

Reaching for Gold: Frequent-Flyer Status Incentives and Moral Hazard *

A. Yeşim Orhun

Tong Guo

Andreas Hagemann

University of Michigan

Duke University

University of Michigan

October 25, 2021

Abstract

We study how frequent-flyer program members change their purchase behaviors as they progress towards achieving elite status. Using data from a leading U.S. airline, we empirically test the theoretical prediction that travelers' switching costs vary dynamically with their progress towards attaining status. We show evidence for increased switching costs as the consumer approaches the target pace of point accumulation required to attain status. These switching costs reflect changes in booking behavior with the airline: Travelers become more likely to choose the airline even when it is less appealing than its competitors, and to pay higher prices than they otherwise would. These responses are reduced when travelers accumulate points at a rate substantially ahead of the target pace. The increase in switching costs is more pronounced for consumers at a hub of the airline and for business travelers. Moreover, we document a stronger willingness-to-pay response when consumers are less likely to shoulder the ticket costs themselves because they are traveling for business. This response suggests that asymmetric incentives induced by business travel explains much of the heterogeneity between business and leisure travelers, and moral hazard may be responsible for a large part of the profitability of frequent-flyer status incentives.

Keywords: Loyalty Programs, Tiered Status Incentives, Moral Hazard.

*The authors thank seminar participants at the Federal Trade Commission, Georgia Institute of Technology, Harvard University, New York University, McGill University, University of Michigan, UTD FORMS conference, and Washington University at St. Louis for their helpful comments. We also thank an anonymous airline for sharing the data.

1 Introduction

Many loyalty programs across the service sector offer additional benefits to customers whose purchases with the company qualify them for an ‘elite status’. The airline industry provides a canonical example: elite status on frequent-flyer programs.¹ These programs commonly feature a deadline and a point threshold structure such that the traveler needs to accumulate a certain number of flown miles within the calendar year to attain status. Travelers with an elite frequent-flyer status enjoy perks such as increased probability of upgrades, lower baggage fees, more generous baggage allowances, priority check-in, and priority boarding every time they fly with that airline after attaining status.

Frequent-flyer programs have been hypothesized to help airlines price discriminate, create behavioral loyalty, and increase profits by inducing switching costs (for a review, see [Shugan, 2005](#)). In cases where the decision-maker does not shoulder the entire cost of travel (e.g., business trips) but enjoys the benefits, frequent-flyer programs have also been theorized to create a moral hazard problem ([Basso et al., 2009](#); [Klemperer and Png, 1986](#); [Levine, 1986](#); [Tretheway, 1989](#); [Shugan, 2005](#)). Despite a rich theoretical discussion regarding the impact of frequent-flyer programs and the central role of consumer response to reward incentives in these discussions, empirical evidence on consumer response to the incentives created by tiered status awards of frequent-flyer programs is lacking.

In this paper, we empirically examine whether and how travelers change their booking behaviors in order to achieve elite status, and whether their response differs by the likelihood that they are personally shouldering the ticket costs. We analyze the transactional database of a leading U.S.-based airline’s frequent-flyer program, including the histories and point accumulations of 3.5 million frequent-flyer program members during the 2010 and 2011 point-earning cycles. Our results show that consumers become more willing to accept higher ticket prices and/or to accept less appealing flight options as they make progress towards their status goal. We find that the increase in the willingness to sacrifice current utility for the possibility of future rewards is stronger for consumers whose home airport is a hub of the airline, and for business travelers. We further show that when travelers are less likely to be shouldering the ticket costs personally, they are more likely to respond to these incentives by increasing their willingness to pay. This asymmetry in incentives explains much of the heterogeneity across business and leisure travelers, and is consistent with the previously hypothesized role of asymmetric incentives created by business travel in frequent-flyer programs.

¹Frequent-flyer programs offer both tier-based and frequency-based rewards. Approximately half of the travelers surveyed by a 2010 Gallup Survey belonged to at least one frequent-flyer program. This paper focuses on the tier-based rewards aspect.

To motivate our empirical approach, Section 3 explains the structure of the status rewards and presents a stylized dynamic optimization model to establish hypotheses regarding how consumers' optimal booking policy may vary with two key state variables: the consumer's point accumulation to date and the time left to acquire additional points. A statistic that captures both these state variables, and one that we employ in our empirical analyses, is the consumer's point accumulation pace compared to the target pace that results in the necessary total points by the end of the year. The model illustrates that the amount of immediate consumption utility the traveler is willing to sacrifice is low when she is considerably behind on her point accumulation compared to the target pace. Her willingness to sacrifice current utility increases as her point accumulation pace gets closer to the target pace needed to attain status and may decline again when she is considerably ahead of schedule in point accumulation. These changes in willingness to sacrifice current utility for future benefits translate to changes in willingness to pay or willingness to book with the airline even when it does not offer the most attractive flight options. We formally discuss the identification strategy that allows us to identify the existence and impact of these dynamic incentives from systematic changes in consumers' booking behaviors in Section 3.3.

Section 4 details our data which come from the bookings database of a leading U.S. airline and the Airline Origin and Destination Survey Databank 1B (DB1B).² As with any other company database, our data are rich at the consumer level but do not include information on whether the member traveled with another airline, took a train, or did not travel when the member did not purchase from the airline. We discuss the identification assumptions that allow us to causally infer changes in consumer switching costs from changes in consumers' booking patterns with one airline. First, we assume that product offerings and prices on a given route, at a given time, do not respond to an individual consumer's purchase history and point accumulation with the airline. Second, we assume that the travel opportunities arise exogenously. Third, we assume that after controlling for weekly seasonality, variations in a traveler's preferences over attributes of air travel (e.g., price, convenience, service quality) are independent of changes in her progress towards status attainment. Thus, we use the co-movement of systematic changes in the selection of booking attributes that appear in the company database and the changes in the individual-level state variables to infer how travelers' preferences for the airline versus other travel options change as they make progress towards attaining elite status. The richness of the member panel allows us to control for common variations over time in travel preferences and flight options, as well as individual preference heterogeneity that may otherwise produce a

²The DB1B database provides a 10% random sample of airline passenger tickets on domestic routes. The data are reported at the quarterly level and allow us to calculate market shares of each airline across domestic routes.

spurious correlation between point accumulation and purchase behaviors. This approach provides a solution to the challenge of identifying changes in consumer valuations with detailed transaction data from only one company in a competitive landscape, a common issue in customer relationship management research.

Our evidence for how travelers' willingness to forgo current utility responds to the pace of the progress they make towards attaining status comes from two specifications, as detailed in Section 5. First, we examine changes in travelers' willingness to book with the airline when it offers less appealing flight options relative to its competition. We find that travelers are more likely to book with the airline on routes where the airline has lower appeal than its competitors when they are close to the target pace. Second, we examine changes in prices that travelers pay relative to the prices others pay on the same flight, effectively controlling for all product features such as the route, timing, directness, and appeal of the flight. We find that loyalty program members on average show an 8% increase in price paid compared to others taking the same flight when they are close to the target pace. In both specifications, we document non-monotonic responses to progress. If travelers' point accumulation over time is significantly behind or significantly ahead of the pace required to attain status, they are less likely to forgo current utility to fly with the target airline. In addition, members' increased willingness to forgo utility to fly with the airline resets when they transition from the end of one point-accumulation cycle to the beginning of another. These non-monotonic empirical patterns are consistent with the predictions of the optimal response to tiered incentives, but are inconsistent with predictions of alternative explanations based on spurious serial correlation, learning, and/or habit formation from cumulative experience with the airline. In sum, our results suggest that tiered status rewards are successful in creating significant incentives for travelers to forgo current utility in favor of booking with the airline.

These findings reflect average responses to status incentives. We examine heterogeneity in consumer response along two different dimensions. First, we explore whether the home airport of the consumer being a hub of the airline changes consumer response to status incentives. Consumers at a hub of the airline are more able to substitute into the airline's flights as their switching costs increase because the airline offers service across a wider selection of routes at a hub. Relatedly, these consumers may have more opportunities to enjoy the perks of attaining status. Indeed, we find that consumers at the airline's hub have a stronger response to status incentives. Second, we examine whether leisure and business travelers differ in the extent to which they modify their purchase behaviors as they progress towards status. The airline industry widely recognizes that business travelers are less price sensitive and travel more frequently on average than leisure

travelers. Therefore, we may expect differential responses to status incentives across the two segments. Indeed, we find that business travelers increase their willingness to pay more aggressively as they make progress towards status.

We then examine how switching costs vary when travelers are less likely to be shouldering the ticket costs themselves. Noting that individuals purchasing the ticket are less likely to shoulder the expenses related to business travel than leisure travel, we examine differences in the extent to which travelers react to status incentives when they are likely to be traveling for business versus leisure. In line with theoretical predictions, results reported in Section [5.2.3](#) show that how much a member’s willingness to pay increases as she makes progress towards status depends on the purpose of the trip. In particular, we see a larger increase in a given member’s willingness to pay when she travels for business than when she travels for leisure. This asymmetry in incentives explains much of the heterogeneity across business and leisure travelers in the price response. Moreover, more than one third of the overall increase in the average consumer’s willingness to pay in the marketplace can be attributed to booking instances where the consumer is not likely to be shouldering the cost of the ticket herself. If travelers had to pay out of pocket, our estimates suggest that companies would save at least 7% of their travel costs.

In sum, our paper makes three main contributions. First, it provides evidence for the degree of sacrifice consumers are willing to make to travel with an airline as they make progress towards attaining status in its loyalty program, and how consumer responses vary across the leisure and business traveler segments. Second, it contributes to the debate about the extent to which asymmetric incentives induced by business travel are responsible for the success of frequent-flyer programs. Third, it addresses the challenge of identifying consumer valuation changes when the database covers consumers’ transactions with only one company in a competitive landscape. We propose a strategy to infer changes in consumer valuations from the changes in booking characteristics that get differentially selected into the database, exploiting the fact that the supply of competitive options does not respond to an individual’s point accumulation. We hope that this approach proves useful in assessing the impact of other customer relationship or loyalty programs.

