \stackrel{\text{def}}{=} 0$.

iii.) If the rewards are concave, both the optimal LP value and the manager's reward increase in the reward potential, i.e., $L_t^\star(y)$ *and* $y - L_t^\star(y)$ *are increasing in* y.

The result illustrates that a policy dependent on a mixture of profits and cash flows is structurally identical to a profit-dependent policy. In view of this equivalence, all the qualitative insights derived in our previous discussion in Sections 4 and 5 directly apply here, as well. From a quantitative standpoint, however, our next result elicits a dependence of LP value on ξ.

Corollary 3. *Under linear reward functions* $f_t(x) = x$*, the optimal LP value* $L_t^\star$ *decreases in* ξ.

The result shows that, when rewards are linear, the LP value is always decreasing (increasing) in ξ, i.e., as the focus shifts on the profits (cash flows). This matches the intuition that a manager focusing more on cash flows would have a tendency to ignore the firm's liabilities, and thus operate under increased leverage, through larger LP-related deferred revenue.

Although the results and insights for a general reward mixture parallel our earlier findings, it is worth emphasizing an important special case that differs qualitatively. This is summarized in the following result.

Lemma 5. *When the manager's rewards depend only on cash flows, i.e.,* $\xi = 0$*, the optimal policies are independent of the choice of reward functions* f_t*, and are given by Lemma 4(ii.) for* $\xi = 0$.

The lemma shows that when the manager's rewards are entirely tied to the firm's cash flows (i.e., $\xi = 0$), the presence of concave distortions in the reward function carries *no impact* on the operational policies.

7 Conclusions, Limitations, and Future Directions

We studied the problem of optimally setting the value of points for firms operating loyalty programs. We proposed a dynamic model of a firm that sells a single type of product over a discrete time horizon, and awards points that can later be exchanged by customers for additional products. Reporting of cash flows and revenues was subject to standard IFRS guidelines, including the deferred revenue method for accounting for loyalty points. We focused on the pricing policies in cash and points for a manager who would maximize expected discounted profits, and also considered the effects of important managerial considerations.

We showed that the manager's optimal pricing policies mimic classical policies for replenishing and pricing inventory: the value of loyalty points acts as "inventory," which, together with cash and point prices, needs to be adjusted according to a "base-stock, list price" policy. Comparative analysis suggested that managers faced with higher discount factors or lower redemption servicing costs would maintain higher base-stock levels, i.e., more valuable LPs.

When important managerial considerations such as taxation, income smoothing, or risk aversion, were accounted for, our analysis showed how the deferred revenue associated with loyalty points could act as a buffer against uncertainty in operating performance, providing a new rationale for the existence of loyalty programs, as means of hedging financial performance. In particular, we found that managers would follow base-stock policies dependent on operating performance, increasing (decreasing) the value of points under strong (weak) operating performance, to the extent that sales cannibalization due to point redemptions could occur. Comparative analysis showed that, when facing greater uncertainty or longer planning horizons, managers would tend to increase the valuation of points, reinforcing the LP value's role as a buffer.

Finally, we studied several extensions to our model, where the firm's operations involved more complex dynamics, or the manager's objective was tied to both profits and cash flows, and showed that our findings were robust.

7.1 Limitations and Future Directions

Although our model captured the high-level considerations facing managers in charge of setting point values, our framework nonetheless has some limitations, which we now revisit in an attempt to outline fruitful directions for future research.

First, we note that our framework modeled sales and redemptions through aggregate demand response functions, as is typically the case in the literature (see, e.g., Talluri and van Ryzin (2005) and Simchi-Levi et al. (2004)), without explicitly capturing the choice faced by an individual customer, or the heterogeneity of the consumer base. In the context of airlines or hotels, for instance, the individual consumer valuation and its dependency on time could become relevant for the purposes of setting booking limits and point requirements for a particular flight or hotel night stay. Thus, including an explicit model capturing the consumers' strategic purchase decision and the choice of payment method (cash vs. point) could constitute a very interesting direction for future research. Such a model could also allow exploring potential departures from rationality when customers transact with the point currency, such as cases where they suffer from "money illusion" (Shafir et al. 1997).

Second, firms running loyalty programs often provide substitutable products and services in practice, and thus compete with rivals. Furthermore, maintaining their reward platforms often requires entering relationships with various other third-party firms that may also act strategically, to their own benefit. For instance, while a financial services firm provides credit cards to its customers, it also enters agreements with participating merchants—where such cards can be used—as well as third parties—where such points could be redeemed. These considerations warrant several interesting directions for future research, including a more detailed model that captures competition and important third-party interactions.

Finally, our model highlighted a new role for a loyalty program, as a buffer against poor financial performance, and a potential tool for engaging in hedging and earnings smoothing. In this sense, the degree to which managerial compensation is based on profits can carry a direct impact on the firm's (cash and point pricing) policies. This suggests future directions for both analytical and empirical research, examining the extent to which managerial incentives or accounting practices impact the value of loyalty points.

References

Afeche, P., Araghi, M. and Baron, O. (2015), Introduction to special issue on marketing and operations management interfaces and coordination, Technical report, University of Toronto. Working paper.

Aflaki, S. and Popescu, I. (2014), 'Managing retention in service relationships', *Management Science* **60**, 415–433.

AHLA (2014), 'Open letter to Treasury Secretary Jacob Lew'. Retrieved on March 17, 2015 from http://www.ahla.com/uploadedFiles/LoyaltyProgramLew.pdf.

Alaska Airlines (2008), 'Alaska Airlines 10K form for 2008'. Retrieved on June 16, 2016 from http://www.sec.gov/Archives/edgar/data/766421/000119312509033073/d10k.htm.

Babich, V. (2010), 'Independence of capacity ordering and financial subsidies to risky suppliers', *Manufacturing & Service Operations Management* **12**(4), 583–607.

Baghaie, M., Cohen, M., Perakis, G., Rizzo, L. and Sun, X. (2015), 'The effects of social media on online retail', working paper.

Ben-Tal, A. and Teboulle, M. (2007), 'An old-new concept of convex risk measures: The optimized certainty equivalent', *Mathematical Finance* **17**(3), 449–476.

Berry, J. (2015), 'The 2015 loyalty census, big numbers, big hurdles, colloquy report'.

Bertsekas, D. P. (2001), *Dynamic Programming and Optimal Control*, Athena Scientific, Belmont, MA.

Beyer, A. (2009), 'Capital market prices, management forecasts, and earnings management', *The Accounting Review* **84**(6), 1713–1747.

Bitran, G. and Caldentey, R. (2003), 'An overview of pricing models for revenue management', *Manufacturing & Service Operations Management* **5**(3), 203–229.

