Loyalty Program Liabilities and Point Values

So Yeon Chun Dan A. Iancu Nikolaos Trichakis*

Abstract

Problem definition. Loyalty programs (LP) introduce a new currency—the points—through which customers transact with firms. Such points represent a promise for future service, and their monetary value thus counts as a liability on the issuing firms' balance sheets. Consequently, adjusting the value of points has a first order effect on profitability and performance, and emerges as a core operating decision. We study the problem of optimally setting the points' value in view of their associated liabilities.

Academic / Practical Relevance: Firms across numerous industries increasingly utilize LPs. The sheer magnitude of LPs coupled with recent changes in accounting rules have turned the associated liabilities into significant balance sheet items, amounting to billions of dollars. Managers (from CFOs to CMOs) struggle with the problem of adjusting the points' value in view of these liabilities. Academic work is primarily aimed at understanding LPs as marketing tools, without studying the liability angle.

Methodology. We develop a multi-period model and use dynamic programming techniques and comparative statics analysis.

Results: We show that the optimal policies depend on a new financial metric, given by the sum of the firm's realized cash flows and outstanding deferred revenue, which we refer to as the *profit potential*. The total value of loyalty points is set to hit a particular target, which increases with the profit potential. We find that loyalty programs can act as buffers against uncertainty, with the value of points increasing (decreasing) under strong (weak) operating performance, and increasing with uncertainty.

Managerial Implications: Setting the point values and adjusting operating decisions in view of LP liabilities should be done by tracking the firm's profit potential. Loyalty programs can act as hedging tools against uncertainty in future operating performance, which provides a new rationale for their existence, even in the absence of competition.

*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 MIT Sloan School of Management, (ntrichakis@mit.edu).

1 Introduction

Originally designed as marketing tools for rewarding customers, loyalty programs have expanded dramatically in size and scope during recent years, with total memberships in the U.S. reaching 3.3 billion in 2014 (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%).

In a typical "point-based" loyalty program (LP), members earn points for purchases from the issuing firm, and can redeem accumulated points for rewards, such as additional products, services or even cash. For consumers, points thus effectively become a new currency, often carrying substantial value — for instance, approximately 14 trillion miles worth more than $700B were outstanding in 2005 (Economist 2005), and the annual reward value issued in the US alone exceeded $48B in 2015 (Gordon and Hlavinka 2015).

For the issuing firm, however, loyalty points represent a promise for future service, and their value thus constitutes a *liability*. The sheer magnitude of LPs coupled with recent changes in accounting rules[1] have turned these liabilities into significant balance sheet items — for instance, at the end of 2015, they amounted to $3.9B for Delta Airlines and $2.6B for Marriott International, or 10% and 25% of their respective total liabilities (Delta Airlines 2015, Marriott 2015). As such, it is easy to see that the value of points can dramatically impact firms' earnings and profitability.

In view of this, setting the value of loyalty points emerges as a key operating decision for the firm. In practice, this is usually done by changing the point requirements for redemptions, or by adjusting the exchange rates for converting points into cash. For instance, Marriott changes the point requirements for a free night stay at its properties on an annual basis, by re-categorizing the properties and/or adjusting the points required for each category (Schlappig 2016, Marriott 2017). In addition, Marriott also alters point values on a daily basis by, e.g., changing the available inventory of rooms for redemption.

Understanding how point values should be set optimally in view of the liabilities they generate for the firm is the main focus of our paper. To elaborate, we first discuss the unique accounting standards governing the calculation of LP-related liabilities. Under rules recently set by the International Financial Reporting Standard (IFRS) — which will become mandatory in the U.S. in 2018 — a firm is required to treat any points issued in connection with a cash sale as a separate

[1] In the U.S., loyalty points have been traditionally accounted for using an incremental cost method, under GAAP. The new rules, which were issued jointly by the Financial Accounting Standards Board and the International Accounting Standards Board under "IFRS 15 Revenue from Contracts with Customers" in 2014, are already a required standard in Europe, Canada, and Australia, and will be required in the U.S. starting in January 2018. This new standard results in significantly larger liabilities than the incremental cost method; e.g., following Chapter 11 reorganization, Delta Airlines switched to the IFRS standard, increasing its LP liabilities from $412M to $2.4B.

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 upon the initial sale. However, when points are redeemed or expire, the firm can *recognize* a corresponding amount from its deferred revenue liabilities, 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. In particular, the total value of a firm's LP-related deferred revenue is 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 known as the *redemption rate*).

In view of these rules, it can be seen how changes in points' value impact profitability due to the deferral process: increasing (decreasing) the value of a point translates into more (less) deferred revenue, which directly hurts (improves) profits/earnings. For firms with billions of outstanding points, even small changes pertaining to loyalty points can thus have a *first-order* impact. For instance, according to Delta Airline's 10-K statement for 2015, "A hypothetical 10% increase in [mile value] would decrease [revenue] by approximately $48 million, as a result of an increase in the amount of revenue deferred" (Delta Airlines 2015). Such changes are not only hypothetical, but do in fact arise. In 2008, Alaska Airlines decided to shorten its points' expiration date from three years to two. This change reduced the total value of its points and the associated deferred revenue, enabling the airline to claim an additional $42.3M in revenue, and reduce its consolidated net losses for the year by a staggering 24% (Alaska Airlines 2008).

The Alaska Airlines example also highlights how reducing point values (and thus their 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 valuation decisions. Similarly, since the deferral process influences the revenue taxable year of inclusion, taxation can also become another important managerial consideration influencing loyalty point valuations (AHLA 2014).

This discussion prompts several natural research questions. How should a firm's manager adjust the value of loyalty points, in view of the liabilities they create? And how is this operational decision impacted by important considerations, such as taxation or earnings smoothing incentives, cost of capital, shocks or volatility in the operating performance, or consumers' perception of the value of the firm's loyalty points? To the best of our knowledge, despite the practical importance of the questions—as recognized by a wide range of industry white papers (see, e.g., Oracle 2008, SAS 2012, Ernst&Young 2014)—little or no academic work has been done on the topic.

We address these questions by developing a dynamic model of a firm that sells a single product,

and awards customers who purchase in cash with points that can be redeemed for additional 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 horizon. The cash sales and the amount of products purchased through redemptions depend on both the cash prices and point requirements, in potentially non-monotonic ways (so as to capture increased sales due to loyalty effects, but also possibly cannibalization due to increased redemptions). The manager's goal is to maximize expected discounted rewards tied to profits. We consider concave reward functions, so as to capture the effects of tax considerations, earnings smoothing incentives, or risk aversion, as well as linear reward functions, so as to capture gross profit maximization.

Our findings and contributions. Our paper is the first to study how to dynamically adjust the monetary value of loyalty points in view of the liabilities they generate for the issuing firm. We formulate the manager's decision problem as a dynamic program (DP) with a high-dimensional state, which includes the number of outstanding loyalty points, the cash price, and the point requirement (or the exchange rate for points into cash). We show that under two mild assumptions concerning consumers' rationality and the firm's accounting practices, the DP state collapses into one variable, given by the sum of the firm's realized cash flows and outstanding deferred revenue. This new financial metric, which we refer to as the *profit potential*, emerges as a key summary of the firm's performance, and a critical driver of decisions concerning loyalty points.

The reformulation also allows characterizing the manager's optimal policy. We find that the core managerial decision concerning loyalty points is their *total outstanding value*, which should be set to hit a particular target. Once this target is set—depending on the observed profit potential—and the cash price is optimally adjusted, the optimal point requirement and the exchange rate for points into cash can be inferred using the balance of outstanding points and the expected redemption rate.

We characterize the dependency of the manager's decisions on several important factors. For instance, we find that the value of points and the associated deferred revenue liability is increasing with the profit potential, but at a slow enough rate to allow reported earnings to also increase. Managerial incentives such as taxation, income smoothing or risk aversion prove critical: in their absence, point values can be set independently of the profit potential. Under a higher marginal tax rate or risk aversion, point values could either increase[2] or decrease, depending on whether the firm's profit potential exceeds or falls short of certain milestones, which decrease over time.

Finally, we confirm that our main findings are robust under several extensions, namely when the firm runs more complex operations (selling multiple products, carrying inventory, updating multiple

[2]Throughout the paper, we use "increasing," "decreasing," etc. in a weak (i.e., non-strict) sense.

decisions more frequently, using complex LP designs), and when the manager's rewards are tied to both cash flows and profits. In addition, to test the validity of our two modeling assumptions, we also present and analyze a micro-founded model of consumer behavior, where individuals (with different point balances and perceptions about the point-cash exchange rate) choose whether to purchase products, and whether to use cash or points.

Managerial insights. Our structural results show that in order to make operating decisions in view of loyalty program liabilities, managers need to keep track of the firm's *profit potential*, and use it to set targets for the total value of points, subsequently adjusting point requirements and/or the exchange rates for points into cash to meet these targets. This also suggests a succinct interaction and a potential decentralization of decisions across the firm's offices, as the treasury office may be more prominently involved with the former part of the process, while the operations or marketing groups may cater to the latter.

Our finding that the value of points increases with the profit potential, while ensuring that reported earnings also increase, highlights an entirely new function for loyalty programs. Namely, the deferred revenues associated with an LP can act as a *buffer* or *hedging tool* against uncertainty in operating performance: when performance is strong (leading to large profit potentials), it is optimal for managers to inflate the points' value so as to defer a larger portion of the revenue for future access; when performance is weak, it is optimal for managers to deflate the points' value, so as to recognize more deferred revenue and boost current earnings. This provides a potential explanation for the Alaska Airlines example, and a new rationale for loyalty programs: while traditionally viewed as a means for softening competition (Kim et al. 2004), such programs may be beneficial even without competition, due to their hedging capability.

Several of our subsequent findings become significantly more transparent in view of the (new) role of the LP as a hedging tool against uncertainty. Faced with higher market volatility, leading to more variable revenues/costs/cash flows, managers should enlarge the value of points, thus ensuring a larger pool of deferred revenues to tap into in future times of need. Similarly, managers with longer planning horizons should ensure larger buffers of deferred revenue are available, and should lower the magnitude of such buffers over time, reflective of their diminishing remaining benefits. Lastly, managers who discount the future less and/or have access to cheaper sources of capital should maintain larger buffers of deferred revenue (associated with points), as the inherent time value loss becomes less hurtful.