2 Related Literature

Both tier-based and frequency-based loyalty programs are widespread practices in the service industry. Frequency-based programs take on the form of “collect X points, get a reward,” and tier-based programs take on the form of “collect X points, qualify for a membership tier.” Marketing literature has devoted

considerable attention to measuring the impact of frequency-based reward programs on several indicators of loyalty, such as increased customer retention (e.g., Verhoef, 2003; Lewis, 2004; Gopalakrishnan et al., 2021), purchase frequency (e.g., Hartmann and Viard, 2008; Gopalakrishnan et al., 2021; Kopalle et al., 2012; Rossi, 2017; Zhang and Breugelmans, 2012), reward redemption (e.g., Lal and Bell, 2003), customer traffic (e.g., Dreze and Hoch, 1998), customer expenditures (e.g., Dreze and Hoch, 1998; Lal and Bell, 2003; Leenheer et al., 2007), and attitudinal measures (e.g., Bolton et al., 2000).³

In contrast to the rich empirical literature on frequency-based reward programs, empirical investigations of how consumers change their behaviors in response to tiered loyalty programs are sparse. Status awards create particular dynamic incentives due to their tier-based nature and deadlines. These dynamic incentives are not unlike those created by tiered bonuses in salesforce compensation structures. Therefore our paper is related to the empirical literature studying the impact of salesforce compensation schemes. We find that travelers become more willing to forgo current utility as their progress is increasing, but scale back their efforts if their point accumulation becomes considerably faster than the pace required to attain their goal. Copeland and Monnet (2009), Misra and Nair (2011) and Chung et al. (2013) predict and/or document similar strategic (and nonlinear) behavior among salespeople choosing optimal effort provision when facing tier-based incentives.

Two papers aim to empirically address related questions in tiered loyalty reward structures. Kopalle et al. (2012) show that loyalty program members' propensity to stay at a hotel chain increases as they progress towards the requirement threshold for attaining elite status with the sponsor hotel. In the context of the airline industry, Dreze and Nunes (2011) argue that travelers are more likely to fly with the airline after achieving the highest status tier. Our paper adds to this body of work by documenting how travelers' willingness to pay and the characteristics of the trips they are willing to take with the airline change as they progress towards attaining status with an airline, and departs from the focus on purchase frequencies. This departure is driven by two factors. First, we are interested in responses to prices and characteristics, as they allow us to study the trade-offs members become more likely to make. Second, the assumption that enables inference based on changes in booking frequencies is unlikely to hold in our setting. In order to identify the underlying changes in a member's valuations by changes in the *frequency* of a member's bookings with the airline, we would have to assume that a member's probability of travel is independent of her past (cumulative) travel in a given year. Such an assumption is likely to be violated in air travel, because most travelers have

³Breugelmans et al. (2015) offer a comprehensive discussion of the existing research on loyalty programs.

a time and/or monetary budget for travel: If a member travels very frequently for a period of time, she may be less likely to travel as frequently in subsequent periods. Therefore, our inference strategy, which relates changes in the *characteristics* of bookings a member makes with the airline to changes in her valuation, does not rely on this assumption. Instead, it relies on the assumption that a member’s preferences regarding travel attributes (e.g., price, convenience, quality of service) are independent of her past (cumulative) travel in a given year. This inference strategy provides an alternative approach for identifying the causal impact of making progress in a loyalty program on consumer valuations.

Our empirical approach also addresses other inference problems highlighted by previous investigations of loyalty programs. [Gopalakrishnan et al. \(2021\)](#) underscore the importance of controlling for individual preference heterogeneity in order to identify causal responses to a frequency-based loyalty program. The ability to do so is also crucial in order to avoid the potential selection bias arising from systematically different purchasing patterns across members whose past travel accumulations differ. The panel nature of our data allows us to conduct within-member analyses to control for preference heterogeneity in a flexible manner. In the context of the hotel industry, [Kopalle et al. \(2012\)](#) discuss the importance of seasonality in price evolution and allow for quarterly variations. In our empirical approach, we have the flexibility to include weekly controls that account for common demand and supply variations. Controlling for these common unobservables is important to avoid potential confounds arising from structural changes in price distributions and/or willingness to travel with the airline that are correlated with time and therefore with overall point accumulation.

Another goal of our paper is to provide empirical evidence for the role of asymmetric incentives induced by business travel in the context of loyalty programs. Previous work highlights that such asymmetry may create a type of *moral hazard* that airlines may be taking advantage of ([Basso et al., 2009](#); [Klemperer and Png, 1986](#); [Levine, 1986](#); [Trethewey, 1989](#); [Shugan, 2005](#)). [Shugan \(2005, p.189\)](#) summarizes this view: “These situations can induce potential agency problems when so-called loyalty programs target the decision rather than the payer. These programs might provide direct de facto side payments (akin to the original concept of “kickbacks”) to decision makers.” Of course, this concern is not specific to the airline industry. Asymmetric incentives could be present in other purchase decisions that employees make and employers pay for, such as hotels, car rentals, office equipment, etc. Speaking to this point, [Basso et al. \(2009, p.117\)](#) state “we think it is no coincidence that these very large loyalty programs exist in areas in which a large fraction of purchases are work related.” These authors call for empirical evidence on the extent to which moral hazard

induced by asymmetric incentives inherent to business travel is a driver of behavior in loyalty programs.

In this regard, Rossi (2017) is closest to our work. In the context of gasoline purchases, Rossi (2017) provides suggestive evidence of moral hazard by documenting a group of travelers who extract more value from one dollar's worth of rewards than from one dollar spent for gasoline. While this finding is suggestive of moral hazard, Rossi points out the importance of knowing whether the segment of consumers with such valuations are business travelers, which his empirical context does not allow. In our empirical context, we can proxy for the purpose of the trip, and characterize member heterogeneity based on members' propensity to travel for business. More importantly, we can identify how a member's purchase behaviors vary across business and leisure trips, and how this difference evolves as the member makes progress towards attaining status. Therefore, we can identify the impact of asymmetric incentives created by business travel on members' responses separately from preference heterogeneity.

We also contribute to the literature on frequent-flyer programs. A large body of theoretical work debates the impact of frequent-flyer programs on airline competition and profitability. For example, Kim et al. (2001) and Klemperer (1987, 1995) argue that loyalty programs give airlines more market power. In contrast, Caminal and Claiici (2007), Caminal and Matutes (1990), and Kopalle and Neslin (2003) suggest that loyalty programs increase competition and erode airline profits. These conclusions rest on assumptions regarding the impact of frequent-flyer programs on consumers' switching costs, yet empirical investigations of this impact have been sparse. Lederman (2007, 2008) examines market price reactions to airline alliance expansions to infer that the value consumers derive from frequent-flyer programs increases in the program's network due to the increased ability to earn and redeem frequency-based rewards. Two survey-based studies report that frequent-flyer program membership influences respondents' airline choice (Prousaloglou and Koppelman, 1995) and willingness to pay (Hess et al., 2007). We add to this literature by examining individuals' actual booking behaviors to infer how consumers' switching costs change as they make progress towards attaining elite status.

3 Industry Background and a Dynamic Decision Model with Status Incentives

3.1 Background: Frequent Flyer Programs and Status Rewards

The U.S. airline industry accounts for more than 10 million American jobs and 5 cents of every dollar of U.S. GDP. Currently, the airlines with the largest market share in the U.S. domestic markets are American (18.2%), Southwest (18.1%), Delta (16.8%), United (14.9%), JetBlue (5.5%), Alaska (5.1%) and Spirit Airlines (3.5%)^[4] In 1981, American Airlines launched what many consider to be the first airline loyalty program. It was quickly followed by Delta Airlines, United Airlines, and Continental Airlines. Twenty-five years after this launch, the world’s frequent-flyer programs boasted more than 180 million members, 120 million of whom were U.S. residents.

Frequent-flyer programs offer two simultaneous reward systems: frequency rewards (reward miles) and tiered rewards (elite status). Members earn reward miles or points either by flying or by spending with a program-affiliated company (e.g., a credit card, hotel, etc.) that has partnered with an airline. The reward miles can be redeemed for airline tickets, seat upgrades, purchases with affiliated programs and companies, etc. By 2005, a total of 14 trillion un-redeemed frequent-flyer points had been accumulated by travelers worldwide.^[5] Airlines periodically devalue reward miles by increasing the number of required miles for free tickets and other rewards, and by making it harder to redeem them.^[6] Perhaps because of the growing accumulation of un-redeemed miles in members’ accounts, frequent travelers have become more interested in achieving status. For example, Forbes argues that “[o]riginally, the goal was to accrue free travel, but that has changed dramatically in today’s aviation landscape, and most frequent-flyers these days are more interested in status than the occasional free ticket” (Larry Olmsted, Forbes, January 23, 2013). In the hotel industry, which offers similar loyalty programs, [Kopalle et al. \(2012\)](#) also find that a large number of the loyalty program members are motivated mainly by achieving status rather than by collecting reward points.

Our study focuses on this tiered elite status aspect of frequent-flyer programs.

⁴Statistics retrieved on September 1, 2018 from [Statista](#)

⁵These statistics are obtained from [webflyer.com](#). According to The Economist (<https://www.economist.com/leaders/2005/01/06/in-terminal-decline>), airlines sold frequent-flyer points to credit-card firms at an average of just under 2 cents a mile in 2005. Their value when used to buy a ticket or to upgrade to business class ranged between 1 cent and over 10 cents per mile. The Economist used the mid-point of this range to calculate the global stock of frequent-flyer miles outstanding in 2005 to be worth \$700 billion U.S. dollars. Outstanding redeemable miles are so large that, by another calculation offered by [Bhaskara \(2015\)](#), Delta’s aggregate liability amounted to \$3.9 billion, United’s to \$4.9 billion, and American’s to \$2.6 billion.

⁶For a recent discussion, please see this recent article: <http://www.latimes.com/business/la-fi-frequent-flier-programs-20170914-story.html>

The frequent-flyer programs of all legacy U.S. airlines feature a tiered status structure. Each calendar year, members must accumulate a required amount of travel between January 1 and December 31 to achieve status. If a member attains the required points before the end of the year, she can start enjoying status benefits at the moment of status attainment and continue to do so for the following calendar year. The three leading legacy airlines, American, United and Delta Airlines, all offer four tiers, with the first three tiers having the same status-qualifying point requirements of 25,000, 50,000, and 75,000 miles. The entry-level status tier (tier 1, from here on) is the Gold tier in American Airlines and the Silver tier in Delta and United Airlines.