Boyd, S. and Vandenberghe, L. (2004), *Convex Optimization*, Cambridge University Press.

Brunger, W. G. (2013), 'How should revenue management feel about frequent flyer programs?', *Journal of Revenue and Pricing Management* **12**(1), 1–7.

Caldentey, R. and Haugh, M. (2009), 'Supply contracts with financial hedging', *Operations Research* **57**(1), 47–65.

Chen, X., Sim, M., Simchi-Levi, D. and Sun, P. (2007), 'Risk aversion in inventory management', *Operations Research* **55**(5), 828–842.

Chen, X. and Simchi-Levi, D. (2004), 'Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The finite horizon case', *Operations Research* **52**(6), 887–896.

Chun, S. Y. and Ovchinnikov, A. (2015), Strategic consumers, revenue management and the design of loyalty programs, Technical report, Georgetown University Working Paper.

Chung, H., Chun, S. Y. and Ahn, H. S. (2015), Pricing when consumers buy with loyalty-points, Technical report, Georgetown University Working Paper.

Cohen, D. A., Dey, A. and Lys, T. Z. (2008), 'Real and accrual-based earnings management in the pre- and Post-Sarbanes-Oxley periods', *The Accounting Review* **83**(3), 757–787.

Coughlan, A. and Shulman, J. (2010), Creating superior value by managing the marketing–operations management interface, *in* A. M. Tybout and B. J. Calder, eds, 'Kellogg on Marketing', second edn, John Wiley & Sons, Hoboken, NJ, chapter 20, pp. 393–408.

Dechow, P. M. (1994), 'Accounting earnings and cash flows as measures of firm performance: The role of accounting accruals', *Journal of Accounting and Economics* **18**, 3–42.

Dechow, P. M., Sloan, R. G. and Sweeney, A. P. (1995), 'Detecting earnings management', *The Accounting Review* **70**(2), pp. 193–225.

DeFond, M. L. and Park, C. W. (1997), 'Smoothing income in anticipation of future earnings', *Journal of Accounting and Economics* **23**(2), 115 – 139.

Delta Airlines (2014), 'Delta Airlines management incentive plan for 2014'. Retrieved on June 26, 2016 from https://www.sec.gov/Archives/edgar/data/27904/000002790414000003/dal12312013ex1017.htm.

Dong, L. and Tomlin, B. (2012), 'Managing disruption risk: The interplay between operations and insurance', *Management Science* **58**(10), 1898–1915.

Dorotic, M., Verhoef, P. C., Fok, D. and Bijmolt, H. A. T. (2014), 'Reward redemption effects in a loyalty program when customers choose how much and when to redeem', *International Journal of Research in Marketing* **31**(4), 339 – 355.

Economist, T. (2005*a*), 'Frequent flier miles: In terminal decline?'. Retrieved on October 12, 2015 from http://www.economist.com/node/3536178.

Economist, T. (2005*b*), 'Frequent-flyer miles: Funny money'. Retrieved on October 12, 2015 from http://www.economist.com/node/5323615.

Elmaghraby, W. and Keskinocak, P. (2003), 'Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions', *Management Science* **49**(10), 1287–1309.

Ernst&Young (2014), 'IFRS15 revenue from contracts with customers: A closer look at the new revenue recognition standard', White paper. Retrieved on March 19, 2016 from http://www.ey.com/Publication/vwLUAssets/Applying_IFRS:_A_closer_look_at_the_new_revenue_recognition_standard_(June_2014)/$FILE/Applying-Rev-June2014.pdf.

Federgruen, A. and Heching, A. (1999), 'Combined pricing and inventory control under uncertainty', *Operations Research* **47**(3), 454–475.

Fisher, I. (1928), *The Money Illusion*, Martino Fine Books.

Fox, H. (1980), 'Top executive compensation: 1978', New York: The Conference Board.

Fudenberg, D. and Tirole, J. (1995), 'A theory of income and dividend smoothing based on incumbency rents', *Journal of Political Economy* **103**(1), pp. 75–93.

Gordon, N. and Hlavinka, K. (2015), 'Liability talk, buried treasure, the 2011 forecast of U.S. consumer loyalty program points value, swift exchange and COLLOQUY'.

Graham, J., Harvey, C. and Rajgopal, S. (2005), 'The economic implications of corporate financial reporting', *Journal of Accounting and Economics* **40**(1-3), 3–73.

Healy, P. M. (1985), 'The effect of bonus schemes on accounting decisions', *Journal of Accounting and Economics* **7**, 85–107.

Healy, P. M. and Wahlen, J. M. (1999), 'A review of the earnings management literature and its implications for standard setting', *Accounting horizons* **13**(4), 365–383.

Ittner, C. D., Larcker, D. F. and Rajan, M. V. (1997), 'The choice of performance measures in annual bonus contracts', *The Accounting Review* **72**(2), pp. 231–255.

Kasznik, R. (1999), 'On the association between voluntary disclosure and earnings management', *Journal of Accounting Research* **37**(1), pp. 57–81.

Kim, B. D., Shi, M. and Srinivasan, K. (2004), 'Managing capacity through reward programs', *Management Science* **50**(4), 503–520.

Kopalle, P. K., Sun, Y., Neslin, S. A., Sun, B. and Swaminathan, V. (2012), 'The joint sales impact of frequency reward and customer tier components of loyalty programs', *Marketing Science* **31**(2), 216–235.

Kouvelis, P., Li, R. and Ding, Q. (2013), 'Managing storable commodity risks: The role of inventory and financial hedge', *Manufacturing & Service Operations Management* **15**(3), 507–521.

Lambert, R. (1984), 'Income smoothing as rational equilibrium behavior', *The Accounting Review* **59**(4), 604–618.

Lederman, M. (2007), 'Do enhancements to loyalty programs affect demand? the impact of international frequent flyer partnerships on domestic airline demand', *RAND Journal of Economics* pp. 1134–1158.

Leuz, C., D., N. and Wysocki, P. D. (2003), 'Earnings management and investor protection: an international comparison', *Journal of Financial Economics* **69**(3), 505 – 527.

Lewis, M. (2004), 'The influence of loyalty programs and short-term promotions on customer retention', *Journal of Marketing Research* **41**(3), 281–293.

Li, L., Shubik, M. and Sobel, M. J. (2013), 'Control of dividends, capital subscriptions, and physical inventories', *Management Science* **59**(5), 1107–1124.

Liu, Y. (2007), 'The long-term impact of loyalty programs on consumer purchase behavior and loyalty.', *Journal of Marketing* **71**(4), 19 – 35.