Interestingly, our results suggest that managers who are faced with an increased marginal tax or who are more risk averse should not necessarily reduce the value of points. Instead, they should

either increase or decrease this value, depending on whether the firm's profit potential exceeds a specific target (i.e., the firm "has gains") or not (i.e., the firm "has losses"). Under a higher tax rate, managers with gains should increase the points' value and the deferred revenue, saving some of the gains for the future, while managers with losses should reduce the points' value by recognizing more deferred revenue, thus reducing current losses. Managers should set these targets internally, and the targets should be lowered over time, and increased with the marginal tax rate. A similar behavior should be followed under increased earnings smoothing incentives or risk aversion.

Finally, we would like to acknowledge that the consumers' response to (either positive or negative) point value adjustments could have important implications. For example, media backlash followed the sudden announcement of plans by the U.K. grocery chain Tesco to reduce its Clubcard rewards vouchers' value in January 2018, which prompted the firm to delay the changes. Carefully modeling all features of consumer behavior is outside the scope of the current paper. Herein, we only partially account for consumer response through the loyalty and cannibalization effects discussed in §3, and through the analysis in §5.3.

1.1 Literature Review

Our paper is 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, 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 made across several industries from a "quantity-

based" to a "spending-based" design, and study the impact of strategic customer behavior on the firm's profit and consumer surplus. 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 effectively charge higher prices, and that the switch to a "spending-based" design could be profitable. In contrast, our paper focuses on the interplay between the loyalty points' value and the liabilities this generates for a firm.

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) for reviews of this topic. Our model assumes that the firm does not engage in any accounting or reporting manipulation; this puts our work closer to *real* earnings management, which is the practice of altering earnings by changing operational decisions. Such practices have reportedly increased following the Sarbanes-Oxley Act in 2002, when "firms switched [...] to managing earnings using real methods, possibly because these techniques, while more costly, are likely to be harder to detect" (Cohen et al. 2008). Healy and Wahlen (1999) and Fudenberg and Tirole (1995) discuss several operational levers that can be used to manage earnings, including sales acceleration, changes in shipment schedules or the delay of maintenance activities. Roychowdhury (2006) finds empirical evidence that price discounts, overproduction, and reductions in discretionary spending are also used in practice to alter earnings. There is also a sparse literature in operations management that discusses earnings management. Lai et al. (2011) show how managers can use channel stuffing (i.e., the practice of shipping excess inventory to the downstream channel) to report higher sales and influence the investors' valuation of the firm. Also related are several recent papers that show how firms can alter their inventory levels by over or under-ordering, in order to signal demand information to outside investors (see, e.g., Lai et al. 2012, Schmidt et al. 2015, and Lai and Xiao 2017). We contribute to this literature by showing how a firm's loyalty program can also serve as a tool for earnings management, and specifically earnings smoothing.

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). On the one hand, 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. On the other hand, Dorotic et al. (2014*a*) and Kopalle et al. (2012*a*) 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

Our main model only includes the critical ingredients needed for capturing the key drivers underpinning managerial decisions on loyalty point values in view of their liabilities. This allows us to derive optimal policies and structural insights in a general-purpose and industry-independent setting. In §5, we show that our insights are robust by considering more realistic operational models, more explicit models for consumer choice, and more general managerial compensation schemes. We discuss limitations in §6, where we outline fruitful directions for future research.

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. The planning horizon allows capturing an employment contract with a finite duration, but also a possibly longer, firm-level horizon (with $T \to \infty$). 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 marginal 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*. Equivalently, since any point requirement induces a monetary value (i.e., an exchange rate) of $\theta_t = \frac{p_t}{q_t}$ for one point, we can also consider the decisions as the *cash price* p_t and the *point value* θ_t. In view of our choice of a period, one can think of p_t and θ_t as average values or targets enforced during a subsequent fiscal quarter or year (see §5.2 for a model where these decisions are made more frequently).

During period t, the firm's customers buy $\tilde{s}_t$ products in cash, and acquire $\tilde{r}_t$ products by

redeeming points. Both the *cash sales* $\tilde{s}_t$ and *redemptions* $\tilde{r}_t$ are random, and depend on the cash price p_t, the point value θ_t, the number of outstanding points w_t, and exogenous noise $\tilde{\varepsilon}_t$. We make no assumptions concerning the monotonicity of these dependencies—in particular, we allow $\tilde{s}_t$ and $\tilde{r}_t$ to either decrease or increase with p_t and θ_t, and only require that $\tilde{\varepsilon}_t$ be independent across time. For convenience of notation, we use s_t (r_t) to denote the realizations of $\tilde{s}_t$ ($\tilde{r}_t$); we omit showing the explicit dependencies of $\tilde{s}_t$ and $\tilde{r}_t$ for now, but return to discuss them extensively in §3.

The firm awards points to its customers at a fixed rate of λ points for every dollar spent; this results in a total of $\lambda\, p_t\, s_t$ new points issued in connection with the realized cash sales in period t. In contrast, redemptions result in $q_t r_t$ points deducted from customer accounts, so that the balance of outstanding points at the end of period t (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* during 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 at 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,

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

The redemption rate g_t, which is estimated by the firm,[3] depends on (p_t, θ_t, w_t); we revisit this

[3] In practice, to estimate the redemption rate, firms usually rely on historical program experience, as well as consideration of enacted program changes (see, e.g., American Airlines 2016).

dependency in detail in §3. 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 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. (For simplicity, we take the initial w_1, L_1, p_1, θ_1 as fixed.)

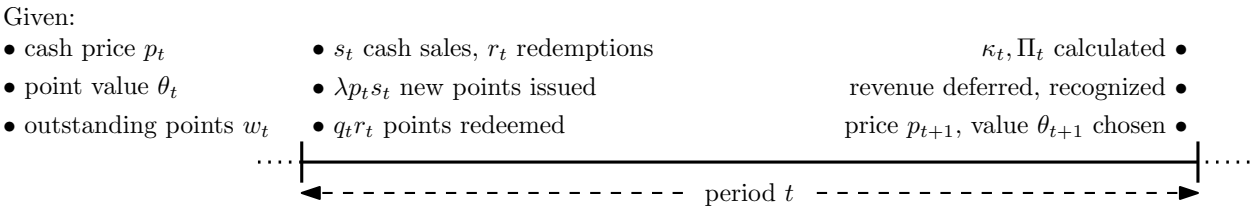

Figure 1: Timeline of events during period t.

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

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

We assume that the firm and its manager do not manipulate any of their estimates or the resulting reported accounting metrics, including the redemption rate, revenues, profits, etc.

The manager's decision problem. The manager obtains a reward $f_t(\Pi_t)$ tied to the firm's profits, where f_t is a concave, increasing function. The manager's problem is to select a policy for setting the cash price and point value, $\{p_t, \theta_t\}_{t=2}^{T+1}$, so as to maximize his cumulative, expected discounted rewards over the given time-frame, i.e., $\sum_{t=1}^{T+1} \alpha^t\,\mathbb{E}[f_t(\Pi_t)]$. Here, $\alpha \in (0, 1]$ is a discount factor, and we take $\Pi_{T+1} := \kappa_{T+1} + L_{T+1}$, i.e., all deferred revenue is recognized at the end of the terminal period. Studying a strictly concave reward function f_t allows us to capture several practical managerial considerations, which become particularly pertinent in the context of LP management:

- *Taxation.* Post-tax profit can be expressed as a concave function of pre-tax profit (Smith and Stulz 1985). While taxation is often ignored in the operations literature, such a simplification can be problematic in our case, due to the significant effect of LP point values on post-tax profits.[4] According to U.S. Income Tax Law, the taxable year of inclusion of

[4]To further illustrate the significance of taxation related to LPs, 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

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 Π_t.

- *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 and references therein). 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).
- *Risk aversion.* Managers are often averse to risks, a preference that can be adequately reflected through a concave utility function (see, e.g., Pratt 1964, Smith and Stulz 1985, etc).

When f_t is linear, we also recover the classical objective of maximizing the firm's (pre-tax) profits.

3 Dynamic Programming Formulation and Optimal Policy

The manager's decision problem can be formulated as a stochastic dynamic program (DP) (Bertsekas 2001). A sufficient state is given by the triple (w_t, p_t, θ_t), since the random cash sales $\tilde{s}_t$, the redemptions $\tilde{r}_t$, and the redemption rate g_t depend on it. The manager seeks an admissible policy for setting $\{p_t, \theta_t\}_{t=2}^{T+1}$, i.e., a policy that is adapted to the available information. With J_t denoting the manager's value function at the beginning of period t, the Bellman recursion can be written as:

$$
\begin{aligned}
J_t(w_t, p_t, \theta_t) &= \mathbb{E}_{\tilde{\varepsilon}_t}\left[\max_{p_{t+1},\theta_{t+1}\geq 0}\Big(f_t\big(\Pi_t\big) + \alpha\, J_{t+1}(w_{t+1}, p_{t+1}, \theta_{t+1})\Big)\right] && (5)\\
\Pi_t &= \kappa_t + L_t - L_{t+1} && \forall\, t \in \{1,\dots,T\}\\
\kappa_t &= p_t\, \tilde{s}_t - c\, \tilde{r}_t, && \forall\, t \in \{1,\dots,T+1\}\\
L_t &= w_t\, \theta_t\, g_t, && \forall\, t \in \{1,\dots,T+1\}\\
w_{t+1} &= w_t + \lambda\, p_t\, \tilde{s}_t - q_t\, \tilde{r}_t, && \forall\, t \in \{1,\dots,T\},
\end{aligned}
$$

where $J_{T+1}(w_{T+1}, p_{T+1}, \theta_{T+1}) = \mathbb{E}\big[f_{T+1}\big(\kappa_{T+1} + L_{T+1}\big)\big]$. 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 redemptions, s_t and r_t, respectively (see our

Priority Guidance Plan. While this action only indicated that the law may undergo changes (without specifying 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)

discussion in §2).

As stated, the problem is not readily amenable for analysis due to the high-dimensional state, and the non-linear dynamic evolution. Note that this is the case even in our stylized operational setting — where the firm does not carry any inventory, sells a single product, etc. — and would only be compounded by a more realistic firm model. Fortunately, it turns out that the following mild assumptions enable tractability.

Assumption 1. The expected cash sales $\mathbb{E}_{\tilde{\varepsilon}_t}[\tilde{s}_t]$, the expected redemptions $\mathbb{E}_{\tilde{\varepsilon}_t}[\tilde{r}_t]$, and the redemption rate g_t depend on (w_t, θ_t) only through the product $w_t \cdot \theta_t$.