Our analyses focus on the impact of progress towards tier 1 on consumer choices. We make this choice for two main reasons. First, tier 1 is the only relevant goal for a large majority of travelers. Second, rules governing tier 1 help us identify consumer valuation dynamics in response to progress towards attaining status. In particular, points attained during one cycle do not roll over to the next cycle if the member fails to achieve tier 1.⁷ The industry defines point-earning cycles based on calendar year and the timing of tier attainment. Given our focus on tier 1, the point-earning cycles can be defined as follows: At the beginning of 2010, travelers in our data are at the beginning of the 2010 cycle. If they fail to achieve tier 1 in 2010, they lose their accumulated points at the end of 2010 and begin a new point-earning cycle on January 1, 2011. If they attain tier 1 in 2010, they begin accumulating points for the 2011 point-earning cycle immediately upon obtaining tier 1 status, because additional points over 25,000 roll over to the next year's account.

Next, we present a dynamic consumer decision model that incorporates the main structural elements of the tiered status rewards. Focusing on the all-or-nothing structure of tier 1, the model features a fixed number of time periods to accumulate the necessary points, and number of additional periods in which the benefits can be enjoyed if status is attained. Point accumulation is based on traveling with the airline, reflecting the rules in place during the data period we study.⁸ Finally, consumers are assumed to purchase from the airline on a route that maximizes their utility.

⁷This all-or-nothing nature of tier 1 means that failing to achieve tier 1 is associated with the largest opportunity costs across all tiers. The fact that for higher tiers travelers can hold onto their points if they fail to achieve these tiers creates additional dynamics that interfere with our research question. For example, a traveler who expects not to travel much in the following year can purposefully stop traveling with the airline in focus (TA, from here on, to protect the identity of the airline) short of achieving tier 3 and roll over her miles to the next cycle, because she may value achieving tier 2 in both years more than achieving tier 3 the first year and falling back to being without status the following year. Finally, airlines differ in their rules governing higher tiers, so we cannot easily discuss these tiers without revealing TA's identity.

⁸At the time of our data, TA rewarded status-qualifying points based only on miles flown with TA. In particular, travelers were not expected to meet spending threshold and could not accumulate status-qualifying points based on purchases with affiliated credit cards, hotels, etc., although they could earn reward miles with such purchases. At different points in time, airlines also allowed limited ways in which travelers could earn tier-qualifying points through other types of purchases. In addition, Delta added revenue-based elite status requirements in January 2014, United in March 2015, and American Airlines in August, 2016. Our data come from 2010-2011, years before these changes were announced. In Section [6](#) we discuss how the new spending thresholds likely impact consumers, airlines and employers in light of our findings.

3.2 Dynamic Consumer Decision Model

3.2.1 Consumer Utility

The consumer derives immediate consumption utility from flying with airline j on route r at time t . Let the deterministic part of this utility be $\bar{U}_{rtj} = X_{rj}\beta - \gamma p_{rtj}$. Here the consumption utility depends on airline-route specific attributes X_{rj} , which may include how convenient the schedule of the airline is, how many direct flights the airline offers on route r , how new the aircraft is, how good the airline's service is, etc. Consumption utility also depends on price p_{rtj} of airline j .

Because our data is from the target airline, we focus on drivers of the booking decisions with the target airline alone by reducing the competitive options to a sufficient statistic – namely, the utility maximizing competitive option. Let j' denote the target airline, and j^* denote the airline with the highest \bar{U}_{rtj} for the purchase occasion (r, t) . We can define the expected attractiveness gap of the target airline with respect to the consumer's best option as

$$\Delta_{rt} = \bar{U}_{rtj'} - \bar{U}_{rtj^*} = (X_{rj'} - X_{rj^*})\beta - \gamma(p_{rtj'} - p_{rtj^*}).$$

This gap is zero if the target airline offers the best consumption utility and is negative if there is a better option. The immediate consumption utility consumer i would be forgoing by booking the target airline instead of her best option is

$$\Delta_{irt} = \Delta_{rt} + \varepsilon_{irt},$$

where the idiosyncratic shock ε_{irt} is median-zero, identically and independently distributed (iid) over booking occasions (r, t) and consumers i , and additively separable from the average attractiveness gap Δ_{rt} .

3.2.2 Dynamically Optimal Decisions with Tiered Status Incentives

In this section, we build intuition for the optimal consumer behavior under tiered incentives by specifying the consumer's decision problem and proving that a consumer's willingness to book with the airline first increases then declines with her pace. In what follows, we suppress subscripts i and r for simplicity.

The environment of the model closely mirrors the industry details. The time horizon is two calendar years, with time periods $t \in \{1, 2, \dots, T-1, T, T+1, T+2, \dots, T+12\}$ with a point accumulation deadline at the end of the 12th month, denoted with T . In each period t , the traveler makes a binary decision $b_t = \{1, 0\}$

whether to book TA or not for a travel occasion.⁹ Each time the traveler books with TA, she earns k_t points. The point gain k_t is assumed to be iid over t . The traveler attains status if her total cumulative points at T meet or exceed threshold M . If she attains status, she enjoys a status value π each period she flies with TA, starting upon attaining status and into the next year.

The variable K_t denotes the number of points the consumer has at the beginning of period t . The traveler starts the year with zero points, i.e. $K_1 = 0$.¹⁰ The total number of points at time t evolves mechanically until T as

$$K_t = \begin{cases} K_{t-1} + k_t & \text{if } b_{t-1} = 1 \\ K_{t-1} & \text{if } b_{t-1} = 0. \end{cases} \quad (1)$$

The timing of information revelation is as follows. At the beginning of each period t , the traveler learns her travel need realization among the possible routes. This determines the consumption utility differential Δ_t as well as the k_t for period t 's decision problem. The traveler has rational expectations over k_t and future realizations of Δ_t , which are assumed to be iid over t . The traveler is also aware of the amount of time left before the calendar year $T - t$ and the total accumulation of miles to date K_t . Therefore, she can assess how well she is doing with respect to attaining status by the end of the year. By construction, the chances that $K_{T+1} \geq M$ weakly decreases in t for any given K_t and weakly increases in K_t for any given t .

If the consumer attains status in period $t \leq T$ (such that $K_t \geq M$ and $K_T \geq M$), additional utility π is enjoyed each period t and after, including the entire second calendar year. However, the consumer does not get to enjoy π in the second calendar year if M is not reached before the deadline T (such that $T < t \leq T + 12$ and $K_t < M$). This difference is expressed in the consumer's discounted, time-separable objective function

$$G_t(K_t, \Delta_t, k_t) = \begin{cases} b_t(\Delta_t + 1(K_t \geq M)\pi) + \rho E[V_{t+1}(K_t + b_t k_t, \Delta_{t+1}, k_{t+1})] & \text{if } t \leq T \\ b_t(\Delta_t + 1(K_T \geq M)\pi) + \rho E[V_{t+1}(K_t + b_t k_t, \Delta_{t+1}, k_{t+1})] & \text{if } T < t \leq T + 12, \end{cases} \quad (2)$$

where ρ is the per-period discount factor and expectations are conditional on information at time t ,

$$E[V_{t+1}(K_t + b_t k_t, \Delta_{t+1}, k_{t+1})] = \sum_n P(k_{t+1} = n) \int_{-\infty}^{\infty} V_{t+1}(K_t + b_t k_t, \Delta_{t+1}, n) p(\Delta_{t+1}) d\Delta.$$

The decision to book with the target airline reflects whether Δ_t is worth forgoing in exchange for in-

⁹Travel occasions arise exogenously. Extending the model to endogenous trip creation is beyond the focus of this paper due to our data not including competitive bookings.

¹⁰Having the initial points be a random variable does not change the conclusions from the model.

creasing the traveler's chances of attaining status. We can write the traveler's problem recursively as a T period stochastic dynamic problem where her value function $V_t(K_t, \Delta_t, k_t)$ and her optimal policy function $b_t(K_t, \Delta_t, k_t)$ are:

$$V_t(K_t, \Delta_t, k_t) = \max_{b_t \in \{0,1\}} G_t(K_t, \Delta_t, k_t) \quad (3)$$

$$b_t(K_t, \Delta_t, k_t) = \arg \max_{b_t \in \{0,1\}} G_t(K_t, \Delta_t, k_t) \quad (4)$$

Examining equation (2), we can see that when the consumer is deciding before T , the consumer books with TA when $\Delta_t > -\lambda_t(K_t, k_t)$, where

$$\lambda_t(K_t, k_t) = \rho(E[V_{t+1}(K_t + k_t, \Delta_{t+1}, k_{t+1})] - E[V_{t+1}(K_t, \Delta_{t+1}, k_{t+1})]). \quad (5)$$

We can think of $\lambda_t(K_t, k_t)$ as the *switching cost* induced by dynamic incentives arising from status rewards. It captures the amount of current utility the consumer is willing to forgo in the current period in exchange for the possibility of attaining status by the end of the year. The optimal policy function can be characterized by $\lambda_t(K_t, k_t)$ as it defines the threshold of utility differential the consumer is willing to take under her optimal booking policy.

Proposition 1. $\lambda_t(K_t, k_t)$ is always non-negative.