Lu, X. and Su, X. (2015), Revenue management with loyalty programs, Technical report, The Wharton School, University of Pennsylvania. Working paper.

Maglaras, C. and Meissner, J. (2006), 'Dynamic pricing strategies for multiproduct revenue management problems', *Manufacturing & Service Operations Management* **8**(2), 136–148.

Mathies, C. and Gudergan, S. P. (2012), 'Do status levels in loyalty programmes change customers' willingness to pay?', *Journal of Revenue and Pricing Management* **11**(3), 274–288.

McCaughey, N. C. and Behrens, C. (2011), Paying for Status? - The effect of frequent flier program member status on air fare choice, Monash economics working papers, Monash University, Department of Economics.

Oracle (2008), 'Powering the customer-centric airline with loyalty management', White paper. Retrieved on March 19, 2016 from http://www.oracle.com/us/products/applications/siebel/051262.pdf.

Ovchinnikov, A., Boulu-Reshef, B. and Pfeifer, P. E. (2014), 'Revenue management with lifetime value considerations: Balancing customer acquisition and retention spending for firms with limited capacity', *Management Science* **Published online in Articles in Advance 13 Mar 2014**.

Petruzzi, N. and Dada, M. (1999), 'Pricing and the newsvendor problem: A review with extensions', *Operations Research* **47**, 183–194.

Porteus, E. L. (2002), *Foundations of Stochastic Inventory Management*, Stanford Business Books.

Pratt, J. W. (1964), 'Risk aversion in the small and in the large', *Econometrica: Journal of the Econometric Society* pp. 122–136.

PwC (2015), 'Loyalty program strategy and liability management', PricewaterhouseCoopers Website. Retrieved on March 19, 2016 from http://www.pwc.com/us/en/actuarial-insurance-management-solutions/property-casualty/loyalty-program-strategy-liability-management.jhtml.

SAS (2012), 'Facing the challenges of building loyalty and retention: The new strategic imperative', White paper. Retrieved on March 19, 2016 from http://www.sas.com/resources/asset/Strategic_Imperative.pdf.

Shafir, E., Diamond, P. and Tversky, A. (1997), 'Money illusion', *The Quarterly Journal of Economics* **112**(2), 341–374.

Simchi-Levi, D., Chen, X. and Bramel, J. (2004), *The Logic of Logistics: Theory, Algorithms, and Applications for Logistics and Supply Chain Management*, 2 edn, Springer.

Smith, A. and Sparks, L. (2009), '"It's nice to get a wee treat if you've had a bad week": Consumer motivations in retail loyalty scheme points redemption', *Journal of Business Research* **62**(5), 542 – 547.

Smith, C. and Stulz, R. (1985), 'The determinants of firms' hedging policies', *Journal of Financial and Quantitative Analysis* **20**(4), 391–405.

Sun, Y. and Zhang, D. (2014), A model of customer reward program with finite expiration term, Technical report, University of Colorado Boulder.

Talluri, K. T. and van Ryzin, G. J. (2005), *The Theory and Practice of Revenue Management*, Springer Science and Business Media, Inc., New York, NY 10013.

Tang, C. S. (2010), 'A review of marketing-operations interface models: From co-existence to coordination and collaboration', *International Journal of Production Economics* **125**(1), 22–40.

Tang, C. and Teck, T.-H. (2004), 'Introduction to special issue on marketing and operations management interfaces and coordination', *Management Science* pp. 429–430.

Taylor, G. A. and Neslin, S. A. (2005), 'The current and future sales impact of a retail frequency reward program', *Journal of Retailing* **81**(4), 293 – 305.

Topkis, D. M. (1998), *Supermodularity and Complementarity*, Princeton University Press.

Trueman, B. and Titman, S. (1988), 'An explanation for accounting income smoothing', *Journal of Accounting Research* **26**, 127–139.

Verhoef, P. C. (2003), 'Understanding the effect of customer relationship management efforts on customer retention and customer share development', *Journal of Marketing* **67**(4), 30–45.

Xu, X. and Birge, J. R. (2008), 'Operational decisions, capital structure, and managerial compensation: A news vendor perspective', *Engineering Economist* **53**, 173–196.

A Notation

Below is a table with the notation used throughout the paper for easy reference,

s	cash sales
r	redemptions or point sales
p	cash price
θ	point value
λ	points issued per dollar spent
c	per-unit redemption servicing cost
q	point requirement or point price
κ	cash flow
w	outstanding points
g	redemption rate
L	value of the loyalty program or deferred revenue
Π	profit (net income / earnings)
y	reward potential
*	notation for optimal quantities
$\bar{}$	notation for expected values.

B Concavity of κ_t

Our analysis requires the firm's cash flow $\kappa_t \stackrel{\text{def}}{=} p_t\, s_t - c\, r_t$ to be concave in (p_t, L_t), for any realization of the noise. For simplicity, we focus our discussion on the case when

$$\begin{aligned} s_t(p, L) &= \alpha_t^s\, S(p, L) + \beta_t^s, \\ r_t(p, L) &= \alpha_t^r\, R(p, L) + \beta_t^r, \end{aligned} \tag{15}$$

where $S(p, L) \stackrel{\text{def}}{=} \mathbb{E}[s_t(p, L)]$, $R(p, L) \stackrel{\text{def}}{=} \mathbb{E}[r_t(p, L)]$, and $\alpha_t^s \geq 0, \alpha_t^r \geq 0$, (β_t^s, β_t^r) are independent random noise terms, with mean one (zero, respectively). This parallels classic models in the literature, and subsumes the additive and multiplicative noise models as special cases (see, e.g., Simchi-Levi et al. 2004 and Talluri and van Ryzin 2005). In this case, concavity of κ_t holds if $S(p, L) - c\, R(p, L)$ is jointly concave.

Example 1 (Linear sales and quadratic redemptions)**.** *Suppose that S and R take the functional forms:*

$$\begin{aligned} S(p, L) &= -A_s\, p + B_s\, L + C_s \\ R(p, L) &= A_r\, p^2 + B_r\, p\, L + C_r\, L^2 + g(p) + f(L), \end{aligned}$$

where g and f are arbitrary convex functions, and the coefficients $A_s, B_s, C_s, A_r, B_r, C_r$ satisfy the relations:

$$A_s \geq 0, \quad C_r \geq 0, \quad 4c(A_s + A_r c)C_r \geq (B_s - B_r c)^2.$$

Then, κ_t is jointly concave in (p, L) for any realization of the noise terms.

The conditions in this example hold, provided c is not too small. For instance, taking $c \geq \frac{B_s}{B_r}$ when $4A_r C_r > B_r^2$ is sufficient. Even when these conditions are violated, the concavity requirement still holds in a particular range of values for p and L.