The assumption requires that, on average, the outstanding loyalty points affect the aggregate purchasing and redemption behavior only through their monetary value (i.e., the product $w_t \cdot \theta_t$), rather than individually. To understand this at an intuitive level, 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 Assumption 1, the firm's aggregate sales and redemptions would not be affected, on average.

This assumption is aligned in spirit with standard rationality requirements in finance and economics, which state that rational decision makers should not suffer from "money illusion," i.e., that purchasing decisions should be in terms of the real value of money, instead of the nominal one (Fisher 1928). It is important though to note that our requirement is weaker, in the sense that it does not concern the decision-making of a single individual, but rather the *average, aggregate* outcomes observed by the firm. In fact, we provide further support for this assumption in §5.3, where we discuss a broad class of micro-founded consumer choice models that satisfy it, even when individual consumers care separately about the number of points they have, or have different (and possibly biased) perceptions concerning the value of a single loyalty point.

As a byproduct of Assumption 1, note that the firm's deferred revenue, $L_t = w_t\,\theta_t\,g_t$, also only depends on the state through $(p_t, w_t\,\theta_t)$. Our next assumption imposes a weak requirement on this dependency, by asking L_t to be strictly increasing in the points' total monetary value $w_t\,\theta_t$.

Assumption 2 (Point Liability Increasing in Monetary Point Value). $L_t(p_t, w_t\,\theta_t)$ is strictly increasing in $w_t\,\theta_t$, for any fixed p_t.

Assumption 2 states that the deferred revenue liability L_t associated with the LP points should strictly increase with the actual monetary value of the outstanding points, $w_t\,\theta_t$. In practice, this is a very natural assumption, since the liability calculated and reported for accounting purposes should reflect the points' value, i.e., an increase in the points' value should result in an increase in the liabilities (see §5.3 for an example and further discussion).

In view of Assumption 2, there is a one-to-one relation between L_t and $w_t\,\theta_t$, for any fixed p_t. Thus, all quantities of interest can be expressed as functions of the cash price and LP value, i.e.,

$$\tilde{s}_t,\ \tilde{r}_t,\ g_t,\ \text{and } \kappa_t \text{ are functions of } p_t,\ L_t \text{ and } \tilde{\varepsilon}_t.$$

We make the important remark that our assumptions do not imply any monotonicity of the (expected) cash sales $\mathbb{E}[\tilde{s}_t]$ in the value of LP points L_t. In particular, increasing the value of points could increase cash sales, as more customers purchase the products to earn the more valuable points, but could also decrease them, due to excessive redemptions. Both effects have been observed in practice and documented empirically, under the names of loyalty (Lewis 2004) and cannibalization (Kopalle et al. 2012*a*), respectively. We direct the interested reader to §5.3 for a class of micro-founded models where both effects are present, and additional discussion.

These assumptions allow us to revisit the DP formulation and characterize the manager's optimal policy, as formalized in the next result.

Theorem 1. *For any time $t \in \{1,\dots,T\}$,*
(a) the manager's optimal policy is to set a cash price $p^\star_{t+1}(y_t)$ and a total value of loyalty points $L^\star_{t+1}(y_t)$ that depend on $y_t \overset{\text{def}}{=} \kappa_t(p_t, L_t, \tilde{\varepsilon}_t) + L_t$, and are the optimal actions in the 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 $V_{T+1}(y) = f_{T+1}(y)$;
(b) the optimal value of a point is determined as

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

where $\phi_t(p_{t+1}, \cdot) : [0,\infty) \to [0,\infty)$ is a strictly increasing bijection;
(c) the manager's optimal value function can be written as $J_t(w_t, p_t, \theta_t) = \mathbb{E}_{\tilde{\varepsilon}_t}[V_t(y_t)]$.

According to Theorem 1, the core managerial decisions are the cash price p_{t+1} and the *total* LP value L_{t+1}. These are set as functions of a single (new) state variable y_t, which we henceforth refer to as the firm's *profit potential*, given by the sum of the firm's cash flow κ_t and current LP value L_t. According to equation (6), the manager seeks to optimally split the profit potential y_t into reported earnings, $y_t - L_{t+1}$, and total value of loyalty points in the next period, L_{t+1}. While the former quantity generates immediate rewards, the latter is "invested" in the future, impacting the

profit potential $\tilde{y}_{t+1} = \kappa_{t+1}(p_{t+1}, L_{t+1}, \tilde{\varepsilon}_{t+1}) + L_{t+1}$ through direct and indirect channels, due to the cash flow κ_{t+1}. This also highlights the fundamental tradeoffs faced by the manager in setting a high value of loyalty points, which sacrifices immediate profits and incurs time-value loss, but may improve future cash flows, e.g., by increasing sales through a loyalty effect.

The manager's problem is also reminiscent of a classical tradeoff in operations management, when pricing and adjusting inventory levels under uncertainty. Specifically, with the deferred revenue associated with the LP playing the role of inventory, it can be seen that maintaining a more valuable LP (i.e., "holding more inventory") involves an immediate sacrifice in profits (i.e., "ordering costs"), and incurs time-value loss (i.e., "overage/holding" costs). On the other hand, maintaining an undervalued LP (i.e., "holding less inventory") generates opportunity costs from missed sales, due to a weak loyalty effect (i.e., underage/backlogging costs). Different from standard inventory models though, the level of inventory here can be adjusted downwards (i.e., akin to inventory disposal), and the price charged may be either higher or lower, depending on the strength of the loyalty effect.

According to (b), once the price and total LP value are determined, the manager can infer a corresponding point value θ_{t+1} that preserves the consistency of all (financial accounting) calculations. This is done through the invertible map ϕ_{t+1}, and requires knowledge of the number of outstanding points w_{t+1}. Note that although the calculation involves tracking w_{t+1}, it does not complicate the manager's decision problem, which involves solving the one-dimensional DP in (6).

4 Comparative Statics and Managerial Implications

The compact characterization of the manager's optimal policy allows us to examine how the core decisions—and, critically, the total value of loyalty points—are influenced by several important considerations, such as shocks in the firm's cash flows, variability in cash flows, the structure of the manager's reward function, the discount rate, or the cost of redemptions. For tractability purposes, we make the following technical assumption concerning the functional form of the firm's cash flows κ_t, effective throughout our subsequent analysis.

Assumption 3: $\kappa_t(p_t, L_t, \tilde{\varepsilon}_t) = \bar{\kappa}_t(p_t, L_t) + \sigma\tilde{\varepsilon}_t$, where $\bar{\kappa}_t(p_t, L_t) \stackrel{\text{def}}{=} \mathbb{E}[\kappa_t(p_t, L_t, \tilde{\varepsilon}_t)]$ is concave in (p_t, L_t), the noise terms $\tilde{\varepsilon}_t$ have zero mean and unit variance, and $\sigma \geq 0$.

The concavity of $\bar{\kappa}_t$ parallels classic requirements in the literature (e.g., Petruzzi and Dada 1999, Federgruen and Heching 1999, and Talluri and van Ryzin 2005), and reflects that the incremental benefits of loyalty are diminishing and/or the incremental costs of loyalty liability are increas-

ing. The assumption of additive noise provides a simple and intuitive way to quantify variability, through the standard deviation σ (see, e.g., Federgruen and Heching 1999, Chen and Simchi-Levi 2004). While many of our results continue to hold under more general noise models, we adopt this parameterization to streamline the analysis.

4.1 Impact of Profit Potential and Loyalty Points Acting as a Buffer Against Uncertainty

Our first result discusses how the value of loyalty points is affected by the firm's operating performance, as summarized in the profit potential y_t.[5]

Theorem 2. *For any time $t \in \{2, \ldots, T+1\}$, the optimal value of loyalty points increases in the profit potential, at a rate smaller than 1. In particular,*
(a) the LP value $L_t^\star(y_t)$ is increasing in y_t;
(b) the firm's reported profit $\Pi_t^\star = y_t - L_t^\star(y_t)$ is increasing in y_t.

Theorem 2 derives an important new insight: it shows that the deferred revenue associated with loyalty points can act as a revenue buffer against poor performance, and can thus be used to smoothen a firm's earnings. To illustrate this, suppose that, ceteris paribus, the firm's operating cash flows κ_t increase (decrease) by an amount of Δ, e.g., due to stronger (weaker) sales or a decrease (increase) in costs. Consequently, the profit potential y_t would also increase (decrease) by Δ, leading the manager — according to Theorem 2 — to increase (decrease) the LP value by an amount *less than* Δ, and to report earnings that are increased (decreased) by *less than* Δ. Effectively then, managers faced with stronger operating performance would defer a larger part of the revenue for future access, while managers faced with mediocre performance 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 firm competition.

The result provides a possible explanation for the Alaska Airlines example discussed in the introduction: when experiencing reduced cash flows (e.g., due to unexpectedly large fuel costs, as in Alaska's case), a firm can reduce the value of its loyalty points (e.g., by reducing their expiration date), which allows it to recognize additional revenue and compensate for the losses partially.

The findings also have interesting implications for the firm's customers, suggesting that they always "share the pain and the gain" with the firm. That is, improved operating performance

[5] Theorem 2 continues to hold without Assumption 3, provided that any realization of the cash flows $\kappa_t(p, L, \tilde{\varepsilon}_t)$ is concave in (p, L), with no further restrictions on the dependency on $\tilde{\varepsilon}_t$.

always induces more valuable future promises for the loyal customers, through an inflated LP, as well as larger immediate profits/earnings for the firm (and larger rewards for the firm's manager).

4.2 Impact of Variability

Our results demonstrated that managers 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 cash flows may critically drive the value of the firm's loyalty points.

To isolate the effect of variability, we analyze a problem with stationary primitives, i.e., where $\tilde{s}_t \equiv \tilde{s}$, $\tilde{r}_t \equiv \tilde{r}$, and $\tilde{\varepsilon}_t$ are i.i.d. To avoid uninteresting cases, we also assume that the problem parameters are such that it is optimal to offer a loyalty program (i.e., $L_t^\star = 0$ is not an optimal solution). Let $V_t(y, \sigma)$ denote the value function in (6) when the standard deviation of cash flows is σ, and let $L_t^\star(y, \sigma)$ denote the optimal LP value. The following result distills the impact of variability on the LP value.

Theorem 3. *Suppose that the model primitives are stationary. Let* $\rho(L) \stackrel{\text{def}}{=} \max_p \bar{\kappa}(p, L)$. *If* f' *and* ρ' *are convex, then for all* $t = 2, \ldots, T+1$ *and for all* y,
(a) the optimal LP value $L_t^\star(y, \sigma)$ *is increasing in* σ,
(b) the value function $V_t(y, \sigma)$ *is decreasing in* σ.