Proof. It suffices to prove that $\kappa \mapsto E[V_t(K_t + \kappa, \Delta_t, k_t)]$ is weakly increasing in κ . The proof is by induction. In period $T + 12$, $\kappa \mapsto E[V_t(K_t + \kappa, \Delta_t, k_t)]$ is weakly increasing by construction. For the induction step with $1 < t < T + 12$, suppose first that $E[V_{t+1}(K_{t+1} + \kappa, \Delta_{t+1}, k_{t+1})]$ is weakly increasing in κ . It then follows from equation (3) that $\kappa \mapsto V_t(K_t + \kappa, \Delta_t, k_t)$ must be increasing as well. Monotonicity of expectations implies $\kappa \mapsto E[V_t(K_t + \kappa, \Delta_t, k_t)]$ is weakly increasing in κ . \square

Furthermore, examining the structure of $\lambda_t(K_t, k_t)$ in equation (5), it is immediately apparent that switching costs are by definition zero when the continuation values do not vary with the current period's decision. Expected future valuation is independent of whether the consumer buys in the current period under three different scenarios: (1) when the consumer has already attained status ($K_t \geq M$), (2) if the consumer has failed to reach status by the end of the year ($T < t \leq T + 12$ and $K_t < M$), or (3) when the probability of attaining status is zero because the number of periods left in the first calendar year are too few to accumulate the necessary points to attain status. The third condition depends on the specifics of the

k_t distribution and M , and will become clearer in the expression of the expected valuation in period T . We highlight the first two conditions as we rewrite the optimal policy for clarity as follows:

$$b_t(K_t, \Delta_t, k_t) = \begin{cases} 1 & \text{if } \Delta_t + \pi > 0, \text{ when } [K_t \geq M \text{ and } K_T \geq M] \\ 1 & \text{if } \Delta_t > 0, \text{ when } [T < t \leq T + 12 \text{ and } K_t < M] \\ 1 & \text{if } \Delta_t + \rho(E[V_{t+1}(K_t + k_t, \Delta_{t+1}, k_{t+1})] - E[V_{t+1}(K_t \Delta_{t+1}, k_{t+1})]) > 0, \\ & \text{when } [t \leq T \text{ and } K_t < M] \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

Given the optimal policy, we can take expectations over $V_t(K_t, \Delta_t, k_t)$ in equation (3) to describe $E[V_t(K_t, \Delta_t, k_t)]$ when $t \leq T$ as

$$\begin{aligned} E[V_t(K_t, \Delta_t, k_t)] &= \sum_n P(k_t = n) \int_{-1(K_t \geq M)\pi - \lambda_t(K_t, k_t)}^{\infty} \left(\Delta_t + 1(K_t \geq M)\pi + \rho E[V_{t+1}(K_t + k_t, \Delta_{t+1}, k_{t+1})] \right) p(\Delta_t) d\Delta \\ &\quad + \int_{-\infty}^{-1(K_t \geq M)\pi - \lambda_t(K_t, k_t)} \rho E[V_{t+1}(K_t, \Delta_{t+1}, k_{t+1})] p(\Delta_t) d\Delta. \end{aligned}$$

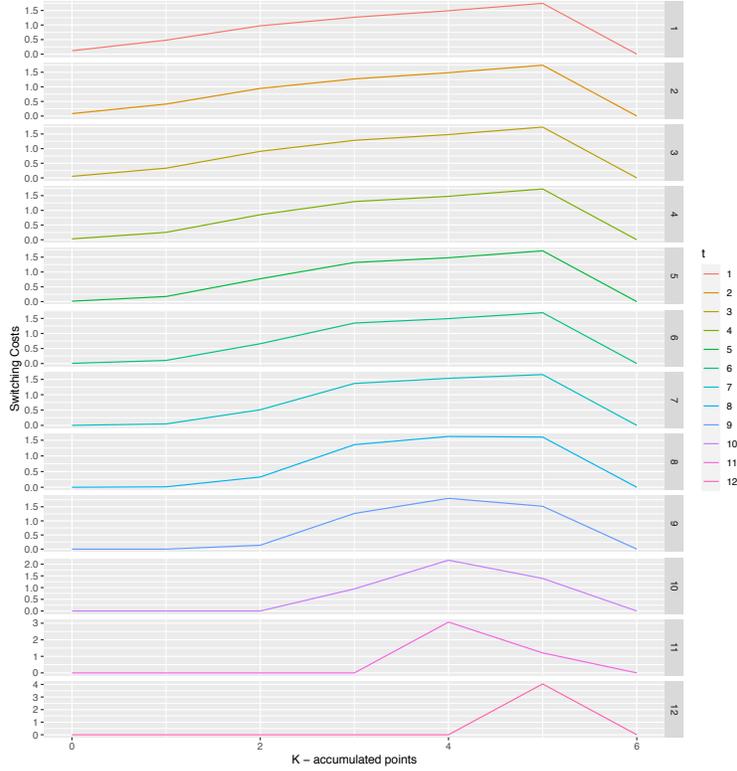
(Please also see Appendix A.2 for more intuition and how this expression simplifies for different regions of K_t and t .)

To investigate how switching costs under the optimal purchasing policy change with K_t and t , we solve the consumer's problem by backwards induction from period T , where the expected value is

$$E[V_T(K_T, \Delta_T, k_T)] = \begin{cases} \frac{\rho^{13}-1}{\rho-1} \int_{-\pi}^{\infty} (\Delta_t + \pi) p(\Delta_t) d\Delta & \text{if } K_T \geq M \\ \frac{\rho^{13}-1}{\rho-1} \int_0^{\infty} \Delta_t p(\Delta_t) d\Delta & \text{if } K_T < M - k_T \\ \int_{-\lambda'}^{\infty} (\Delta_t + E[V_{T+1}(M, \Delta_{T+1}, 0)]) p(\Delta_t) d\Delta \\ \quad + \int_{-\infty}^{-\lambda'} E[V_{T+1}(0, \Delta_{T+1}, 0)] p(\Delta_t) d\Delta & \text{if } M > K_t \geq M - k_T \end{cases} \quad (7)$$

and

$$\lambda' = \rho(E[V_{T+1}(M, \Delta_{T+1}, 0) - E[V_{T+1}(0, \Delta_{T+1}, 0)]) = \rho \left(\frac{\rho^{12}-1}{\rho-1} \right) \left(\int_{-\pi}^{\infty} (\Delta_t + \pi) p(\Delta_t) d\Delta - \int_0^{\infty} \Delta_t p(\Delta_t) d\Delta \right)$$



Notes: In our simulations, we draw the utility differential realizations from a normal distribution, $\Delta_t \sim N(\bar{\Delta}, \sigma^2)$. For this simulation, we set $\sigma = 2$, $\bar{\Delta} = -2.5$. The consumer has the chance to earn $k_t = 1$ in each period if she buys. The discount factor $\rho = 0.95$ and the additional utility the consumer gets from flying with status is $\pi = 2$.

Figure 1: Switching costs vary over accumulated points and time

is the difference between the continuation values of reaching M at T and not reaching it. Note that λ' is the same for all values of k_T satisfying $M > K_t \geq M - k_T$, since all such k_T can get the consumer over the limit M . Therefore, in this range the consumer trades off whether the current utility sacrifice is worth the utility differential of achieving status versus not achieving it. Future expectation depends only on whether the consumer reaches M or not by the end of period T , because flights taken after T do not contribute points towards status attainment.

In Figure 1, we plot how λ_t varies with K_t .¹¹ The switching cost $\lambda_t(K_t, k_t)$ characterizes the optimal policy at each (t, K_t, k_t) . For this illustration, we consider a simple scenario where the consumer earns 1 point each time she flies with TA and needs to accumulate $M = 6$ points to attain status. We see that the switching costs first increase, then decrease as the consumer accumulates more points. This result is driven by the change in the marginal impact of an additional point on the probability of achieving status by the end of the year. When the consumer has very few points compared to the total amount she needs to accumulate,

¹¹We also include a plot of how switching costs vary with time for a fixed K_t in Appendix A.1

the chances of attaining status are low. As the proportion of points increases, so do her chances. When she has a lot of points, the chances of being able to accumulate points before the deadline is so high that it is no longer impacted much by whether the consumer buys today or waits, which causes the switching costs to decline.

Notice, however, that the inflection point depends on the time period. Intuitively, the degree to which point accumulation impacts the probability of achieving status depends on *when* those points are accumulated. A consumer who has accumulated half of the necessary points by the period before the deadline is unlikely to reach status in the remaining time. On the other hand, a consumer who has the same number of points earlier on has a decent chance of attaining status. Because a forward-looking consumer responds to her chances of attaining status, the timing of point accumulation matters for the optimal policy given K_t .

We therefore construct a measure of progress that directly takes timing into account. This measure compares the consumer’s average rate of earning up to t with a constant rate of point accumulation that would achieve M by the end of the calendar year at T . We define

$$progress_t = \frac{K_t - (M/T \cdot t)}{(M/T) \cdot t}.$$

When the consumer starts with zero accumulated points, she has a progress value of -1 . If she accumulates the exact number of points required for a steady speed of progress towards reaching M by the end of the year at T , the progress measure is zero. If the member is “behind” schedule, the progress measure is negative. If the member is “ahead,” the progress measure is positive. Therefore, a member’s progress metric can be low even if she has a large amount of accumulated points but there is little time left to the end of the year. Similarly, it can be high if the member accumulated a modest amount of points early in the year.

Figure 2 plots the progress measure and the accumulated points for a particular consumer in our data who achieved status. In our empirical application, consumers need to fly 25,000 miles within the calendar year, which amounts to (approximately) 480 points per week. As demonstrated in Figure 2, progress is positive when the points earned by the person accelerate beyond (approximately) 480 points per week, which is tracked by the 45-degree line. Progress is negative (positive) when the point accumulation falls below (rises above) the 45-degree line. This figure demonstrates the non-monotonic changes in progress over time, while point accumulation increases monotonically. This is a desirable feature. While point accumulation is necessarily increasing as the traveler books more trips within the calendar year, the level of progress depends

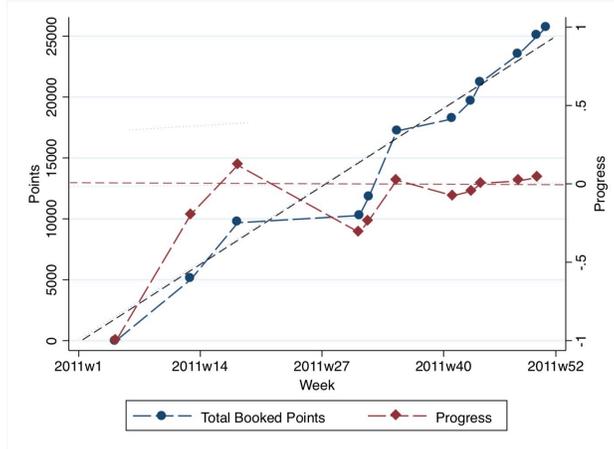


Figure 2: Point earnings and progress of a traveler.