Example 2 (Separable sales and price lower bound)**.** *Suppose that R is jointly convex, and*

$$S(p, L) = -A_s(p) + \left(C_s - e^{-B_s L}\right)$$

where $A_s(\cdot)$ is convex increasing, and $e^{B_s L} p\big(2A_s'(p) + pA_s''(p)\big) \geq 1$. Then, κ_t is jointly concave in (p, L), for any noise realization.

These conditions hold whenever suitable lower bounds exist on the cash price p (and L).

Remark 1. *Instead of requiring the revenue function to be concave in prices, it is common in the literature to require it to be concave in the planned demand (see, e.g., Talluri and van Ryzin 2005, Simchi-Levi et al. 2004 and references therein). More precisely, with s_t, r_t given by (15), one would require that, for any value of $\bar{L}$, the function $S(\cdot, \bar{L})$ is continuous and strictly decreasing, with an inverse function $S^{-1}(\cdot, \bar{L})$ such that $\bar{s}\, S^{-1}(\bar{s}, \bar{L})$ is concave in $\bar{s}$. It is known that such requirements are readily satisfied for several functional forms of $S(\cdot, \bar{L})$, such as linear, exponential, iso-elastic or logit (Simchi-Levi et al. 2004). This allows switching the decision variables from (p_t, L_t) to $(\bar{s}_t, L_t)$, where $\bar{s}_t$ denotes the planned (average) sales level. We could adopt a similar approach in our model, and require $\mathbb{E}[\kappa_t]$ to be concave in $(\bar{s}_t, L_t)$. However, to retain more clarity in the exposition and connect more directly with the manager's core operational decisions (the cash and point prices), we prefer maintaining (p_t, L_t) as the decisions.*

C Proofs

Proof of Theorem 1. Note that the representation holds in period $t = T+1$, since $J_{T+1}(w_{T+1}, p_{T+1}, \theta_{T+1}) = \mathbb{E}[f_{T+1}(\kappa_{T+1} + L_{T+1})] \stackrel{\text{def}}{=} \mathbb{E}[V_{T+1}(y_{T+1})]$. Also, since f_{T+1} is concave increasing, so is V_{T+1}.

Assume the representation holds at time $t+1$, and consider the Bellman recursion (5) at time t:

$$\begin{aligned} J_t(w_t, p_t, \theta_t) &= \mathbb{E}\Big[\max_{p_{t+1}, \theta_{t+1}} \Big\{ f_t\big(\kappa_t + L_t - L_{t+1}(w_{t+1}, p_{t+1}, \theta_{t+1})\big) + \alpha\, J_{t+1}(w_{t+1}, p_{t+1}, \theta_{t+1}) \Big\}\Big] \\ &= \mathbb{E}\Big[\max_{p_{t+1}, \theta_{t+1}} \Big\{ f_t\big(y_t - L_{t+1}(w_{t+1}, p_{t+1}, \theta_{t+1})\big) + \alpha\, \mathbb{E}[V_{t+1}(y_{t+1})] \Big\}\Big] \\ &= \mathbb{E}\Big[\max_{p_{t+1}, L_{t+1}} \Big\{ f_t\big(y_t - L_{t+1}\big) + \alpha\, \mathbb{E}[V_{t+1}(y_{t+1})] \Big\}\Big]. \end{aligned}$$

The last step is justified by recalling the *No Money Illusion* and *Redemption Increasing in Value* assumptions (see Section 3). These ensure that y_{t+1} only depends on (p_{t+1}, L_{t+1}) and the random noise in period $t+1$, and that one can equivalently maximize over L_{t+1} instead of θ_{t+1}. The latter follows since w_{t+1} is known and fixed at the time when the decisions (p_{t+1}, θ_{t+1}) are taken, and $\theta_{t+1} = \frac{\phi_{t+1}(p_{t+1}, L_{t+1})}{w_{t+1}}$, where $\phi_{t+1}(p_{t+1}, \cdot) : [0, \infty) \to [0, \infty)$ is a strictly increasing bijection, for any fixed p_{t+1}. This also shows how one can recover the optimal prices $(p^\star_{t+1}, \theta^\star_{t+1})$, proving part (i).

To prove part (ii), note that κ_{t+1} being concave (by our assumption) implies $y_{t+1} = \kappa_{t+1}(p_{t+1}, L_{t+1}) + L_{t+1}$ is concave in (p_{t+1}, L_{t+1}), for any value of the noise in period $t+1$. Since V_{t+1} is concave increasing, we readily have that the function $f_t(y - L_{t+1}) + \alpha\mathbb{E}[V_{t+1}(y_{t+1})]$ is jointly concave in (y, p_{t+1}, L_{t+1}). As such, the partial maximization in (6) will preserve concavity, so that $V_t(y)$ will remain concave (see Boyd and Vandenberghe 2004). Furthermore, since f_t is increasing, V_t will also be increasing. □

Proof of Lemma 1. We prove by induction on t. Note that the decomposition (8) holds trivially for $t = T+1$. Assume it also holds for t, so that $V_t(y) = y + K_t$, where K_t is a constant. Consider the Bellman recursion at $t-1$:

$$
\begin{aligned}
V_{t-1}(y) &= \max_{p_t, L_t} \Big\{ y - L_t + \alpha\, \mathbb{E}\big[V_t(y_t)\big] \Big\} \\
&= \max_{p_t, L_t} \Big\{ y - L_t + \alpha \left(\mathbb{E}[\kappa_t(p_t, L_t)] + L_t - K_t \right) \Big\} \\
&= y - \alpha \cdot K_t + \max_{p, L} \big\{ \alpha\, \mathbb{E}[\kappa_t(p, L)] - (1-\alpha) L \big\}.
\end{aligned}
$$