Part (a) contains a potentially surprising insight: that a manager faced with increased variability in cash flows should actually *increase* the value of loyalty points, 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 interpreting the LP as a "safety stock" (of cash) held in anticipation of future fluctuations in performance, with a larger stock being preferable under increased uncertainty. Part (b) confirms the intuition that a manager derives less value under increased variability—an intuitive consequence of Jensen's inequality.

We note that the conditions in the theorem are not overly restrictive. Convexity of f' is a reasonable assumption, satisfied by the vast majority of commonly used utility functions, including the entire family of HARA utilities. Convexity of ρ' is a more technical condition, introduced for tractability. Note that it is readily satisfied for the important class of quadratic revenue models, i.e., under linear price impact.

4.3 Impact of Taxation, Earnings Smoothing and Risk Aversion

To examine the role of taxation or income smoothing incentives, it is instructive to first analyze the case when reward functions are linear, taken without loss as $f_t(\Pi) = \Pi$, $\forall t \in \{1, \ldots, T+1\}$. In this important special case, which corresponds to the goal of maximizing the firm's *gross* expected discounted profits, the manager's policy changes substantially, as formalized in the next result.

Theorem 4. *When reward functions are linear,*[6] *the optimal value of loyalty points L_t and the optimal cash price p_t are set independently of the firm's profit potential:*

$$(p_t^\star, L_t^\star) \in \operatorname*{arg\,max}_{p \geq 0, L \geq 0} \big\{ \alpha \cdot \mathbb{E}[\, \kappa_t(p, L, \tilde{\varepsilon}_t)\,] - (1-\alpha) \cdot L \big\} \tag{7a}$$

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

The result suggests that when the manager's objective is to maximize gross profits, he should set the LP value and the cash price *independently* of the firm's current operating performance, i.e., of the profit potential y_t. This lies in contrast with Theorem 1, and highlights the impact of incentives due to, e.g., taxation, earnings smoothing incentives, or risk aversion on the role and value of the LP. In the absence of such effects, the LP's role as an income smoothing buffer is dimished, and the optimal value of points reflects a simpler trade-off: L_t is chosen to balance the time-value loss with the potential improvements in cash flows (e.g., due to an increased loyalty effect), and the cash price p_t is always set to maximize the resulting expected cash flows.

Given the central role of such incentives, it is then natural to also ask how *changes* would impact the value of points. For instance, if such incentives arise from tax considerations, would a firm faced with a larger tax burden prefer to lower the value of its LP points? Similarly, if incentives arise due to risk aversion, would a more risk-averse manager prefer to lower the LP-related liabilities by maintaining less valuable LPs?

To address such questions, we consider the following family of parameterized reward functions:

$$f_t(\Pi) = \begin{cases} \gamma \cdot (\Pi - \hat{\Pi}) + \hat{\Pi}, & \Pi \leq \hat{\Pi}, \\ \Pi, & \Pi > \hat{\Pi}, \end{cases} \quad \text{for all } t = 1, \ldots, T, \tag{8}$$

where $\gamma \geq 1$. For $\gamma = 1$, we recover the case of linear rewards. As γ increases, the effects of concavity become more pronounced. Such piece-wise linear rewards/utilities have been studied

[6] Assumption 3 is not required for this result.

before in the literature (see, e.g., Ben-Tal and Teboulle 2007). Although other parameterizations are clearly possible, this choice renders tractability and remains suitable for capturing important effects such as taxation (e.g., increasing γ is a substitute for increasing the marginal tax rate for profits above $\hat{\Pi}$[7]) or risk aversion (e.g., γ captures 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 (8). The next result summarizes the impact of γ on the firm's LP.

Theorem 5. *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 the tax rate or the degree of risk aversion should cause managers to maintain less valuable LP points, Theorem 5 shows that this is not always the case. In fact, this intuition is reversed when the firm is faced with a "sufficiently good" performance, as determined by the current profit potential y exceeding a certain threshold $\hat{y}_t$. In such cases, increasing the marginal tax rate (or risk aversion) would lead to a *larger* LP-related liability and a *larger* value for the loyalty points.

To understand the effect, note that $\hat{y}_t$ can be thought of as an *adjusted* target that the manager sets internally. Profit potentials above (below) this target are then considered "gains" (respectively, "losses"). When the firm currently has "gains," an increased tax rate (or risk aversion) would result in more deferred revenue and a larger value of points, so as to hedge against future losses. On the contrary, if the firm currently faces "losses," an increased tax rate (or risk aversion) would result in less deferred revenue and a lower value of points, so as to mitigate the present losses. This also shows that managers of otherwise identical firms could respond quite differently to increased tax rates (or risk aversion), depending on the firms' current financial prospects, which would also generate different benefits for the firms' (loyal) consumers.

Finally, the threshold $\hat{y}_t$ being decreasing in t and increasing in γ suggests that managers would set higher targets early in their planning horizon, and as they become more risk averse.

[7] To see this, consider the alternative parameterization where the slope is 1 for $\Pi \leq \hat{\Pi}$, and δ for $\Pi > \hat{\Pi}$. Increasing γ is equivalent to decreasing δ, i.e., a higher marginal tax rate.

4.4 Impact of Time and Planning Horizon

Our next result characterizes the effect of time and the planning horizon on the firm's LP.

Lemma 1. *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 profit 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 would thus tend to inflate the point values early on. To understand this, recall that LPs carry a positive effect on future performance by facilitating the ease of hedging against uncertainty, through the larger buffer of deferred revenues. This capacity diminishes as fewer time steps remain, which induces the manager to reduce the point values over time. The result also suggests that, ceteris paribus, there may be a positive relationship between the length of a manager's planning horizon (e.g., the length of the employment contract) and the value of the firm's LP.

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 an inventory unit decreases over time (see, e.g. Talluri and van Ryzin 2005).

4.5 Impact of Discount Factor and Cost of Capital

Our next result highlights the dependency of the LP value on the discount factor.

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

Under larger discount factors, the (inaccessible) deferred revenue associated with the loyalty program incurs a smaller opportunity cost, leading managers to prefer larger LP values. Viewed under a financial lens, the result also suggests that firms facing a lower cost of capital (i.e., higher α) will tend to operate under higher leverage, by increasing their LP-related liabilities.

4.6 Impact of Loyalty Program on Cash Price

Our first result characterizes the impact of running an LP on the cash price charged by the firm.

Lemma 3. *If the expected cash flow $\bar{\kappa}_t(p_t, L_t)$ 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 suggests that whether managers should charge lower or higher cash prices when operating an LP critically depends on whether loyalty and price have complementary or substitutable effects on the cash flow. For instance, in contexts where the loyalty effect does not strictly decrease the customers' willingness to pay, $\bar{\kappa}_t$ is likely to be supermodular. Several recent empirical papers confirm this to be the case in the travel and hospitality industry (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, $\bar{\kappa}_t$ is likely to be submodular, so having (more valuable) loyalty programs would warrant lower cash prices.

4.7 Impact of Redemption Cost

We next study how the manager's decisions depend on the redemption cost c.

Lemma 4. *If the expected cash flow $\bar{\kappa}_t(p_t, L_t)$ is supermodular in (p_t, L_t) and the expected redemptions are increasing in the cash price and the LP value (i.e., $\frac{\partial \mathbb{E}[\tilde{r}_t]}{\partial p_t} \geq 0, \frac{\partial \mathbb{E}[\tilde{r}_t]}{\partial L_t} \geq 0$), 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 decreasing in 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 LP 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.

5 Extensions and Robustness Checks

We now extend our model in several important dimensions. First, we consider firms with more complex operational models: selling multiple products, carrying and replenishing inventory, operating LPs with more complex schemes for awarding points (§5.1), updating their prices or point values more frequently (§5.2), and compensating managers based on both cash flows and profits (§5.4). For each setting, we confirm that our main results and insights pertaining to loyalty point

valuation remain unchanged. Lastly, we introduce a micro-founded model where consumers choose whether to purchase products, and whether to use cash or points (§5.3); apart from validating our assumptions about the aggregate cash sales and redemption behavior, this also allows us to derive a number of new managerial insights concerning the impact of certain behavioral biases on point values and cash prices.

5.1 More Complex Operating Model

We 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 possibly 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 $x_t \in \mathbb{R}^{n_x}$, and the firm's manager takes a set of constrained actions $a_t \in \mathcal{A}(x_t) \subseteq \mathbb{R}^{n_a}$ corresponding to operating decisions. The firm's operations during period t generate total sales of $\tilde{s}_t \in \mathbb{R}^m$ for the m sold 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 $\tilde{\varepsilon}_t$. The firm's manager obtains a reward f_t tied to the firm's profit during the period, and seeks an operating policy $\{a_t\}_{t=2}^{T+1}$ 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, the firm's 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 $q_t = \frac{p_t}{\theta_t} \in \mathbb{R}^m$ for the m products. During period t, the firm now observes cash sales of $s_t \in \mathbb{R}^m$ and redemptions of $r_t \in \mathbb{R}^m$, and correspondingly issues $\Lambda_t(s_t, p_t)$ new points and retracts $r_t^\top q_t$ points. Furthermore, a (potentially random) fraction $\tilde{\xi}_t \in [0, 1]$ of the unused points expires during the period, so that $w_{t+1} = (1 - \tilde{\xi}_t)(w_t - r_t^\top q_t) + \Lambda_t(s_t, p_t)$. As a result of the sales, the firm's 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 state x_t, on the decisions a_t, on the outstanding points w_t, and the monetary value θ_t, and are affected by exogenous randomness $\tilde{\varepsilon}_t$. As before, the firm's manager seeks a policy for setting the operating decisions and monetary point values $\{a_t, \theta_t\}_{t=2}^{T+1}$ that maximizes his cumulative, discounted rewards.

It is worth noting that our base model is a special case of this more general framework, with

$x_t = \emptyset$, $a_t = p_t$, $\Lambda_t(s_t, p_t) = \lambda p_t s_t$, $\tilde{\xi}_t = 0$, and $\kappa_t = p_t s_t - c r_t$. This framework can capture firms with more complex dynamics, such as retailers/manufacturers deciding replenishment and/or production quantities and selling prices, or airlines and hotels adjusting booking limits to manage capacity. It also allows different LP designs—e.g., awarding points based on sales volume, $\Lambda_t(s_t, p_t) = \lambda s_t$, or a mix of volume and cash expenditures—and it allows some of the points to expire.