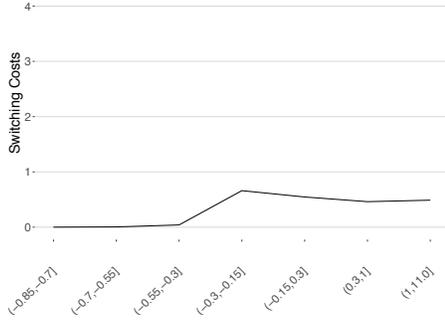
on whether the member is keeping up the pace towards the goal.

The optimal booking policy varies systematically as a function of progress. Figure 3 plots how λ_t on average varies with $progress_t$ at different parameter combinations. Overall, it is easy to see that switching costs first increase with progress as the probability of attaining status increase. In some cases, after progress is large enough such that the chance of status attainment is high and the additional point accumulation does not substantially impact this chance, switching costs can also decline in progress.

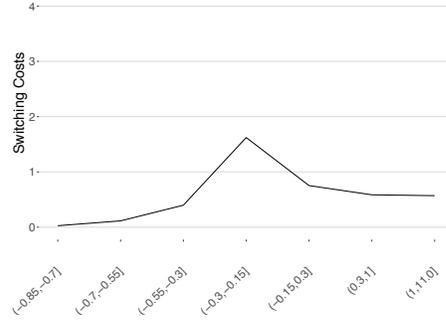
The comparison between the right and left columns of Figure 3 illustrates that when the average utility discrepancy of TA from its competitors is large (a more negative $\bar{\Delta}$), switching costs under the optimal purchase policy are lower. While the overall level of switching costs are lower everywhere, the relative drop in switching costs as progress becomes positive is less pronounced because the chances of the consumer continuing to buy in future periods is low.

The comparison between rows 1 and 2 illustrates that when the per-period utility of status π is larger, switching costs are higher, and the relative drop in switching costs as progress becomes positive is more pronounced, as the consumer expects a high probability of purchase in the future. As can be seen from the comparison of 2nd and 3rd rows, increasing the discount factor ρ has a similar effect, as the consumer values attaining status in the future more when the discount factor is larger.

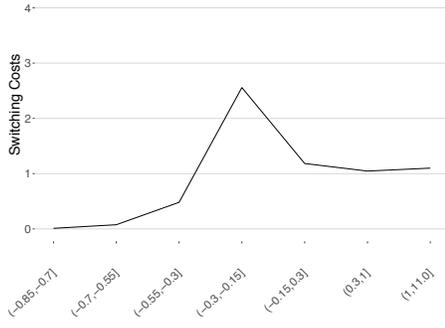
We replicate these comparative statics in Figure A.2 in Appendix A.1 using a model that allows for uncertainty in point accumulation. We examine a case where the consumer may have the opportunity to earn 0 or 4 points with 3/4 and 1/4 chance in each period, keeping expected points in each period equal to 1. Comparing Figure 3 and Figure A.2, we can see that the uncertainty depresses switching costs. It also



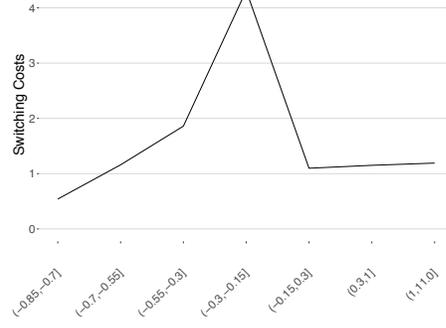
(a) $\bar{\Delta} = -2.5, \pi = 1, \rho = 0.95$



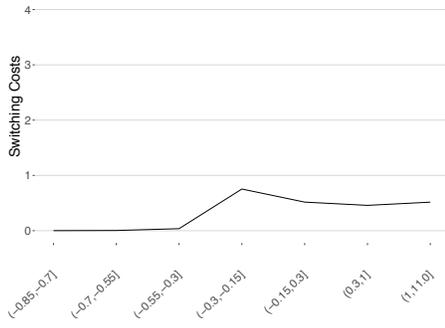
(b) $\bar{\Delta} = -1.5, \pi = 1, \rho = 0.95$



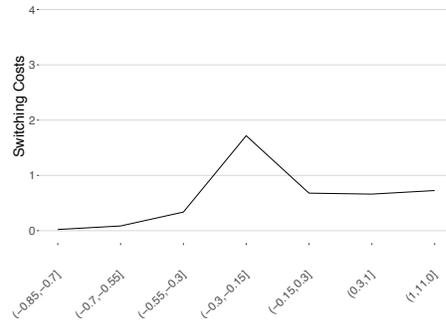
(c) $\bar{\Delta} = -2.5, \pi = 2, \rho = 0.95$



(d) $\bar{\Delta} = -1.5, \pi = 2, \rho = 0.95$



(e) $\bar{\Delta} = -2.5, \pi = 2, \rho = 0.8$



(f) $\bar{\Delta} = -1.5, \pi = 2, \rho = 0.8$

Notes: In our simulations, we draw the utility differential realizations from a normal distribution, $\Delta_t \sim N(\bar{\Delta}, \sigma^2)$. For this simulation, we set $\sigma = 2$. The consumer has the chance to earn $k_t = 1$ in each period if she buys. The discount factor is ρ and the additional utility the consumer gets from flying with status is π . We vary $\bar{\Delta}$, π and ρ in each subgraph.

Figure 3: Switching costs vary over progress.

widens the region of progress for which switching costs are elevated. Even when consumers are ahead of pace, they cannot afford to let go of point accumulation opportunities as easily. When they are considerably behind pace, they still have some chance of catching up because of the possibility of getting 4 points in the future. Therefore, switching costs start to increase at lower levels of progress and remain high at higher levels of progress because the prospect of being able to reach status is noisier.

3.3 Empirical Hypotheses and Inference

As the previous section demonstrates, we expect forward-looking consumers who care about attaining status to increase their willingness to forgo current consumption utility in order to book with TA as their progress increases from largely negative to less negative values. Importantly, if consumers' purchase decision is not impacted by status incentives, there is no systematic change in their purchase decisions as a function of progress. Therefore, the first hypothesis we test is whether or not consumers become more likely to forgo current utility (i.e., their switching cost λ_t increases) as their progress increases in the negative region. We also explore whether or not consumers become less likely to forgo current utility (i.e., their λ_t decreases) as their progress increases beyond a threshold, although the model indicates that under some parameter combinations the decline may not be detectable.

A second hypothesis we test is whether the consumers' willingness to forgo current utility is impacted by whether they are fully shouldering the ticket costs or not. In the previous section, we have cast the consumer decision problem as trading off the current utility gap Δ_{rt} and the switching cost λ_t . When the consumer does not fully shoulder the ticket costs, her Δ_{rt} is modified to $\Delta_{rt} + \alpha\gamma(p_{rtj'} - p_{rtj^*})$ where α denotes the proportion of the ticket price the consumer is not paying. Therefore, the consumer books when $\Delta_{rt} + \lambda_t + \alpha\gamma(p_{rtj'} - p_{rtj^*}) + \varepsilon_{irt} > 0$, which allows her to accept Δ_t realizations that are lower than she would if she were paying for the ticket herself. In the case of business travel, the employer shoulders more of the costs and the employee reaps almost all the status benefits.¹² Prior literature has theorized this asymmetry creates a moral hazard problem where the consumer books more expensive tickets than they otherwise would when traveling for business (Basso et al., 2009; Klemperer and Png, 1986; Levine, 1986; Tretheway, 1989; Shugan, 2005). In this paper, we examine whether consumers' increased willingness to pay is more pronounced when they travel for business than when they travel for leisure for the same level of progress.

¹²The employer may also benefit from some of the perks like reduced baggage fees. We thank an anonymous reviewer for pointing this out.

3.3.1 Inference Strategy

Switching costs are unobservable and need to be inferred from consumer choices. The consumer choice data we employ includes all booking records of TA from January 2010 to December 2011.^[13] This database includes frequent-flyer identification numbers and each member’s frequent-flyer program status at the time of booking, which allows us to construct a panel data set of bookings and to track members’ point accumulation at any given time. The database also includes flight descriptors, such as departure date, operating carrier, flight number, origin, and destination; booking details, such as purchase channel, fare class, booking date, revenues accrued by TA. Like most CRM databases, the database of TA is rich at the consumer level but does not have data on the consumer’s consideration set at the time of the decision. We augment this data with the quarterly count of carriers and their market shares on each of the routes that TA serves, obtained from the Airline Origin and Destination Survey Databank 1B (DB1B).^[14]

Next we explain how we use this data to create moments for identification in the absence of competitive individual choice data. Note that as λ_t increases, consumers become more likely to book at lower utility differential (Δ_{rt}) levels than they otherwise would. In other words, as their switching costs increase, consumers should become more willing to book trips with TA even when it is less appealing than its competition. A corollary of this assertion is that members should become more willing to pay higher prices to fly with TA. To examine the impact of within-person changes in progress towards status on purchase behavior, we conduct two main investigations that parallel these observations. First, we test whether members become more likely to book with TA on routes where TA has lower relative market share (reflecting lower general appeal). Second, we test whether the consumer becomes more likely to pay a higher price differential with respect to other passengers on the same flight. Both strategies have in common that airlines’ flight options and prices do not respond to an individual traveler’s purchase history and point accumulation with TA. As a result, after controlling for demand and supply conditions common to all travelers, systematic changes in the characteristics of a traveler’s bookings with TA associated with changes in her progress towards tier 1 can inform us about the impact of progress on her booking choices.

In what follows, we explain how we can identify changes in the underlying willingness to stick with the airline as progress increases. Our arguments build on the dynamic consumer decision model introduced in the

¹³The data also cover some of 2012. However, the 2012 data are not complete because the data were pulled in October 2012 and therefore cover tickets sold by that time. Therefore, we use only the transactions for flights departing in 2010 and 2011 for our analyses.

¹⁴This database provides a 10% random sample of airline passenger tickets on domestic routes. The data are reported at the quarterly level.

previous subsection. We start by formally restating and strengthening some conditions on the idiosyncratic shock ε_{irt} and the points gain k_r on a route r . These conditions ensure that frequent flyers on a given route and time are comparable and that points gain does not vary systematically with market conditions.

(IID) The shock ε_{irt} is iid across (i, r, t) from a distribution F that is supported on \mathbb{R} , has median zero, and has a density function that is bounded away from infinity. The points gain k_r is iid across r .