As such, letting $(p^\star_t, L^\star_t) \in \arg\max\big\{\alpha\, \mathbb{E}[\kappa_t(p, L)] - (1-\alpha)L\big\}$, one can readily see that the cash and point price can be obtained according to (7a) and (7b), respectively, and we have

$$
\begin{aligned}
V_{t-1}(y) &= y - \alpha \cdot K_t + \alpha\, \mathbb{E}[\kappa_t(p^\star_t, L^\star_t)] - (1-\alpha) L^\star_t \\
&= y - \alpha \cdot \left[L^\star_{t+1} + \sum_{k=t+1}^{T+1} \alpha^{k-t} \big[\, \mathbb{E}[\kappa_k(p^\star_k, L^\star_k)] + (L^\star_k - L^\star_{k+1}) \,\big] \right] + \alpha\, \mathbb{E}[\kappa_t(p^\star_t, L^\star_t)] - (1-\alpha) L^\star_t \\
&= y - L^\star_t + \sum_{k=t}^{T+1} \alpha^{k-t+1} \big[\, \mathbb{E}[\kappa_k(p^\star_k, L^\star_k)] + (L^\star_k - L^\star_{k+1}) \,\big],
\end{aligned}
$$

which completes the proof of the inductive step. □

Proof of Corollary 1. By Lemma 1, $(p^\star_t, L^\star_t) \in \arg\max\big\{\alpha\, \mathbb{E}[\kappa_t(p, L)] - (1-\alpha)L\big\}$. Since κ_t is jointly concave in (p, L), the first order optimality condition (FOC) with respect to L yields

$$
\frac{\partial \mathbb{E}[\kappa_t]}{\partial L}\bigg|_{(p^\star_t, L^\star_t)} = \frac{1-\alpha}{\alpha} \geq 0, \tag{16}
$$

which proves part (i). To prove (ii), note that the FOC imply:

$$\left.\frac{\partial \mathbb{E}[p_t s_t]}{\partial p_t}\right|_{(p_t^\star, L_t^\star)} = c \left.\frac{\partial \mathbb{E}[r_t]}{\partial p_t}\right|_{(p_t^\star, L_t^\star)} \geq 0,$$
$$\left.\frac{\partial \mathbb{E}[p_t s_t]}{\partial L_t}\right|_{(p_t^\star, L_t^\star)} = c \left.\frac{\partial \mathbb{E}[r_t]}{\partial L_t}\right|_{(p_t^\star, L_t^\star)} + \frac{1-\alpha}{\alpha} \geq 0. \qquad \square$$

Proof of Lemma 2a. By the FOC in (16), since $\frac{1-\alpha}{\alpha}$ is decreasing in α and $\frac{\partial \mathbb{E}[\kappa_t]}{\partial L}$ is decreasing in L (since $\mathbb{E}[\kappa_t]$ is concave), we have that increasing α would lead to larger values of $L_t^\star$. $\square$

Proof of Lemma 2b. Note that $\frac{\partial^2 \mathbb{E}[\kappa_t(p_t, L_t)]}{\partial c\, \partial p_t} = -\frac{\partial \mathbb{E}[r_t]}{\partial p_t} \leq 0$ by our assumption, so that $\mathbb{E}[\kappa_t(p_t, L_t)]$ is submodular in (c, p_t), and the optimal price $p_t^\star$ will be decreasing in c. A similar argument applies to optimal LP value $L_t^\star$. Lastly, if g_t does not depend on p_t, then we have $L_t^\star = w_t \theta_t\, g_t(w_t\, \theta_t)$, where g_t is increasing. Since w_t is fixed, increasing c, which leads to lower $L_t^\star$ (by the argument above), would also lead to lower $\theta_t^\star$. $\square$

Proof of Lemma 2c. If $\mathbb{E}[\kappa_t(p_t, L_t)]$ is supermodular (submodular) in (p_t, L_t), then the set of maximizers for the problem $\arg\max_p \mathbb{E}[\kappa_t(p, L_t)]$ is increasing (decreasing) in L_t (Topkis 1998), proving the claim. $\square$

Proof of Theorem 2. Consider the main Bellman recursion in Theorem 1 at time $t-1$, and note that it can be rewritten as:

$$V_{t-1}(y) = \max_{L_t} \phi_t(y, L_t), \tag{17a}$$

$$\phi_t(y, L) \stackrel{\text{def}}{=} f_{t-1}\big(y - L\big) + \alpha\, G_t(L) \tag{17b}$$

$$G_t(L) \stackrel{\text{def}}{=} \max_{p_t \geq 0} \mathbb{E}\Big[V_t\big(\kappa_t(p_t, L) + L\big)\Big]. \tag{17c}$$

Here, G_t is concave, since V_t is concave increasing, and $\kappa_t(p, L)$ is jointly concave in its arguments. Also, ϕ_t is jointly concave, and supermodular in (y, L) on the lattice $\mathbb{R}_+^2$, since f_{t-1} is concave (see Topkis 1998).

To prove part (a), note that the maximizer in (17a), $L_t^\star(y)$, must be increasing in y, since ϕ_t is supermodular. Furthermore, by changing variables into $x \stackrel{\text{def}}{=} y - L_t$, problem (17a) can be rewritten as $V_t(y) = \max_x \big[f_{t-1}(x) + \alpha\, G_t(y - x)\big]$. As before, the maximand in the latter problem is supermodular in (x, y) on the lattice $\mathbb{R}_+^2$, since G_t is concave. Therefore, $x^\star(y) = y - L_t^\star(y)$ is increasing in y.

To prove part (b), let $p_t^\star(y) = \arg\max_p \bar{\kappa}_t(p, L_t^\star(y))$. By Assumption 4, we have that $\kappa_t(p_t^\star, L_t^\star) + L_t^\star \geq \kappa_t(p, L_t^\star) + L_t^\star$ holds almost surely, for any p, and thus, since V_t is increasing,

$$\mathbb{E}\Big[V_t\big(\kappa_t(p_t^\star, L_t^\star) + L_t^\star\big)\Big] \geq \mathbb{E}\Big[V_t\big(\kappa_t(p, L_t^\star) + L_t^\star\big)\Big],$$

so that $p_t^\star(y)$ is optimal. The result for $\theta_t^\star(y)$ readily follows from the general result in Theorem 1.

To prove part (c), consider the first-order condition (FOC) yielding $L_t^\star$ in (17a):

$$\begin{aligned} f_{t-1}'(y - L_t^\star) &= \alpha\, G_t'(L_t^\star) \\ \text{(by the Envelope Theorem)} \quad &= \alpha\, \mathbb{E}\Big[\Big(1 + \frac{\partial \kappa_t(p,L)}{\partial L}\Big)\, V_t'\big(\kappa_t(p,L) + L\big)\Big]\Big|_{(p_t^\star, L_t^\star)} \\ \text{(by Assumption 4)} \quad &= \alpha\, \Big(1 + \frac{\partial \bar{\kappa}_t(p,L)}{\partial L}\Big)\, \mathbb{E}\Big[V_t'\big(\kappa_t(p,L) + L\big)\Big]\Big|_{(p_t^\star, L_t^\star)}. \end{aligned}$$

Since f_{t-1} is strictly increasing, and V_t is increasing, we must have that $1 + \frac{\partial \bar{\kappa}_t(p,L)}{\partial L}\big|_{(p_t^\star, L_t^\star)} > 0$, which completes the proof. □

Proof of Theorem 3. We first prove the following useful intermediate results:

(I) $\dfrac{\partial V_t(y,\sigma)}{\partial y}$ is convex in y for all $t = 1, \ldots, T$ and $\sigma \geq 0$.