As in our analysis in §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:

$$J_t(x_t, w_t, a_t, \theta_t) = \mathbb{E}\Big[\max_{\substack{a_{t+1}\in\mathcal{A}(x_t)\\ \theta_{t+1}\geq 0}} \Big(f_t\big(\kappa_t + L_t - L_{t+1}\big) + \alpha\, J_{t+1}(x_{t+1}, w_{t+1}, a_{t+1}, \theta_{t+1})\Big)\Big], \tag{9}$$

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 (9), 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.

Under our earlier assumptions that (1) customers' aggregate choices are only impacted by the points' value on average, and (2) that the point liability is increasing in the monetary point value, we can run the same argument as in §3 to conclude that s_t, r_t, and κ_t only depend on $(x_t, a_t, L_t, \tilde{\varepsilon}_t)$, and we have $x_{t+1} = X_t(x_t, a_t, L_t, \tilde{\varepsilon}_t)$, $\kappa_t = K_t(x_t, a_t, L_t, \tilde{\varepsilon}_t)$, for some functions X_t, K_t.

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

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

i.) the manager's optimal policy is to set operating decisions $a^\star_{t+1}(y_t, x_{t+1})$ and a total value of loyalty points $L^\star_{t+1}(y_t, x_{t+1})$ that only depend on the firm's current profit potential $y_t := \kappa_t + L_t$ and on the state x_{t+1}, and are optimal actions in the recursion:

$$V_t(y_t, x_{t+1}) = \max_{a_{t+1}\in\mathcal{A}(x_{t+1}),\, L_{t+1}\geq 0} \Big[f_t\big(y_t - L_{t+1}\big) + \alpha\, \mathbb{E}\big[V_{t+1}(y_{t+1}, x_{t+2})\big]\Big]. \tag{10}$$

Furthermore, the optimal value of a point is determined as $\theta^\star_{t+1} = \frac{\phi_{t+1}(x_{t+1}, a_{t+1}, L_{t+1})}{w_{t+1}}$, where $\phi_{t+1}(x_{t+1}, a_{t+1}, \cdot) : [0, \infty) \to [0, \infty)$ is an increasing bijection, and the manager's optimal value function is given by $J_t(x_t, w_t, a_t, \theta_t) = \mathbb{E}[V_t(y_t, x_{t+1})]$.

ii.) the optimal value of loyalty points increases in the profit potential at any given firm state, i.e., $L^\star_t(y, x)$ is increasing in y, for any x.

iii.) if reward functions are linear, the LP value $L_t^\star$ and the operating decisions $a_t^\star$ are set independently of the firm's profit potential, as optimal actions in the recursion:

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

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 profit potential y_t between realized profit during the present period, $y_t - L_{t+1}$, and LP value invested in the future.

Part *ii.*) mirrors Theorem 2(a), and reveals that 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., $L_{t+1}^\star$ 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. Furthermore, if we also required X_t and K_t to be jointly concave, we can readily check that $y - L_t^\star(y)$ is increasing in y (for any fixed operating state x), mirroring the results in Theorem 2(b).

Finally, part *iii.*) parallels Lemma 4, and illustrates that managerial considerations such as taxation, earnings smoothing or risk aversion play a critical role, and in their absence point values and operating decisions would be unaffected by the firm's realized profit potential.

5.2 Frequent Updating of Prices and Point Values

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), $i \in \{1,\ldots,N\}$, 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.[8] Provided that the expected cash flow achieved during period t—when maximizing over price $\boldsymbol{p}_t \stackrel{\text{def}}{=} (p_{t,1},\ldots,p_{t,N})$ 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 be the case if the expected cash flows achieved in every micro-period (t,i)

[8] 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.

were concave in the firm's decisions in that period. To see this, note that $\mathbb{E}_{\tilde{\varepsilon}_t}[\kappa_t(p_t, L_t, \tilde{\varepsilon}_t)] \stackrel{\text{def}}{=} \max_{\boldsymbol{p}:\, \boldsymbol{e}^\top \boldsymbol{p}/N = p_t} \sum_{i=1}^{N} \mathbb{E}[\kappa_{t,i}(p_{t,i}, L_t, \tilde{\varepsilon}_{t,i})]$ remains jointly concave in (p_t, L_t) if $\mathbb{E}[\kappa_{t,i}]$ are jointly concave, so that all our results carry over. Similar arguments can also be employed to address more frequent updates of the point value θ_t.

5.3 Micro-founded Consumer Choice Model

Our model captured the firm's sales and redemptions through aggregate response functions $\tilde{s}_t$ and $\tilde{r}_t$, respectively, which were required to satisfy certain assumptions (see §3). This section introduces a more refined model that accounts for the purchasing behavior of individual consumers, who choose whether to buy a product and whether to use cash or points. Using this model, we then show that the resulting aggregate sales and redemptions are consistent with our earlier assumptions (see Lemma 5), which confirms the robustness of our model and results so far. This micro-founded model also enables us to examine how the manager's decisions depend on potentially relevant aspects of consumer behavior, such as a bias in the perceived point value.[9]

To keep notation simple, we suppress time dependency in this section. We consider a population of N consumers, indexed by $i \in \{1, \ldots, N\}$. Consumer i (she) has a random valuation for the firm's product $\tilde{v}_i \geq 0$, is endowed with a random number of loyalty points $\tilde{\omega}_i \geq 0$, and perceives each point to be worth $\tilde{\gamma}_i \theta$ monetary units, where $\tilde{\gamma}_i \geq 0$ is a random bias governing whether she underestimates, overestimates or exactly agrees with the point value θ set by the firm. Our model thus captures consumer heterogeneity in multiple dimensions, including the willingness to pay, the point balance, and—more importantly—the perceived value of a point. The latter assumption builds on several empirical findings in the marketing literature. Liston-Heyes (2002) and Basumallick et al. (2013) demonstrate that the perceived point value may differ substantially across consumers and may exhibit systematic biases, e.g., due to specific marketing techniques used or due to consumers' cognitive limitations (e.g., not having all the information or computational abilities to correctly assess changing values). Furthermore, Kivetz and Simonson (2002*a*,*b*) suggest that this perceived value may also vary depending on the individual effort required to obtain the point reward, or on the guiltiness of hedonic (instead of utilitarian) consumption.

The ith consumer observes the cash price p and point requirement $q = p/\theta$, and considers a purchase as follows. If she does not have enough points to redeem for the product, i.e., $\tilde{\omega}_i < q$, then

[9] We note that the choice of currency—i.e., cash versus points—is in itself a new research area, with several recent papers devoted solely to the topic (see, e.g., Chun and Hamilton 2017). As such, our goal in introducing such a choice model here is primarily to illustrate the robustness of our main findings; a detailed study of consumer behavior is a very interesting topic, but one that arguably lies outside the main scope of our paper.

she purchases with cash if and only if her valuation exceeds the charged cash price, i.e., $\tilde{v}_i > p$. On the other hand, if she has enough points, then she considers a purchase in either points or cash—whichever is less costly to her. Specifically, given that a point purchase would have a perceived cost of $\tilde{\gamma}_i \theta q$, she considers a point purchase if $\tilde{\gamma}_i \theta q \leq p$, and a cash purchase otherwise. In the former (latter) case, she opts to purchase if and only if her valuation exceeds her perceived cost (the cash price), i.e., $\tilde{v}_i > \tilde{\gamma}_i \theta q$ ($\tilde{v}_i > p$). Note that, by allowing heterogeneity in perceived point values, the ith consumer may purchase in either points or cash, depending on whether $\tilde{\gamma}_i$ is below or above 1, respectively. As a side benefit, this feature also allows our model to capture realistic aspects of choice behavior, such as the cannibalization of regular cash sales due to excessive product redemptions, documented in Kopalle et al. (2012*b*) and Dorotic et al. (2014*b*). To capture the effect that propensity to consume with the firm may increase as one has access to more valuable rewards or points (see, e.g., Lewis 2004, Liu 2007, Kopalle et al. 2012*b*), we assume that the i-th consumer has a random loyalty threshold $\tilde{\xi}_i \geq 0$, and considers a purchase with the firm if and only if her perceived value of her accumulated rewards is "high" enough, i.e., $\tilde{\gamma}_i \theta \tilde{\omega}_i \geq \tilde{\xi}_i$.

With $\tilde{s}_i$ and $\tilde{r}_i$ denoting the indicators of whether the ith customer makes a cash or a point purchase, respectively, we then have that

$$\tilde{s}_i = \mathbb{1}\{\tilde{\gamma}_i \theta \tilde{\omega}_i \geq \tilde{\xi}_i, \tilde{\omega}_i < q, \tilde{v}_i > p\} + \mathbb{1}\{\tilde{\gamma}_i \theta \tilde{\omega}_i \geq \tilde{\xi}_i, \tilde{\omega}_i \geq q, \tilde{\gamma}_i \theta q > p, \tilde{v}_i > p\},$$
$$\tilde{r}_i = \mathbb{1}\{\tilde{\gamma}_i \theta \tilde{\omega}_i \geq \tilde{\xi}_i, \tilde{\omega}_i \geq q, \tilde{\gamma}_i \theta q \leq p, \tilde{v}_i > \tilde{\gamma}_i \theta q\}.$$

The firm's aggregated cash sales and redemptions thus become $\tilde{s} = \sum_{i=1}^{N} \tilde{s}_i,\ \tilde{r} = \sum_{i=1}^{N} \tilde{r}_i$.

We assume that $\tilde{v}_i$, $\tilde{\omega}_i$, $\tilde{\gamma}_i$, and $\tilde{\xi}_i$ are independent across different customers, and (respectively) identically distributed, with c.d.f. F_v, F_ω, F_γ, F_ξ, and p.d.f. f_v, f_ω, f_γ, and f_ξ. With regard to the point distribution F_ω, we assume it is strictly increasing and $F_\omega(x) = H\left(\frac{x}{\mathbb{E}\tilde{\omega}}\right), \forall x$, for some function H. This assumption is purely of technical nature, and can be verified for many practically relevant distributions, including exponential, uniform on $[0, U]$, Pareto, and lognormal.

With this micro-founded model for consumer choice, we are now ready to prove our first result concerning the expected aggregate sales and redemptions. Let $\bar{H}(x) = 1 - H(x)$.