An immediate consequence of this assumption is that the switching cost $\lambda_t(K_{it}, k_r)$ at t depends on past idiosyncratic shocks and past k_r 's only through the miles to date K_{it} of individual i , which makes switching cost independent of the current idiosyncratic shock ε_{irt} . Also notice that k now has subscript r instead of t . The reason is that we now explicitly consider different routes. This notation emphasizes that the points gain k_r may vary across routes but individuals traveling on the same route at the same time will gain the same number of points. The time aspect of these points gains is secondary here and therefore suppressed in the notation. Furthermore, Assumption **(IID)** ensures that the average utility gap Δ_{rt} is the median of $\Delta_{irt} = \Delta_{rt} + \varepsilon_{irt}$ and we view Δ_{rt} as fixed for given (r, t) in the following.

Our main source of exogenous variation is the condition that consumers only choose airlines but not routes. This is an identifying assumption that lets us compare consumers across different market conditions without explicitly modeling the mechanism that chose route r for the consumer at t .

(IA1) The need to travel on route r at time t is exogenously determined.

Given assumption **(IA1)**, consumer i at time t books from target airline j' if and only if $I_{irt} := 1(\Delta_{rt} + \varepsilon_{irt} + \lambda(K_{it}, k_r) \geq 0) = 1$. In the sample, the market share of j' on route r relative to the best alternative j^* is therefore $\sum_{i=1}^{n_{rt}} I_{irt} / n_{rt}$, where n_{rt} is the number of consumers on that route. Under assumption **(IID)**, the market share of j' relative to the best alternative j^* in the population is $P(I_{rt} = 1) = P(\Delta_{rt} + \varepsilon_{rt} + \lambda(K_t, k_r) \geq 0)$, where we drop the i subscript because this market share does not depend on the individual. An immediate consequence of this representation is that if we observe two relative market shares $\mu_1 \neq \mu_2$, then the corresponding utility gaps must satisfy $\Delta_{r_1 t} \neq \Delta_{r_2 t}$ on any routes r_1 and r_2 with respective market shares μ_1 and μ_2 in order to not violate assumption **(IID)**. As we will now discuss, a systematic change in λ translates to a systematic change in the threshold of Δ admitted by consumers, which has implications for the distribution of TA's market share among the routes booked with TA.

We show that among the bookings consumer make with TA, switching costs of the consumers booking these trips must in expectation decrease in the market share of TA on the booked routes. The intuition is

that as their switching costs increase, consumers become more likely to book at lower utility differential Δ_{rt} levels than they otherwise would, which translates to consumers becoming more willing to book trips with TA even when it has low market appeal. Proposition 2 formalizes this intuition. Notice that the expectation in the proposition averages over the shocks ε_{irt} and the points gain k_r that influence the buying decision $I_{irt} = 1$ but holds the accumulated points K_{it} of the consumer fixed.

Proposition 2. *Suppose (IID) and (IA1) hold. Consider two markets with with corresponding market shares $\mu_1 < \mu_2$ at time t . Then $E[\lambda_t(K_{it}, k_{r_1}) \mid K_{it}, I_{ir_1t} = 1] \geq E[\lambda_t(K_{it}, k_{r_2}) \mid K_{it}, I_{ir_2t} = 1]$.*

Proof. By assumption (IID), we can without loss of generality drop i subscripts throughout this proof. Suppose $-\varepsilon_{rt} - \lambda_t(K_t, k_r) \sim F_t$. By construction, the market shares satisfy $\mu_1 = F_t(\Delta_{r_1t})$ and $\mu_2 = F_t(\Delta_{r_2t})$. If $\mu_1 < \mu_2$, then we must have $\Delta_{r_1t} < \Delta_{r_2t}$. Suppose j' is booked on each route, that is, $\Delta_{rt} + \varepsilon_{rt} + \lambda_t(K_t, k_r) \geq 0$ in both situations. Then the expectation E_k over k_r satisfies

$$E_k[\lambda_t(K_t, k_{r_1}) \mid K_t, \Delta_{r_1t} + \varepsilon_{r_1t} + \lambda_t(K_t, k_{r_1}) \geq 0] \geq E_k[\lambda_t(K_t, k_{r_1}) \mid K_t, \Delta_{r_2t} + \varepsilon_{r_1t} + \lambda_t(K_t, k_{r_1}) \geq 0]$$

by Lemma 1 in Appendix A.3. Integrate both sides with respect to ε_{1t} . Then

$$\begin{aligned} E_\varepsilon E_k[\lambda_t(K_t, k_{r_1}) \mid K_t, \Delta_{r_1t} + \varepsilon_{r_1t} + \lambda_t(K_t, k_{r_1}) \geq 0] &\geq E_\varepsilon E_k[\lambda_t(K_t, k_{r_2}) \mid K_t, \Delta_{r_2t} + \varepsilon_{r_1t} + \lambda_t(K_t, k_{r_1}) \geq 0] \\ &= E_\varepsilon E_k[\lambda_t(K_t, k_{r_2}) \mid K_t, \Delta_{r_2t} + \varepsilon_{r_2t} + \lambda_t(K_t, k_{r_2}) \geq 0], \end{aligned}$$

where the inequality follows from monotonicity of expectations and the equality follows from identical distributions across r and independence of ε_{rt} and $\lambda_t(K_t, k_r)$. By independence, $E_\varepsilon E_k$ can be viewed as a single expectation over both (ε_{rt}, k_r) and the desired result follows. \square

We can conclude from the proposition that the average switching costs $C_{K_{it}}(\mu) := E[\lambda_t(K_{it}, k_r) \mid K_{it}, I_{irt} = 1]$ by consumers with given miles endowment K_{it} as a function of the relative market share μ has a well defined inverse $c \mapsto \mu_{irt}(c, K_{it}) := C_{K_{it}}^{-1}(c) = \inf\{m \in \mathbb{R} : C_{K_{it}}(m) \geq c\}$ and $\mu_{irt}(E[\lambda_t(K_{it}, k_r) \mid K_{it}, I_{rt} = 1], K_{it})$ is the relative market share of the target airline j' on a given route r faced by a consumer i at time t . Because that market share is fixed for the consumer, changes in K_{it} imply that, for a given relative market share, average costs must change as well in a potentially nonlinear fashion. To capture this nonlinearity in the empirical analysis, we will approximate $g(K_{it})$ with a simple nonparametric binning technique.

We can similarly characterize how the prices paid by consumers vary in expectation over her endowment of points K_{it} . If we focus on a single flight at time t , we can again drop the route subscript and represent the utility gap as $\Delta_{it} = (X_{j'} - X_{j^*})\beta - \gamma(p_{tj'} - p_{tj^*}) + \varepsilon_{it}$. Because only the buying decision $\Delta_{it} + \lambda_t(K_{it}, k) > 0$ matters for the analysis, we can normalize $\gamma = 1$. In order to characterize the prices p_{it} paid by the consumer, we then make a simple identifying restriction that interprets, without loss of generality, the shock ε_{it} as variation in prices on the same flight.

(IA2) Individual prices can be represented as $p_{it} = p_{tj'} - \varepsilon_{it}$.

Provided [\(IID\)](#) and [\(IA2\)](#) hold, we can build expectations over p_{it} conditional on miles to date K_{it} and booking the flight. This expectation is $E[p_{it} \mid K_{it}, I_{it} = 1] = E[p_{it} \mid K_{it}, \Delta_{it} + \lambda(K_{it}, k) > 0]$. The price enters Δ_{it} linearly but, as K_{it} changes, $\lambda(K_{it}, k)$ varies in a potentially nonlinear and nonmonotone fashion. This truncates the support of p_{it} and influences the buying decision. More precisely, Lemma [1](#) in Appendix [A.3](#) implies that $\lambda_t(K_{it}, k)$ and the expected price in fact must move in the same direction. This immediately gives the following result.

Proposition 3. *Suppose [\(IID\)](#) and [\(IA2\)](#) hold. If $\lambda(K_{it}, k) < \lambda(K'_{it}, k)$ for two miles endowments K_{it}, K'_{it} , then $E[p_{it} \mid K_{it}, I_{it} = 1] \leq E[p_{it} \mid K'_{it}, I_{it} = 1]$. The result remains true if both inequalities are reversed.*

We implement this result in our empirical strategy by describing observed prices p_{it} relative to other prices on the same flight as a nonlinear function of K_{it} similar to the one described above for market shares. To explain how we translate this inference strategy to the particulars of our empirical context, the next section provides data details.

4 Data

Our data come from the transaction database of TA and Airline Origin and Destination Survey Databank 1B (DB1B). In their raw forms, the transaction data is at the booking-consumer-segment level and the DB1B data is at the coupon or ticket levels. All data construction details are included in the Appendix [A.4](#). In this section, we summarize important characteristics of the data and report relevant summary statistics, using industry-specific nomenclature of trips, routes and segments. A route is a pair of departure and destination cities. A one-way trip involves one route. A round-trip involves an outbound and an inbound route. If a route is direct, it involves one segment, otherwise it involves more than one segment.

Table 1: Yearly Member Booking Behaviors.

					Percentiles				
	Mean	Std	Min	Max	5p	25p	Median	75p	95p
Number of booking occasions	3.06	1.88	2	34	2	2	2	3	7
Number of route bookings	5.85	3.88	2	66	3	4	4	6	13
Total status points	9,756	7,056	500	35,000	2,430	4,534	7,477	12,935	25,130
Percentage, leisure travel	.66	.40	0	1	0	.33	1	1	1
Percentage, TA.com booking	.46	.42	0	1	0	0	.5	1	1

Summary statistics are constructed using data aggregated to individual-year observations (N=4,252,723).

Our sample consists of 3,489,102 active members who made bookings on at least two different points in at least one of the two years but did not accumulate more than 35,000 status-qualifying miles in a given year. Table 1 reports consumers’ purchase patterns. The average member in the sample books 5.9 routes over 3 booking occasions per year, and accumulates around 9,756 status-qualifying miles per year; however, there is a long tail of members who book more frequently and accumulate many more points. Members achieve tier 1 status 5% of the time. The average member takes leisure trips 66% of the time, but 18% of members book only business trips and 47% book only leisure trips. Members vary in their propensity to book on TA’s website across trips: The average member books on TA.com 46% of the time, but 27% of members only book on TA.com website and 36% never do. Our analyses will account for cross-member heterogeneity in a flexible manner.