To ease notation, we omit denoting the dependence on σ and consider the recursion as in (22a-22c). Also, we omit the argument for some functions that are evaluated repeatedly at the same argument (as are their derivatives). In particular, $L_t^\star$ is repeatedly evaluated at y in the expressions below; thus $L_t^\star$ will denote $L_t^\star(y)$. Similarly, the functions f, V_t and ρ (as well as their derivatives) are evaluated at $y - L_t^\star(y)$, $\rho(L_t^\star) + L_t^\star + \sigma\epsilon$ and $L_t^\star$ respectively. In such instances, we will similarly omit their respective argument; for instance, $f^{(2)} = f^{(2)}(y - L_t^\star)$.

By applying the Envelope Theorem and taking the second order derivative we obtain

$$V_{t-1}^{(3)}(y) = f^{(3)}\left(1 - L_{t,y}^\star\right)^2 - f^{(2)} L_{t,yy}^\star, \tag{18}$$

where $L_{t,y}^\star$ denotes the partial derivative of $L_t^\star$ with respect to y. The necessary and sufficient first-order optimality condition that $L_t^\star$ satisfies can be written as $F_t(y, L) = 0$, where

$$F_t(y, L) \stackrel{\text{def}}{=} -f'(y - L) + \alpha(1 + \rho'(L)\mathbb{E}\big[V_t'(\rho(L) + L + \sigma\epsilon)\big]).$$

The maximand of the recursion, ϕ_t, is strictly concave, hence $F_{t,L}(y, L_t^\star) < 0$. To obtain expressions for the derivatives of $L_t^\star$ we apply the Implicit Function Theorem to the above equation, yielding

$$F_{t,y}(y, L_t^\star) + L_{t,y}^\star F_{t,L}(y, L_t^\star) = 0.$$

Applying the Implicit Function Theorem again we get

$$F_{t,yy}(y, L_t^\star) + L_{t,yy}^\star F_{t,L}(y, L_t^\star) + (L_{t,y}^\star)^2 F_{t,L}(y, L_t^\star) + 2L_{t,y}^\star F_{t,yL}(y, L_t^\star) = 0.$$

By substituting for $L^\star_{t,yy}$ in (18) we get

$$V^{(3)}_{t-1}(y) = \underbrace{f^{(3)}\left(1 - L^\star_{t,y}\right)^2}_{\geq 0\ (f'\ \text{convex})} + \underbrace{\frac{f^{(2)}(L^\star_{t,y})^2}{F_{t,L}(y, L^\star_t)}}_{\geq 0\ (f\ \text{concave})} \left(\underbrace{\alpha\rho^{(3)}\mathbb{E}V'_t + 3\alpha\rho^{(2)}(1+\rho')\mathbb{E}V^{(2)}_t}_{\geq 0\ (\rho\ \text{concave},\ \rho'\ \text{convex, Corollary 2})} + \quad \alpha(1+\rho')\mathbb{E}V^{(3)}_t \right).$$

It suffices to show that $V^{(3)}_t$ is non-negative. To see this, note that for $t = T+1$, $V^{(3)}_{T+1} = f^{(3)} \geq 0$, since f' is convex. Thus, one can use an induction argument to complete the proof.

(II) If X is a continuous random variable with zero mean and $f : \mathbb{R} \to \mathbb{R}$ is a differentiable, strictly concave (convex) and increasing (decreasing), then $\mathbb{E}[Xf'(X)] < 0$ (> 0).

Let h denote the probability density function of X. We have,

$$\begin{aligned}
\mathbb{E}[Xf'(X)] &= \int_{-\infty}^{0} xf'(x)h(x)dx + \int_{0}^{\infty} xf'(x)h(x)dx \\
&< \int_{-\infty}^{0} xf'(0)h(x)dx + \int_{0}^{\infty} xf'(x)h(x)dx && [f \text{ is strictly concave and increasing}] \\
&= -\int_{0}^{\infty} xf'(0)h(x)dx + \int_{0}^{\infty} xf'(x)h(x)dx && [X \text{ is zero mean}] \\
&= \int_{0}^{\infty} x(f'(x) - f'(0))h(x)dx < 0. && [f \text{ is strictly concave}]
\end{aligned}$$

The argument is similar for f being convex and decreasing.

(a) Consider the simplified recursion as in (22a-22c). Using the parameterized expression for the noise term we have for all $t = 1, \ldots, T$, y and $\sigma \geq 0$

$$V_t(y,\sigma) = \max_{L_{t+1}} \Big[f\big(y - L_{t+1}\big) + \alpha\mathbb{E}\big[V_{t+1}(\rho(L_{t+1}) + L_{t+1} + \sigma\epsilon, \sigma)\big]\Big]. \tag{19}$$

We have

$$\begin{aligned}
\frac{\partial V_T(y,\sigma)}{\partial\sigma} &= \alpha\mathbb{E}\big[\epsilon\, f'(\rho(L^\star_{T+1}(y,\sigma)) + L^\star_{T+1}(y,\sigma) + \sigma\epsilon)\big] && [\text{by the Envelope Theorem}] \\
&< 0. && [f \text{ is concave increasing + (II)}]
\end{aligned}$$

To complete the proof via induction, assume that $\frac{\partial V_{t+1}(y,\sigma)}{\partial\sigma} < 0$ for all y and $\sigma \geq 0$. Then

$$\begin{aligned}
\frac{\partial V_t(y,\sigma)}{\partial\sigma} &= \alpha\mathbb{E}\big[\epsilon \frac{\partial}{\partial y} V_{t+1}(\rho(L^\star_{t+1}(y,\sigma)) + L^\star_{t+1}(y,\sigma) + \sigma\epsilon, \sigma)\big] \\
&\quad + \alpha\mathbb{E}\big[\frac{\partial}{\partial\sigma} V_{t+1}(\rho(L^\star_{t+1}(y,\sigma)) + L^\star_{t+1}(y,\sigma) + \sigma\epsilon, \sigma)\big] && [\text{by the Envelope Theorem}] \\
&< \alpha\mathbb{E}\big[\epsilon \frac{\partial}{\partial y} V_{t+1}(\rho(L^\star_{t+1}(y,\sigma)) + L^\star_{t+1}(y,\sigma) + \sigma\epsilon, \sigma)\big] && [\text{induction hypothesis}] \\
&< 0. && [V_{t+1} \text{ is concave, increasing in } L]
\end{aligned}$$

We next prove another useful intermediate result.