Lemma 5. *The expected cash sales, expected redemptions, and redemption rate are given by*

$$\mathbb{E}[\tilde{s}] = N\,\bar{F}_v(p)\int_0^\infty\left(\int_{\xi/p}^\infty\left(\bar{H}\left(\frac{N\xi}{\gamma\theta w}\right)-\bar{H}\left(\frac{Np}{\theta w}\right)\right)f_\gamma(\gamma)d\gamma+\int_1^\infty\bar{H}\left(\frac{N\max\{p,\xi/\gamma\}}{\theta w}\right)f_\gamma(\gamma)d\gamma\right)f_\xi(\xi)d\xi,$$

$$\mathbb{E}[\tilde{r}] = N\int_0^\infty\int_0^1\bar{F}_v(\gamma p)\bar{H}\left(\frac{N\max\{p,\xi/\gamma\}}{\theta w}\right)f_\gamma(\gamma)f_\xi(\xi)d\gamma d\xi,$$

$$g = \frac{Np}{\theta w}\int_0^\infty\int_0^1\bar{F}_v(\gamma p)\bar{H}\left(\frac{N\max\{p,\xi/\gamma\}}{\theta w}\right)f_\gamma(\gamma)f_\xi(\xi)d\gamma d\xi.$$

In particular, $\mathbb{E}[\tilde{s}]$, $\mathbb{E}[\tilde{r}]$ and g are functions of p and $\theta\cdot w$, and Assumption 1 is satisfied. Furthermore, $\theta\cdot w\cdot g$ is strictly increasing in $\theta\cdot w$ for any fixed p, and Assumption 2 is satisfied.

These results confirm that Assumptions 1 and 2 introduced in §3 arise naturally in this class of choice models, providing further support to our main analysis and conclusions. As a side note, the result highlights that even when an individual consumer's choice depends on her own point balance, or when consumers disagree about the point-cash exchange rate, Assumption 1 is nonetheless reasonable for characterizing the *aggregate* outcomes observed by the firm (i.e., expected sales and redemptions, redemption rate), which only depend on the points' monetary value.

This micro-founded model can also be leveraged to study how specific behavioral parameters (such as a bias in point value perception $\tilde{\gamma}$) affect the firm's optimal decisions. We direct the interested reader to §A of the paper's online appendix for more details.

5.4 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, and Ittner et al. 1997).

To capture this feature, 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. Here, ξ and $1-\xi$ can be thought of as weights used in the compensation plan (see, e.g., Delta Airlines 2014).

Lemma 6. *When the manager's rewards depend on x_t,*

i.) the manager's optimal policy is to set a cash price $p^\star_{t+1}(y_t)$ and a total value of loyalty points $L^\star_t(y_t)$ that depend on the firm's profit potential $y_t \stackrel{\text{def}}{=} p_t\,s_t(p_t,L_t)-c\,r_t(p_t,L_t)+\xi\cdot L_t$, and

are optimal actions in the recursion

$$V_t(y) = \max_{\substack{p_{t+1}\geq 0 \\ L_{t+1}\geq 0}} \Big[f_t\big(y - \xi \cdot L_{t+1}\big) + \alpha\, \mathbb{E}\big[V_{t+1}(y_{t+1})\big]\Big], \tag{11}$$

where $V_{T+1}(y) = f_{T+1}(y)$. *Furthermore, the optimal point value can be obtained as*

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

where $\phi_{t+1}(p_{t+1}, \cdot) : [0,\infty) \to [0,\infty)$ *is an increasing bijection.*

ii.) the optimal LP value is increasing in the profit potential at a rate less than 1, i.e., $L^\star_t(y)$ *and* $y - L^\star_t(y)$ *are increasing in* y.

iii.) if f_t *are linear, the optimal cash price and the value of points are set independently of the profit potential, and are given by:*

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

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 discussions directly apply here, as well. From a quantitative standpoint, however, our next result elicits a dependence of LP value on ξ.

Corollary 1. *Under linear reward functions* $f_t(x) = x$, *the optimal LP value* $L^\star_t$ *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, which we summarize next.

Lemma 7. *When the manager's rewards depend only on cash flows, i.e.,* $\xi = 0$, *the optimal policies are independent of the reward function* f_t, *and are given by Lemma 6*(*iii.*) *for* $\xi = 0$.

The lemma shows that when rewards are entirely tied to cash flows ($\xi = 0$), operational policies are again set independently of the firm's profit potential, and considerations such as taxation, earnings smoothing or risk aversion carry no impact.

6 Conclusions, Limitations, and Future Directions

We studied the problem of optimally setting the value of loyalty points in view of their deferred revenue liabilities.

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, our base model dealt with a generic firm, selling a single product, and adjusting cash prices over time. It would be insightful to build a more detailed model that captures the specifics of certain industries (e.g., travel and hospitality vs. financial services vs. retail), and examine more closely how setting point values interacts with other operational features.

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.

Third, 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 value of a firm's loyalty points. This suggests future directions for analytical and empirical research examining how managerial incentives or accounting practices impact the value of points.

Lastly, devaluing the loyalty points may alienate customers and, in some industries, significantly hurt the firm's market share, as in the example of Tesco. Modeling all aspects of consumer behavior and understanding their implications would be an interesting direction for future work.

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', *Mngm. Science* **60**, 415–433.

AHLA (2014), 'Open letter to Treasury Secretary Jacob Lew'. https://tinyurl.com/y7zu4sd3.

Alaska Airlines (2008), 'Alaska Airlines 10K form for 2008'. https://tinyurl.com/ya6wouej.

American Airlines (2016), 'American Airlines 10K form for 2016'. https://seekingalpha.com/filing/3421547.

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

Basumallick, D., Ozdaryal, B. and Madamba-Brown, C. (2013), 'Perceived value of a mile', *Journal of Revenue & Pricing Management* **12**(1), 8–15.

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'.

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

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

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 Hamilton, R. (2017), 'Should I pay for this purchase or redeem points? Effects of loyalty program design on consumer decisions to redeem points', working paper.

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.

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', *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'. https://tinyurl.com/hyfhomn.

Delta Airlines (2015), 'Delta Airlines 10K form for 2015'. https://tinyurl.com/zxzbk4g.

Dorotic, M., Verhoef, P. C., Fok, D. and Bijmolt, H. A. T. (2014*a*), '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.

Dorotic, M., Verhoef, P. C., Fok, D. and Bijmolt, T. H. (2014*b*), '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), 'Frequent-flyer miles: Funny money'. http://www.economist.com/node/5323615.

Ernst&Young (2014), 'IFRS15 revenue from contracts with customers: A closer look at the new revenue recog-

nition standard', White paper. Retrieved on March 19, 2016 from https://tinyurl.com/jcgzwrr.

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'.

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', *Accounting Review* **72**(2), pp. 231–255.

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

Kivetz, R. and Simonson, I. (2002*a*), 'Earning the right to indulge: Effort as a determinant of customer preferences toward frequency program rewards', *Journal of Marketing Research* **39**(2), 155–170.

Kivetz, R. and Simonson, I. (2002*b*), 'Self-control for the righteous: Toward a theory of precommitment to indulgence', *Journal of Consumer Research* **29**(2), 199–217.

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

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

Lai, G., Debo, L. and Nan, L. (2011), 'Channel stuffing with short-term interest in market value', *Management Science* **57**(2), 332–346.

Lai, G. and Xiao, W. (2017), 'Inventory decisions and signals of demand uncertainty to investors', *Manufacturing and Service Operations Management* . Articles in Advance, pp. 1-17.

Lai, G., Xiao, W. and Wang, J. (2012), 'Supply chain performance under market valuation: An operational approach to restore efficiency', *Management Science* **58**(10), 1933–1951.

Lambert, R. (1984), 'Income smoothing as rational equilibrium behavior', *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.

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

Liston-Heyes, C. (2002), 'Pie in the sky? real versus perceived values of air miles', *Journal of Consumer Policy* **25**(1), 1–27.

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.

Marriott (2015), 'Marriott International 10K form for 2015'. https://tinyurl.com/zxzbk4g.

Marriott (2017), '2017 Marriott rewards category/tier update'. https://tinyurl.com/gug6cep.

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, Technical report, Monash University, Department of Economics.

Oracle (2008), 'Powering the customer-centric airline with loyalty management', White paper.

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.

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

Roychowdhury, S. (2006), 'Earnings management through real activities manipulation', *Journal of Accounting and Economics* **42**, 335–370.

SAS (2012), 'Facing the challenges of building loyalty and retention: The new strategic imperative', White paper. Retrieved on March 19, 2016 from https://tinyurl.com/gperr46.

Schlappig, B. (2016), 'Marriott's brutal 2016 hotel category changes'. https://tinyurl.com/zw4zh6m.

Schmidt, W., V, G., R, L. and A, R. (2015), 'Signaling to partially informed investors in the newsvendor model', *Production Operations Management* **24**(3), 383401.

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', *Jour. of Bus. Res.* **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.

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.

Online Supplement for "Loyalty Program Liabilities and Point Values"

A Comparative Statics For Microfounded Model

Our micro-founded model in §5.3 can be used to study how specific behavioral parameters or biases affect the firm's optimal decisions. Specifically, we examine how the customers' point value perception $\tilde{\gamma}$ impacts the optimal cash price and the point value. For simplicity, we consider a uniformly distributed customer valuation $\tilde{v}$, a deterministic loyalty threshold ξ, and a linear reward function f. Based on empirical evidence, we furthermore assume that the point balance $\tilde{\omega}$ is exponentially distributed.[10]

As discussed earlier, empirical evidence suggests that the point value perceived by consumers may systematically differ from the monetary point value set by the firm, due to a variety of cognitive and behavioral effects (Liston-Heyes 2002, Basumallick et al. 2013). To study this bias, we parameterize the consumers' point value perception $\tilde{\gamma}$ as follows:

$$\tilde{\gamma} = \begin{cases} 1 - \Gamma^N & \text{with probability } 1 - z, \\ 1 + \Gamma^P & \text{with probability } z, \end{cases} \tag{13}$$

where $z \in [0, 1]$. A fraction z of customers overvalue points by an additive amount $\Gamma^P \geq 0$, and the remaining $1 - z$ undervalue points by an additive amount $\Gamma^N \in [0, 1]$. We refer to the former (latter) class of customers as *positively (negatively) biased.* For tractability purposes, we analyze the limiting case when $\Gamma^P, \Gamma^N \to 0$.

Lemma 8. *As a larger fraction of consumers are positively biased (i.e., as z increases), the optimal cash price $p^\star$ decreases, and the optimal value of points $L^\star$ increases.*

When a larger fraction of the population is positively biased (i.e., z increases), more consumers would consider purchases with the firm, since the perceived value of points more readily meets or exceeds the loyalty thresholds (i.e., $\tilde{\gamma}\theta\tilde{\omega}_i \geq \tilde{\xi}$). Thus, the result is surprising, as one might expect the firm to *raise* its cash price in response to the increased demand, and perhaps also *lower* the points' monetary value given the higher perceived value by consumers.