We will examine how member’s booking behavior varies with the amount of accumulated points they have at a given time in the year, as summarized by the progress metric which depends on points accumulated and time left until the end of the year. We calculate the accumulation of status-qualifying miles for each member in each point-earning cycle using the point accumulation rules of the TA.¹⁵ We take into account members’ point earnings on all affiliated airlines to account for all qualifying points. Point accumulation accounting can be done based on booked flights or flown flights. Since the consumer knows the flights she booked, booked points better reflect her sense of progress towards the point requirements. Table 2 reports the distribution of cumulative booked points and progress of the members in our sample at the time they make each of their bookings. On average, members’ progress is behind the target pace required to attain status by 54% of that pace. The median booking is associated with an instance where the member is 81% behind the target pace. This is driven by the fact that 33% of observations are associated with the first booking in a point earning

¹⁵Status-qualifying points are equal to base points accumulated on each route multiplied by a factor greater than one for full-fare coach-class and all first-class tickets. The factor is one for discounted coach-class tickets. The base points are equal to the number of miles flown for routes longer than 500 miles. For routes shorter than 500 miles, the base points are equal to 500.

Table 2: Time, Booked Points to Date and Progress Associated with Bookings.

					Percentiles				
	Mean	Std	Min	Max	5 th	25 th	50 th	75 th	95 th
Accumulated Points	4,379	5,379	0	25,000	0	0	2,393	6,474	16,365
Week of Year	28.3	14.7	1	52	3	17	29	41	50
Progress	-.54	1.01	-1	50.9	-1	-1	-.81	-.42	.57

Summary statistics are based on all booking observations in our sample data.

cycle, where accumulated points are zero and progress equals -1 . The distribution statistics indicate a long tail of booking occasions where travelers can be significantly ahead of schedule. The 95th percentile of the progress distribution is an instance where the member is ahead of the target pace required to attain status by 57% of that pace.

The consumer booking decision is at the route level. Each booking is identified by the time of booking t , the time of departure τ , and the route r . Table 3 reports summary statistics of several characteristics of route bookings. Thirteen percent of routes in our sample are international and 59% of all routes are direct. Forty-five percent of the routes are booked directly on TA’s own website. The average distance traveled on a route is 1,412 miles and the average number of status-qualifying points earned is 1,666. This number is higher than the average distance due to the fact that a multiplier larger than one is effective for full-fare coach and all first-class ticket purchases. Only 3.6% of route bookings involve at least one flight in first-class and 65% of routes involve at least one flight with a discount fare in coach-class.

The top panel of Table 3 also reports characteristics of the itineraries each route booking is associated with. On average, consumers book their travel about 36 days before the outbound departure date. Roundtrips account for 79% of the bookings. For these trips, we define trip duration as the number of days between departure and return dates, which is 6 days for the average booking in the sample. In keeping with the industry practice, we classify a roundtrip trip as a leisure trip if it includes a Saturday-night stay-over, and as a business trip otherwise. Leisure trips account for 60% of all roundtrip bookings.

The business/leisure classification captures other booking characteristics that are correlated with business travel. For example, advance booking is 27 days longer for leisure trips than for business trips (p-value $< .001$), and trip length is 5.4 days shorter for business trips than for leisure trips (p-value $< .001$). Also, routes associated with leisure trips are more likely to involve discounted tickets (77% vs. 51%, p-value $< .001$). These patterns suggest that Saturday-night stay-over is a sensible proxy for the purpose of the trip.

The middle panel of Table 3 reports the distribution of ticket revenues broken down by class and inter-

Table 3: Booking Characteristics.

	Mean	Std	Min	Max	Percentiles				
					5 th	25 th	50 th	75 th	95 th
TA database, all bookings by members in the sample. N = 24,862,671									
International	.13	.34	0	1	0	0	0	0	1
Direct	.59	.49	0	1	0	0	1	1	1
Booked on TA.com	.45	.5	0	1	0	0	0	1	1
Distance (miles)	1,412	1,287	11	15,677	345	627	1,005	1,743	4,387
Status Qualifying Points	1,666	1,509	500	20,349	500	750	1,129	1,968	4,773
First-class	.036	.187	0	1	0	0	0	0	0
Discounted (coach-class)	.65	.48	0	1	0	0	1	1	1
Advanced Booking (days)	35.7	41.2	0	352	2	10	22	46	114
Roundtrip	.78	.41	0	1	0	0	1	1	1
(Roundtrip) Trip duration (days)	5.9	9.5	0	329	1	2	4	7	15
(Roundtrip Leisure Trip)	.60	.49	0	1	0	0	1	1	1
TA's Revenue for those bookings, U.S. Dollars									
First-class, Int'l. (N=264,150)	1,757	1,272	1	13,084	384	728	1,509	2,382	4,242
First-class, Domestic (N=634,585)	527	274	1	4,082	212	344	467	633	1,029
Coach-class, Int'l. (N=2,954,326)	428	300	1	9,202	100	229	370	550	948
Coach-class, Domestic (N=21,015,906)	191	124	1	2,936	69	111	157	231	429
Route descriptors from DB1B, domestic bookings. N = 21,221,612									
Number of carriers	3.8	1.4	1	9	2	3	4	5	6
TA's market share	.55	.25	0.01	1	.1	.39	.59	.72	.96
TA monopolist	.04	.19	0	1	0	0	0	0	0
Average one-way fare on route	219.5	63.3	10	2,070	129.2	172.6	214.1	259.3	327.7
TA's average one-way fare on route	224.5	65.2	10	2,150	134	175.9	217.2	266	332.8

The top and middle panels report summary statistics from TA database. The bottom panel reports summary statistics of booking characteristics that are constructed using the DB1B dataset. Therefore, they are only associated with domestic bookings.

national/domestic categories across the 24,862,671 route bookings that are in our final sample. The average coach class revenue per member on a domestic route is \$191 and the average first class revenue is \$527. Note that these revenues are about half of the revenues that TA obtains from a roundtrip itinerary. On international routes, the average revenue is \$1,757 from a first-class route booking and \$428 from a coach-class route booking. We would like to highlight two stark features of the price distribution. First, the first-class revenue distribution has little overlap with the coach-class revenue distribution. Second, price outliers at the low end of the price distribution seem to include revenues from tickets that are partially funded with points or refunds (revenues lower than \$15 account for 0.4% of the observations). For simplicity, we focus on coach-class bookings with positive revenues in our main analyses of these data. In Appendix [A.9](#) we present additional results from data samples that exclude revenues lower than \$15, and from specifications that include first-class prices. Our findings remain the same.

The bottom panel of Table [3](#) reports summary statistics of the competitive landscape on each route at the booking level. Because consumers care about price and convenience (e.g., departure/arrival time, number of connections) when booking a flight, a carrier may be very competitive on one or both of these dimensions and therefore command a large market share on one route, while lagging behind its competitors on another route.

Legacy airlines (American, Delta, and United) operate in a hub-and-spoke system, which initially gave them different strengths across geographies. Over time, both legacy and low-cost carriers (Spirit, JetBlue, and Southwest) have expanded their networks, leading to more overlap in routes served. Still, large differences remain in the market shares of these carriers’ services across different routes. Across domestic bookings, on average, TA has 2.8 competitors and 55% market share. Four percent of the bookings are made on routes that TA serves as a monopolist. On average, TA is slightly more expensive than other carriers on competitive routes.

4.1 Relative market share and price differential

In order to apply our inference strategy to the data, we construct dependent variables that bridge the theory and the nuances of the empirical context. First, we define TA’s *relative market share* on route r in quarter q associated with the departure date of τ as a proxy for its market appeal with respect to its competitors serving the same route. In particular, we calculate the normalized deviation of TA’s market share from the average market share on the route as $RMS_{rq} = (s_{TA,rq} - \overline{s_{rq}}) / \overline{s_{rq}}$ where $\overline{s_{rq}}$ is the average market share and $s_{TA,rq}$ is TA’s market share on the route. Proposition 2 shows the connection between switching cost λ and relative market shares because it holds the number of competitors fixed for simplicity of exposition. Empirically, however, we need to consider the fact that number of active airlines on each route may differ and how it impacts the relative appeal of TA as measured by market shares. For example, on route 1, TA may be the trailing one of two competitors, commanding a 30% market share, and on route 2, it may be the leading airline out of eight competitors, commanding a 25% market share. Even though TA has lower market share in route 2, TA has a higher relative appeal (a higher average Δ) compared to its competitors on route 2 than on route 1. Congruently, on route 1, TA’s *RMS* is -0.4 , reflecting the fact that its market share is 40% lower than the average share on the route. On route 2, TA’s *RMS* is 1, reflecting the fact that its market share is double the average share on the route.

In our empirical analyses, we explore how the *RMS* associated with the bookings of a consumer changes with progress. To that end, we associate route-quarter level RMS_{rq} to each booking based on the quarter of the booking departure date. In this manner, we obtain $RMS_{ib(t\tau)}$ for each booking observation. Table 4 presents summary statistics of $RMS_{ib(t\tau)}$. On average, members make bookings on routes where TA’s market share is 1.56 times the average carrier’s market share, although this relative market share metric spans a large range. At one extreme, TA has negligible market share against a competitor who is a monopolist, i.e.,

Table 4: Additional Variables Constructed for Regressions.

					Percentiles				
	Mean	Std	Min	Max	5 th	25 th	50 th	75 th	95 th
<i>RMS</i> (TA’s relative market share)	1.56	1.32	-.99	6.1	-.49	.53	1.51	2.63	3.73
<i>markup</i> (TA’s relative price)	.03	.1	-.98	8.8	-.11	-.01	.02	.07	.15
<i>RPD</i> (relative price differential)	.35	1.01	-1	150.4	-.43	-.13	.03	.46	2.22
<i>PDPM</i> (price differential per mile)	6	19	-364	1,126	-8	-2	0	7	40
Median coach-class revenue on flight	189	126	1	8,469	71	112	157	219	445
<i>advbook</i> (advance booking differential)	.31	1.23	-1	304	-.85	-.41	0.01	.64	2.37
Median advanced booking on flight	31.4	27.5	0	349	12	17.8	24	34	80

its relative market share is close to -1 . At the other extreme, TA’s market share is 6.1 times the average market share on the route.