(III) $\dfrac{\partial^2 V_t(y,\sigma)}{\partial y \partial \sigma} \geq 0$ for all $t = 1, \ldots, T$, y and $\sigma \geq 0$.

By using the expressions above we get

$$\begin{aligned}
\frac{\partial^2 V_T(y,\sigma)}{\partial \sigma \partial y} &= \frac{\partial}{\partial y} \alpha \mathbb{E}\big[\epsilon\, f'(\rho(L^\star_{T+1}(y,\sigma)) + L^\star_{T+1}(y,\sigma) + \sigma\epsilon)\big] \\
&= \alpha \underbrace{(\rho'(L^\star_{T+1}(y,\sigma)) + 1)}_{\geq 0 \text{ by Corollary } 2} \underbrace{\frac{\partial L^\star_{T+1}(y,\sigma)}{\partial y}}_{\geq 0 \text{ by Theorem } 2(a)} \underbrace{\mathbb{E}\big[\epsilon\, f''(\rho(L^\star_{T+1}(y,\sigma)) + L^\star_{T+1}(y,\sigma) + \sigma\epsilon)\big]}_{\geq 0 \text{ by (II) for } f' \text{ convex, } f \text{ concave}} \\
&\geq 0.
\end{aligned}$$

To complete the proof via induction, assume that $\frac{\partial^2 V_{t+1}(y,\sigma)}{\partial \sigma \partial y} \geq 0$ for all y and $\sigma \geq 0$. Then

$$\begin{aligned}
\frac{\partial^2 V_t(y,\sigma)}{\partial \sigma \partial y} &= \alpha \underbrace{(\rho'(L^\star_{t+1}(y,\sigma)) + 1)}_{\geq 0 \text{ by Corollary } 2} \underbrace{\frac{\partial L^\star_{t+1}(y,\sigma)}{\partial y}}_{\geq 0 \text{ by Theorem } 2(a)} \underbrace{\mathbb{E}\big[\epsilon \frac{\partial^2}{\partial y^2} V_{t+1}(\rho(L^\star_{t+1}(y,\sigma)) + L^\star_{t+1}(y,\sigma) + \sigma\epsilon, \sigma)\big]}_{\geq 0 \text{ by (I), (II)}} \\
&\quad + \alpha \underbrace{(\rho'(L^\star_{t+1}(y,\sigma)) + 1)}_{\geq 0 \text{ by Corollary } 2} \underbrace{\frac{\partial L^\star_{t+1}(y,\sigma)}{\partial y}}_{\geq 0 \text{ by Theorem } 2(a)} \underbrace{\mathbb{E}\big[\frac{\partial^2}{\partial \sigma \partial y} V_{t+1}(\rho(L^\star_{t+1}(y,\sigma)) + L^\star_{t+1}(y,\sigma) + \sigma\epsilon, \sigma)\big]}_{\geq 0 \text{ by the induction hypothesis}} \\
&\geq 0.
\end{aligned}$$

(b) Similarly with (I), the necessary and sufficient first-order optimality condition that $L^\star_t(y,\sigma)$ satisfies can be re-written in this case as $F_t(L,\sigma) = 0$, where

$$F_t(L,\sigma) \stackrel{\text{def}}{=} -f'(y-L) + \alpha(1+\rho'(L))\mathbb{E}\big[\frac{\partial}{\partial y} V_t(\rho(L) + L + \sigma\epsilon, \sigma)\big]).$$

Since the maximand of (19) is strictly concave in L_{t+1}, the partial derivative of F_t with respect to L is negative and we can apply the Implicit Function Theorem to obtain

$$\frac{\partial L^\star}{\partial \sigma} = -\frac{\frac{\partial F_t}{\partial \sigma}\big|_{L^\star}}{\frac{\partial F_t}{\partial L}\big|_{L^\star}}.$$

Thus, it suffices to show that the partial derivative of F_t with respect to σ, evaluated at $L^\star$ is non-negative:

$$\begin{aligned}
\frac{\partial F_t}{\partial \sigma}\bigg|_{L^\star} &= \alpha \underbrace{(\rho'(L^\star_t(y,\sigma)) + 1)}_{\geq 0 \text{ by Corollary } 2} \underbrace{\mathbb{E}\big[\epsilon \frac{\partial^2}{\partial y^2} V_t(\rho(L^\star_t(y,\sigma)) + L^\star_t(y,\sigma) + \sigma\epsilon, \sigma)\big]}_{\geq 0 \text{ by (I), (II)}} \\
&\quad + \alpha \underbrace{(\rho'(L^\star_t(y,\sigma)) + 1)}_{\geq 0 \text{ by Corollary } 2} \underbrace{\mathbb{E}\big[\frac{\partial^2}{\partial \sigma \partial y} V_t(\rho(L^\star_t(y,\sigma)) + L^\star_t(y,\sigma) + \sigma\epsilon, \sigma)\big]}_{\geq 0 \text{ by (III)}} \\
&\geq 0.
\end{aligned}$$

(c) The result follows by arguing as in the proof of Lemma 2c and the monotonicity of $L_t^\star$ with respect to σ established above. □

Proof of Theorem 4. Note that the rewards function is piecewise-linear, thus differentiable almost everywhere, except for a finite number of points. All quantities will thus exhibit the same behavior, i.e., V_t and $L_t^\star$. As a result, exchanging the order of integration and differentiation of V_t would still be possible. Furthermore, to ease exposition, we will use the standard derivative notation to denote either the derivative of a function, or any of its subgradients if it is not differentiable at the point it is evaluated.

(a) Consider the simplified recursion as in (22a-22c). The necessary and sufficient first-order optimality condition that $L_t^\star(y,\gamma)$ satisfies can be written as $F_t(L,\gamma) = f'(y-L)$, where

$$F_t(L,\gamma) \stackrel{\text{def}}{=} \alpha(1+\rho'(L))\mathbb{E}\big[\frac{\partial}{\partial y}V_t(\rho(L)+L+\epsilon,\gamma)\big]).$$

Note that the left-hand side term $F_t(L,\gamma)$ is decreasing in L, since V_t is concave in L, whereas the right-hand side term $f'(y-L)$ is increasing in L. In particular, the right-hand side term takes the value of 1 for $L < y - \hat{\Pi}$, any value between 1 and γ for $L = y - \hat{\Pi}$, and γ for $L > y - \hat{\Pi}$. Consequently, there exist values $\underline{y}_{t+1}$ and $\overline{y}_{t+1}$ such that $L_t^\star(y,\gamma)$ satisfies

(i) $F_t(L_t^\star(y,\gamma),\gamma) = \gamma$, for $y < \underline{y}_t$,

(ii) $L_t^\star(y,\gamma) = y - \overline{y}$, for $\underline{y}_t \le y \le \overline{y}_t$, and

(iii) $F_t(L_t^\star(y,\gamma),\gamma) = 1$, for $y > \overline{y}_t$.