B Proofs

Proof of Theorem 1. The representation holds at $t = T+1$, since $J_{T+1}(w_{T+1}, p_{T+1}, \theta_{T+1}) = \mathbb{E}[f_{T+1}(\kappa_{T+1} + L_{T+1})] \stackrel{\text{def}}{=} \mathbb{E}[V_{T+1}(y_{T+1})]$. Assume this holds at time $t+1$, and consider the Bellman recursion (5) at t:

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

[10] We collected customer point balance data from an industry partner with a large established loyalty program. Our statistical analysis revealed the exponential distribution to provide an excellent fit. The industry partner's management team had expertise with other established loyalty programs, and confirmed our view that the exponential distribution would likely remain appropriate to model consumer point balances more broadly.

The last step is justified by recalling Assumption 1 and Assumption 2 (see §3). These ensure that y_{t+1} only depends on $(p_{t+1}, L_{t+1}, \tilde{\varepsilon}_{t+1})$, 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 L_{t+1} is strictly increasing in $w_{t+1}\theta_{t+1}$ (in view of Assumption 2). In particular, there exists a strictly increasing bijection $\phi_{t+1}(p_{t+1}, \cdot) : [0, \infty) \to [0, \infty)$ so that $\theta_{t+1} = \frac{\phi_{t+1}(p_{t+1}, L_{t+1})}{w_{t+1}}$, 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 (b). Part (c) readily follows. □

Proof of Theorem 2. The Bellman recursion in Theorem 1 for period $t-1$ can be rewritten as:

$$V_{t-1}(y) = \max_{L_t} \phi_t(y, L_t), \tag{14a}$$

$$\phi_t(y, L) \stackrel{\text{def}}{=} f_{t-1}\big(y - L\big) + \alpha\, G_t(L) \tag{14b}$$

$$G_t(L) \stackrel{\text{def}}{=} \max_{p_t \geq 0} \mathbb{E}_{\tilde{\varepsilon}_t}\Big[V_t\big(\kappa_t(p_t, L, \tilde{\varepsilon}_t) + L\big)\Big]. \tag{14c}$$

Since f is concave increasing and $\kappa_t(p, L, \tilde{\varepsilon}_t)$ is jointly concave in (p, L) for any $\tilde{\varepsilon}_t$ (in view of Assumption 3), a simple inductive argument can be used to show that $G_t(L)$ and $V_t(y)$ are concave, $\phi_t(y, L)$ is jointly concave, and V_t and ϕ_t are increasing in y.

To prove (a) and (b), note that ϕ_t is supermodular in (y, L) on the lattice $\mathbb{R}^2_+$, since f_{t-1} is concave. Thus, the maximizer $L^\star_t(y)$ in (14a) must be increasing in y. Furthermore, by changing variables to $x \stackrel{\text{def}}{=} y - L_t$, problem (14a) can be rewritten: $V_t(y) = \max_x \big[f_{t-1}(x) + \alpha\, G_t(y - x)\big]$. The maximand in this problem is supermodular in (x, y) on the lattice $\mathbb{R}^2_+$, since G_t is concave. Thus, $x^\star(y) = y - L^\star_t(y)$ is increasing in y. □

Proof of Theorem 3. In view of Assumption 3, the Bellman recursions at time $t-1$ can be written as:

$$V_{t-1}(y, \sigma) = \max_{L_t} \phi_t(y, L_t, \sigma), \tag{15a}$$

$$\phi_t(y, L, \sigma) = f(y - L) + \alpha\, \mathbb{E}\Big[V_t\big(\rho(L) + \sigma\tilde{\varepsilon}_t + L, \sigma\big)\Big], \text{ where} \tag{15b}$$

$$\rho(L) \stackrel{\text{def}}{=} \max_{p \geq 0} \bar{\kappa}(p, L). \tag{15c}$$

We first prove several useful intermediate results. To ease notation, we omit explicitly showing the dependency on σ here, and we use V'_t to denote $\frac{\partial V_t(y,\sigma)}{\partial y}$. Also, we omit the argument for some functions that are evaluated repeatedly at the same argument (as are their derivatives). In particular, $L^\star_t$ is repeatedly evaluated at y in the expressions below; thus $L^\star_t$ will denote $L^\star_t(y)$. Similarly, the functions f, V_t and ρ (as well as their derivatives) are evaluated at $y - L^\star_t(y)$, $\rho(L^\star_t) + L^\star_t + \sigma\tilde{\varepsilon}_t$ and $L^\star_t$ respectively. In such instances, we will similarly omit their respective argument; for instance, $f^{(2)} = f^{(2)}(y - L^\star_t)$.

Property (O). At optimality, $1 + \rho'(L^\star_t(y)) \geq 0, \quad$ for all $t = 2, \ldots, T+1$ and y.

To prove this, consider the first-order condition (FOC) yielding $L^\star_t$ in (15a). By an application of the Envelope Theorem, we have: $f'(y - L^\star_t) = \alpha\,\big(1 + \rho'(L^\star_t)\big)\,\mathbb{E}\big[V'_t\big(\rho(L^\star_t) + \sigma\tilde{\varepsilon}_t + L^\star_t\big)\big]$. Since f is strictly increasing, and V_t is increasing, we must have that $1 + \rho'(L^\star_t) \geq 0$, completing the proof.

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

We prove this by induction. Note that this holds at $T+1$ since $V_{T+1}(y) = f(y)$, and f' is convex by assumption. For $t < T$, by applying the Envelope Theorem and taking the second order derivative we obtain

$$V_{t-1}^{(3)}(y) = f^{(3)}\left(1 - L_{t,y}^{\star}\right)^2 - f^{(2)} L_{t,yy}^{\star}, \tag{16}$$

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

$$F_t(y, L) \stackrel{\text{def}}{=} -f'(y - L) + \alpha\big(1 + \rho'(L)\big)\mathbb{E}\big[V_t'(\rho(L) + L + \sigma\tilde{\varepsilon}_t)\big]. \tag{17}$$

The maximand ϕ_t in (15a) 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

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

By using this expression to substitute for $L_{t,yy}^{\star}$ in (16), we get:

$$V_{t-1}^{(3)}(y) = \frac{1}{F_{t,L}}\left[f^{(3)}\left(1 - L_{t,y}^{\star}\right)^2 F_{t,L} + f^{(2)}\Big(F_{t,yy} + (L_{t,y}^{\star})^2 F_{t,LL} + 2L_{t,y}^{\star} F_{t,yL}\Big)\right]. \tag{18}$$

We show that this is non-negative, which proves (I). To that end, note that from (17) we have:

$$F_{t,L} = f^{(2)} + \alpha\rho^{(2)}\mathbb{E}[V_t'] + \alpha(1 + \rho')^2\mathbb{E}[V_t^{(2)}] \tag{19a}$$

$$F_{t,y} = -f^{(2)} \tag{19b}$$

$$F_{t,yy} = -f^{(3)} \tag{19c}$$

$$F_{t,yL} = f^{(3)} \tag{19d}$$

$$F_{t,LL} = -f^{(3)} + \alpha\rho^{(3)}\mathbb{E}[V_t'] + 3\alpha\rho^{(2)}(1 + \rho')\mathbb{E}[V_t^{(2)}] + \alpha(1 + \rho')^3\,\mathbb{E}[V_t^{(3)}]. \tag{19e}$$

We now use these to rewrite (18). First, note that the parenthesis in the second term of (18) can be written:

$$\begin{aligned}
&F_{t,yy} + (L_{t,y}^{\star})^2 F_{t,LL} + 2L_{t,y}^{\star} F_{t,yL} \\
&= -f^{(3)} + (L_{t,y}^{\star})^2\Big(-f^{(3)} + \alpha\rho^{(3)}\mathbb{E}[V_t'] + 3\alpha\rho^{(2)}(1 + \rho')\mathbb{E}[V_t^{(2)}] + \alpha(1 + \rho')^3\,\mathbb{E}[V_t^{(3)}]\Big) + 2L_{t,y}^{\star} f^{(3)} \\
&= -f^{(3)}\left(1 - L_{t,y}^{\star}\right)^2 + (L_{t,y}^{\star})^2\Big(\alpha\rho^{(3)}\mathbb{E}[V_t'] + 3\alpha\rho^{(2)}(1 + \rho')\mathbb{E}[V_t^{(2)}] + \alpha(1 + \rho')^3\,\mathbb{E}[V_t^{(3)}]\Big).
\end{aligned}$$

Using this expression and (19a) to replace $F_{t,L}$, we can rewrite (18) as follows:

$$\begin{aligned}
V_{t-1}^{(3)}(y) = \frac{1}{F_{t,L}}\Big\{ & f^{(3)}\big(1-L_{t,y}^\star\big)^2\Big(f^{(2)}+\alpha\rho^{(2)}\mathbb{E}[V_t']+\alpha(1+\rho')^2\mathbb{E}[V_t^{(2)}]\Big) \\
& - f^{(2)}f^{(3)}\big(1-L_{t,y}^\star\big)^2 + f^{(2)}(L_{t,y}^\star)^2\Big(\alpha\rho^{(2)}\mathbb{E}[V_t']+3\alpha\rho^{(2)}(1+\rho')\mathbb{E}[V_t^{(2)}]+\alpha(1+\rho')\mathbb{E}[V_t^{(3)}]\Big)\Big\} \\
= \frac{1}{F_{t,L}}\Big\{ & \underbrace{f^{(3)}\Big[\alpha\rho^{(2)}\mathbb{E}[V_t']+\alpha(1+\rho')^2\mathbb{E}[V_t^{(2)}]\Big]}_{:=A} \\
& + \underbrace{f^{(2)}(L_{t,y}^\star)^2\Big(\alpha\rho^{(3)}\mathbb{E}[V_t']+3\alpha\rho^{(2)}(1+\rho')\mathbb{E}[V_t^{(2)}]+\alpha(1+\rho')^3\,\mathbb{E}[V_t^{(3)}]\Big)}_{:=B}\Big\}.
\end{aligned}$$

To conclude the argument, recall the following properties for the functions of interest:

$$f \text{ is concave and } f' \text{ is convex} \Rightarrow f^{(2)} \le 0, \quad f^{(3)} \ge 0$$

$$\rho \text{ is concave, } \rho' \text{ is convex, and property (O)} \Rightarrow \rho^{(2)} \le 0, \quad \rho^{(3)} \ge 0, \quad (1+\rho') \ge 0$$

$$V_t \text{ increasing and concave, and the induction hypothesis} \Rightarrow V_t' \ge 0, \quad V_t^{(2)} \le 0, \quad V_t^{(3)} \ge 0.$$

The induction and the proof for (I) follow since:

$$\Big\{f^{(2)} \le 0,\ \rho^{(2)} \le 0,\ V_t' \ge 0,\ V_t^{(2)} \le 0\Big\} \Rightarrow F_{t,L} \le 0$$

$$\Big\{f^{(3)} \ge 0,\ \rho^{(2)} \le 0,\ V_t' \ge 0,\ (1+\rho') \ge 0,\ V_t^{(2)} \le 0\Big\} \Rightarrow A \le 0$$

$$\Big\{f^{(2)} \le 0, \{\rho^{(3)} \ge 0, V_t' \ge 0\},\ \{\rho^{(2)} \le 0, (1+\rho') \ge 0, V_t^{(2)} \le 0\}, \{(1+\rho') \ge 0, V_t^{(3)} \ge 0\}\Big\} \Rightarrow B \le 0.$$

Property (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) function, then $\mathbb{E}[Xf'(X)] < 0$ (> 0).