Suppose that $y \le \overline{y}_t$. Then, either $L_t^\star(y,\gamma)$ is constant (case (ii)), or it satisfies the condition in (i). Using the notation as in the proof of Theorem 3, the Implicit Function Theorem yields

$$F_{t,\gamma}(L_t^\star(y,\gamma),\gamma) - 1 + \frac{\partial L_t^\star}{\partial \gamma}F_{t,L}(L_t^\star(y,\gamma),\gamma) = 0, \tag{20}$$

where $F_{t,L}(L_t^\star(y,\gamma),\gamma) < 0$ by the concavity of V_t. Also,

$$\begin{aligned} F_{t,\gamma}(L_t^\star(y,\gamma),\gamma) &= \alpha(1+\rho'(L_t^\star(y,\gamma)))\mathbb{E}\big[\frac{\partial}{\partial\gamma}\frac{\partial}{\partial y}V_t(\rho(L_t^\star(y,\gamma)) + L_t^\star(y,\gamma)+\epsilon,\gamma)\big]) \\ &= \frac{\mathbb{E}\big[\gamma\frac{\partial}{\partial\gamma}\frac{\partial}{\partial y}V_t(\rho(L_t^\star(y,\gamma)) + L_t^\star(y,\gamma)+\epsilon,\gamma)\big])}{\mathbb{E}\big[\frac{\partial}{\partial y}V_t(\rho(L_t^\star(y,\gamma)) + L_t^\star(y,\gamma)+\epsilon,\gamma)\big])} \\ &\le 1. \end{aligned}$$

The second equality above follows by substituting for $\alpha(1+\rho'(L_t^\star(y,\gamma)))$ using the condition in (i). For the inequality, note that at points at which the functions are differentiable (and these are the relevant ones for the expectations above) we have

$$\frac{\partial}{\partial y}V_t(y,\gamma) = f'(y - L_{t+1}^\star(y)). \tag{21}$$

The right-hand side above takes values 1 or γ. As such,

$$\gamma \frac{\partial}{\partial \gamma} \frac{\partial}{\partial y} V_t(y, \gamma) = \gamma \frac{\partial}{\partial \gamma} f'(y - L^\star_{t+1}(y)) \leq f'(y - L^\star_{t+1}(y)) = \frac{\partial}{\partial y} V_t(y, \gamma).$$

Using the bounds $F_{t,L}(L^\star_t(y,\gamma),\gamma) < 0$ and $F_{t,\gamma}(L^\star_t(y,\gamma),\gamma) \leq 1$, equation (20) yields that $\frac{\partial L^\star_t}{\partial \gamma} \geq 0$ for case (i). For case (ii), the inequality still holds as $L^\star_t$ is constant. Thus, $\frac{\partial L^\star_t}{\partial \gamma} \geq 0$ for $y \leq \overline{y}_t$.

To complete the proof, note that for $y \leq \overline{y}_t$ and case (iii), the equivalent of equation (20) is

$$F_{t,\gamma}(L^\star_t(y,\gamma),\gamma) + \frac{\partial L^\star_t}{\partial \gamma} F_{t,L}(L^\star_t(y,\gamma),\gamma) = 0,$$

and it suffices to show that $F_{t,\gamma}(L^\star_t(y,\gamma),\gamma) \geq 0$. This follows however from Corollary 2 and equation (21), since f' is increasing in γ.

(b) By Lemma 3(b), $\frac{\partial}{\partial y} V_t(y,\gamma)$ is decreasing in t. Consequently, $F_t(L,\gamma)$ is also decreasing in t and the result follows.

(c) As we remarked above $F_t(L,\gamma)$ is increasing in γ and the result follows. □

Proof of Lemma 3. We prove both parts together, by backwards induction. To simplify notation, in view of the result in Theorem 4(b) and Assumption 4, the recursion in (17a-17b) can be rewritten as

$$V_{t-1}(y) = \max_{L_t} \phi_t(y, L_t), \tag{22a}$$

$$\phi_t(y, L) = f_{t-1}(y - L) + \alpha\, \mathbb{E}\Big[V_t\big(\rho_t(L) + \varepsilon_t + L\big)\Big], \text{ where} \tag{22b}$$

$$\rho_t(L) \stackrel{\text{def}}{=} \max_{p \geq 0} \bar{\kappa}_t(p, L). \tag{22c}$$

When the data are stationary, i.e., $f_t = u$, $\bar{\kappa}_t = \bar{\kappa}$, and $\rho_t = \rho$, the Envelope Theorem for (22a) yields:

$$V'_t(y) = f'(y - L^\star_{t+1}), \forall\, t \in \{1, \ldots, T\}.$$

Since $V_{T+1}(y) = u(y)$, we readily have that $V'_T(y) = f'(y - L^\star_{T+1}) \geq V'_{T+1}(y) = f'(y)$, since u is strictly concave (so that f' is decreasing). Furthermore, we also have $L^\star_{T+1}(y) \geq L^\star_{T+2}(y) \equiv 0, \forall\, y$. Thus, the properties hold at time $T+1$. Assume they also hold at t, so that $V'_t(y) \geq V'_{t+1}(y)$. Then, consider the FOC for problem (22a) written at time $t-1$, yielding $L_t^\star$, and note that:

$$\begin{aligned}
\frac{\partial \phi_t}{\partial L}\Big|_{L_{t+1}^\star} &= \Big\{ f'(y-L) + \alpha \underbrace{\big(1+\rho'(L)\big)}_{\geq 0,\text{ by (c)}} \mathbb{E}\Big[V_t'\big(\rho(L)+\varepsilon_t + L\big)\Big]\Big\}\Big|_{L_{t+1}^\star} \\
&\geq \Big\{ f'(y-L) + \alpha \big(1+\rho'(L)\big) \mathbb{E}\Big[V_{t+1}'\big(\rho(L)+\varepsilon_t + L\big)\Big]\Big\}\Big|_{L_{t+1}^\star} \\
&= \frac{\partial \phi_{t+1}}{\partial L}\Big|_{L_{t+1}^\star} \\
&= 0.
\end{aligned}$$

As such, it must be that $L_t^\star \geq L_{t+1}^\star$. In turn, this implies that $V_{t-1}'(y) = f'(y-L_t^\star) \geq f'(y-L_{t+1}^\star) = V_t'(y)$, completing the proof of the inductive step. □

Proof of Lemma 5. The proof follows analogously to Lemma 1, and is omitted. □