We prove this for f concave, increasing (the argument for convex, decreasing is similar). Let h denote the probability density function of X. We have:

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

(b) Consider the simplified recursion as in (15a-15c). We have for all $t = 1, \ldots, T$, y and $\sigma \ge 0$:

$$V_t(y,\sigma) = \max_{L_{t+1}}\Big[f\big(y-L_{t+1}\big)+\alpha\mathbb{E}\big[V_{t+1}(\rho(L_{t+1})+L_{t+1}+\sigma\tilde{\varepsilon}_{t+1},\sigma)\big]\Big]. \tag{20}$$

$$\begin{aligned}
\text{We have: } \frac{\partial V_T(y,\sigma)}{\partial\sigma} &= \alpha\mathbb{E}\big[\tilde{\varepsilon}_{T+1}\,f'(\rho(L_{T+1}^\star(y,\sigma))+L_{T+1}^\star(y,\sigma)+\sigma\tilde{\varepsilon}_{T+1})\big] && [\text{by the Envelope Theorem}] \\
&< 0. && [f \text{ is concave increasing} + \text{(II)}]
\end{aligned}$$

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

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

We next prove another useful intermediate result.

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

By using the expressions above we get

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

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

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

(a) 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

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

Since the maximand of (20) 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 $\frac{\partial L^\star}{\partial \sigma} = -\big(\frac{\partial F_t}{\partial \sigma}\big|_{L^\star}\big)/\big(\frac{\partial F_t}{\partial L}\big|_{L^\star}\big)$. Thus, it suffices to show that the partial derivative of F_t with respect to σ, evaluated at $L^\star$ is non-negative:

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

Proof of Theorem 4. We find it helpful to also prove that the value function has the following form:

$$V_t(y) = y - L^\star_{t+1} + \sum_{\tau=t+1}^{T+1} \alpha^{\tau-t} \Big\{ \mathbb{E}_{\tilde{\varepsilon}_\tau} \big[\kappa_\tau(p^\star_\tau, L^\star_\tau, \tilde{\varepsilon}_\tau)\big] + (L^\star_\tau - L^\star_{\tau+1}) \Big\}, \tag{21}$$

We prove all results by induction on t. Note that (21) holds trivially for $t = T+1$. Assume it holds at time t, so that $V_t(y) = y + K_t$, where K_t is a constant. Consider the Bellman recursion at $t-1$:

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

As such, letting $(p^\star_t, L^\star_t) \in \arg\max \big\{ \alpha\, \mathbb{E}_{\tilde{\varepsilon}_t}[\kappa_t(p, L, \tilde{\varepsilon}_t)] - (1-\alpha) L \big\}$, one can see that the cash price and point value follow from (7a) and (7b), respectively, and the induction proof is completed as follows:

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

□

Proof of Theorem 5. Note that the reward function is piecewise-linear, thus differentiable almost everywhere, except for a finite number of points. All quantities of interest (e.g., V_t and $L^\star_t$) will thus inherit this property. As a result, exchanging the order of integration and differentiation of V_t is possible under suitable continuity assumptions on the distribution of $\tilde{\varepsilon}_t$. To ease exposition, we use the standard derivative 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 (15a-15c), where the dependency on σ is now replaced with a dependency on γ. The necessary and sufficient first-order optimality condition that $L^\star_t(y, \gamma)$ satisfies can be written as $F_t(L, \gamma) = f'(y - L)$, where

$$F_t(L, \gamma) \overset{\text{def}}{=} \alpha(1 + \rho'(L)) \mathbb{E}\Big[\frac{\partial V_t(y, \gamma)}{\partial y} \Big|_{(\rho(L) + L + \sigma\tilde{\varepsilon}_t, \gamma)} \Big].$$

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 $\underline{y}_t$ and $\overline{y}_t$ such that $L^\star_t(y, \gamma)$ satisfies

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

(ii) $L^\star_t(y, \gamma) = y - \hat{\Pi}$, 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

$$F_{t,\gamma}(L_t^\star(y,\gamma),\gamma) - 1 + \frac{\partial L_t^\star}{\partial \gamma} F_{t,L}(L_t^\star(y,\gamma),\gamma) = 0, \tag{22}$$

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

$$\begin{aligned} F_{t,\gamma}\big(L_t^\star(y,\gamma),\gamma\big) &= \alpha\Big(1+\rho'\big(L_t^\star(y,\gamma)\big)\Big)\mathbb{E}\Big[\frac{\partial}{\partial\gamma}\frac{\partial}{\partial y}V_t\Big(\rho\big(L_t^\star(y,\gamma)\big) + L_t^\star(y,\gamma) + \sigma\tilde{\varepsilon},\gamma\Big)\Big] \\ &= \frac{\mathbb{E}\big[\gamma\frac{\partial}{\partial\gamma}\frac{\partial}{\partial y}V_t(\rho(L_t^\star(y,\gamma)) + L_t^\star(y,\gamma) + \sigma\tilde{\varepsilon},\gamma)\big]}{\mathbb{E}\big[\frac{\partial}{\partial y}V_t(\rho(L_t^\star(y,\gamma)) + L_t^\star(y,\gamma) + \sigma\tilde{\varepsilon},\gamma)\big]} \le 1. \end{aligned} \tag{23}$$

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

$$\frac{\partial}{\partial y}V_t(y,\gamma) = f'(y - L_{t+1}^\star(y)). \tag{24}$$

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

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

Using the bounds $F_{t,L}(L_t^\star(y,\gamma),\gamma) < 0$ and $F_{t,\gamma}(L_t^\star(y,\gamma),\gamma) \le 1$, equation (22) yields that $\frac{\partial L_t^\star}{\partial\gamma} \le 0$ for case (i). For case (ii), the inequality still holds as $L_t^\star$ is constant. Thus, $\frac{\partial L_t^\star}{\partial\gamma} \le 0$ for $y \le \overline{y}_t$.

To complete the proof, note that for $y > \overline{y}_t$ and case (iii), the equivalent of equation (22) is

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

so it suffices to show $F_{t,\gamma}(L_t^\star(y,\gamma),\gamma) \ge 0$. Recall from (23) that the sign of $F_{t,\gamma}$ is given by $\alpha(1+\rho')\mathbb{E}\Big[\frac{\partial}{\partial\gamma}\frac{\partial}{\partial y}V_t\Big]$. By (24), $\frac{\partial V_t}{\partial y} = f'\big(y - L_{t+1}^\star(y)\big)$; and since f' is trivially increasing in γ and $1+\rho' \ge 0$ by property (O), the result follows.

(b) By Lemma 1(b), $\frac{\partial}{\partial y}V_t(y,\gamma)$ is decreasing in t. Thus, $F_t(L,\gamma)$ also decreases 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 1. We prove both parts together, by backwards induction. Consider the Bellman recursions in (15a)-(15c), where we omit the dependency on σ. The Envelope Theorem for (15a) yields:

$$V_t'(y) = f'(y - L_{t+1}^\star),\ \forall\, t \in \{1,\dots,T\}.$$

Since $V_{T+1}(y) = f(y)$, we readily have that $V_T'(y) = f'(y - L_{T+1}^\star) \ge V_{T+1}'(y) = f'(y)$, since f is strictly concave (so that f' is decreasing). Furthermore, we also have $L_{T+1}^\star(y) \ge L_{T+2}^\star(y) \equiv 0, \forall\, y$. 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 (15a) written at time $t-1$, yielding $L_t^\star$, and note that:

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

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 2. The proof proceeds by induction, in an analogous fashion to Lemma 1. Details are omitted for space considerations, but are available from the authors upon request. □

Proof of Lemma 3. Consider the Bellman recursions in (15a)-(15b). If $\bar{\kappa}_t(p_t, L_t)$ is supermodular (submodular) in (p_t, L_t), then the set of maximizers in problem (15c) is increasing (decreasing) in L_t. Since $L_t^\star \geq 0$, the result follows. □

Proof of Lemma 4. We argue for the case of linear reward. The proof for a concave reward function is similar. Note that $\frac{\partial^2 \mathbb{E}[\kappa_t(p_t, L_t, c)]}{\partial c \, \partial p_t} = -\frac{\partial \mathbb{E}[r_t]}{\partial p_t} \leq 0$, $\frac{\partial^2 \mathbb{E}[\kappa_t(p_t, L_t, c)]}{\partial c \, \partial L_t} = -\frac{\partial \mathbb{E}[r_t]}{\partial L_t} \leq 0$, and $\frac{\partial^2 \mathbb{E}[\kappa_t(p_t, L_t, c)]}{\partial p_t \, \partial L_t} \geq 0$ by our assumption, so that $\mathbb{E}[\kappa_t]$ is supermodular in $(p_t, L_t, -c)$, and the optimal price $p_t^\star$ and LP value $L_t^\star$ will be decreasing in c. Lastly, recall that $L_t^\star = w_t \theta_t \, g_t(p_t, w_t \, \theta_t)$ is increasing in θ_t for any fixed w_t, p_t. Thus, consider increasing c: by the argument above, this would decrease $L_t^\star$ and $p_t^\star$, which would lead to (i) a decrease in the left-hand-side of the equation, and (if g_t decreases in p_t, by our assumption) (ii) an increase in the right-hand-side of the equation. Therefore, with w_t fixed, it must be that θ_t decreases with c. □

The proofs of all remaining results are available upon request.