Table 4 also presents summary statistics of the *relative price differential* (*RPD*) and *price differential per mile* (*PDPM*). These variables capture the price differential the consumer pays compared to other consumers who choose the same flight.¹⁶ In particular, we measure the *RPD* that consumer i pays when booking at time t in comparison to others in the same cabin as $RPD_{ib(t\tau)} = \sum_{f \in F_b} (p_{if} - \bar{p}_f) / \sum_{f \in F_b} \bar{p}_f$ where p_{if} is the revenue associated with individual i ’s ticket, F_b is the set of flights the traveler takes on a booking b and \bar{p}_f is the median coach-class revenue obtained from all tickets sold on the same flight instance f .¹⁷ Note that all medians are constructed using the entire transaction database of TA and therefore reflect the median prices paid by all members and non-members for coach-class tickets on a given flight. This price differential is normalized by the median and can therefore be interpreted as the percentage deviation from the median. Another useful metric is the price differential normalized by distance in miles, which we define as $PDPM_{ib(t\tau)} = \sum_{f \in F_b} (p_{if} - \bar{p}_f) / \sum_{f \in F_b} distance_f$. It measures the additional cents per mile the consumer is paying and is considered here because it aligns with the industry standard for reporting price differentials.

In order to control for the advance booking differential of the consumer compared to others making the same trip, we define $advbook_{ib(t\tau)} = \sum_{f \in F_b} (ab_{if} - \overline{ab}_f) / \sum_{f \in F_b} \overline{ab}_f$ where ab_{if} indicates the number of days in advance member i made the booking and \overline{ab}_f is the median advance booking across all coach-class tickets sold on flight instance f . The average member in our sample books 31% earlier than others on the same flight which corresponds to about 31 days before departure.

¹⁶We could alternatively compare the price the consumer paid to prices on the same route and week. Studying within-flight price differential offers tighter controls for quality differences based on departure/arrival times, aircraft type, directness of service, etc.

¹⁷A flight instance is identified by a carrier, flight number, route and departure date, and time.

4.2 Exploratory Analysis

We first document some data patterns in order to provide guidance and intuition before proceeding to our main analysis. Ideally, we would like to observe the booking behavior of a consumer with different point accumulation levels in the same period t . This is of course impossible because we cannot observe the same consumer at the same time in two counterfactual situations. However, when the consumer takes a long flight, the increase in the total points is quite significant within a short time range. Therefore, the willingness to pay for TA before and after such a flight should be substantially different.

To examine this hypothesis, we focus on the impact of international trips, which on average cover 3 times the distance of domestic trips in the data. We compare how consumers' booking behavior right before and right after their first international trip are different¹⁸ In particular, we restrict the booking period to 60 days on either side and contrast the booking right before and right after the long trip. Focusing on the subset of consumers who had at least one international trip across the two years, we estimate

$$RMS_{ib(t\tau)} = \alpha_i + \beta AfterIntl_{ib} + \nu_\tau + \varepsilon_{ib(t\tau)}, \quad (8)$$

where $RMS_{ib(t\tau)}$ refers to TA's relative marketshare on route r in quarter q associated with a booking that member i made at time t to travel at time τ . While the unit of observation is at the individual booking level, we carry the $(t\tau)$ as a reminder that each booking is associated with two dates: a booking date t and a departure date τ . Because booking decisions are made at time t , we explore whether the progress of the decision-maker at time t impacts her choices. We include week fixed effects ν_τ for departure times of the bookings in order to account for seasonality in travel or possible changes in TA's popularity across time due to changes in service or other marketing activities. In addition, we include intercepts α_i for each individual to control for unobserved heterogeneity in preferences. The variable $AfterIntl_{ib}$ equals one for the trip booked after flying internationally with TA and equals zero for the trip right before the international one. We also estimate

$$P_{ib(t\tau)} = \alpha_i + \beta AfterIntl_{ib} + \varepsilon_{ib(t\tau)} \quad (9)$$

where the dependent variable P is either the relative price differential (RPD) or price differential per mile ($PDPM$). This specification does not include departure week fixed effects because the price differential is defined relative to prices on a flight departing at a particular date. Variations across time in travel seasonality

¹⁸We thank an anonymous reviewer for this suggestion.

and TA’s services are therefore already controlled for.

We find that the point boost from the international trip generates an immediate impact on the price differential the consumer is willing to accept. We find that members pay 8.4% (p-value < 0.001) or 1.24 cents more per mile (p-value = 0.004) on average compared to others on the same trip after the international trip than they did before the international trip. We do not find significant changes in the market share of TA among booked trips, possibly because this approach severely restricts the number of booking instances to 9,876 compared to the 20,319,862 in our main analyses.

In our main analyses, we extend the intuition of this within-person analysis to study how members change their booking behaviors as their progress varies. Our theoretical model predicts that the booking policy of consumers depends on the time left to the end of the year ($T - t$) and the number of points they have accumulated (K_t) at the time of booking. The simple analysis presented above takes advantage of large changes in the number of accumulated points over a short period of time but does not account for differences across consumers in the time they have left until the end of the year. Our main analyses will incorporate this heterogeneity.

5 Consumer Response to Dynamic Incentives Created by Status Awards

In our main analyses, we examine how the market shares and the price differentials that get selected into bookings with TA respond to a consumer’s progress. The theoretical model shows that switching costs increase with progress before possibly declining as they get positive. We empirically test for this relationship in this section.

5.1 Regression Specifications

Our goal is to test whether the propensity of booking TA on routes where TA has substantially lower appeal compared to its competition increases as members make progress toward attaining status. To this end, we estimate

$$RMS_{ib(t\tau)} = \alpha_i + \sum_{k=0}^8 \beta_k I(\text{progress}_{it} \in S_k) + \nu_\tau + \varepsilon_{ib(t\tau)} \quad (10)$$

where $RMS_{ib(t\tau)}$ refers to TA's relative market share on route r in quarter q associated with a booking that member i made at time t to travel at time τ . While the unit of observation is at the individual i and booking b level, we carry the $(t\tau)$ as a reminder that each booking is associated with a booking date t and a departure date τ . Because booking decisions are made at time t , we explore whether the progress of the decision-maker at time t impacts her choices. We again include week fixed effects ν_τ for departure times of the bookings in order to account for common unobservables that impact choices associated with travel at that time. For example, in certain times of the year, TA may have a different service schedule on some routes, increasing their popularity. Or, members may be more likely to travel on certain routes during certain times of the year. We also again include individual intercepts α_i to control for unobserved heterogeneity in preferences.

To capture potential nonlinearities in travelers' responses, we transform the continuous progress metric into a categorical predictors by discretizing it into ranges S_k . Here, S_0 includes the 33% percent of the observations associated with the first booking in a cycle, where progress equals -1 . This serves as the reference level. Other S_k 's are characterized by the following progress values as cutoff points: $-.9, -.8, -.7, -.55, -.3, .1, .6$. These cutoffs correspond to the 40th, 50th, 60th, 70th, 80th, 90th and 95th percentiles of the progress distribution, and the maximum value of progress is the upper bound. The set of coefficients β_k reflect the degree to which a within-person change in progress leads to a within-person change in the willingness to choose TA over its competition on routes characterized by TA's relative market share. We expect β_k to increase in magnitude (becoming more negative) as progress increases¹⁹. However, at some point, when the member feels that she is significantly ahead of the pace that is required to attain status, the member may scale back her efforts.

Similarly, we examine how the price differential members pay in their bookings with TA responds to progress with

$$P_{ib(t\tau)} = \alpha_i + \sum_{k=0}^8 \beta_k I(\text{progress}_{it} \in S_k) + \varepsilon_{ib(t\tau)}, \quad (11)$$

where the dependent variable P is either the relative price differential (RPD) or price differential per mile ($PDPM$). This specification again does not include departure week fixed effects because the price differential is defined relative to prices on a flight departing at a particular date and therefore absorbs all time and flight level unobservables. Individual fixed effects α_i control for unobserved heterogeneity in preferences and the set of coefficients β_k reflects the extent to which a within-person change in progress leads to a within-person

¹⁹We caution the reader that β_k estimates are likely to be biased downwards because our analyses regarding RMS rely on quarterly and aggregate moments from external data. While these moments allow us to understand consumer tradeoffs between TA and its competition even when we do not observe the exact competitive options consumers considered, they also introduce noise. Due to this noise, we expect an attenuation bias in our estimates.

change in the willingness to pay for TA. We expect an increase in the willingness to pay for TA in response to increased progress.

We examine two types of response heterogeneity across groups of consumers: consumers whose home airport is a TA hub vs. not, and business travelers vs. leisure travelers. When exploring heterogeneity across consumer segments, we allow β_k to vary by segment by interacting an indicator variable with our progress metric in specifications (10) and (11). For example, in exploring business and leisure travelers differ in their responses, we capture this heterogeneity with $\beta_{sk} = \beta_{1k}1(\text{BusinessSegment}_i = 1) + \beta_{2k}1(\text{BusinessSegment}_i = 0)$ where a consumer belongs to the Business Segment if her proportion of business travel is at least 1/2.

Finally, we extend our main specifications to explore the changing nature of the decision problem for the same traveler when she is more likely to be shouldering the costs of the ticket. We cannot directly observe whether the traveler is paying for the ticket herself but we can plausibly assume that one is more likely to pay out of own pocket for leisure trips than for business trips. We therefore offer a test for the role of asymmetric incentives by examining whether the same individual's response to progress is more pronounced on business trips. In particular, we estimate

$$RMS_{ib(t\tau)} = \alpha_{iw} + \sum_{k=1}^8 (\beta_{sk} + \mu_k I(w_{ib(t\tau)} = B)) I(\text{progress}_{it} \in S_k) + \nu_\tau + \varepsilon_{ib(t\tau)} \quad \text{and} \quad (12)$$

$$P_{ib(t\tau)} = \alpha_{iw} + \sum_{k=1}^8 (\beta_{sk} + \mu_k I(w_{ib(t\tau)} = B)) I(\text{progress}_{it} \in S_k) + \varepsilon_{ib(t\tau)}, \quad (13)$$

where $w \in \{L, B\}$ indexes the purpose of the booking: *Leisure* or *Business*. Person and trip purpose specific fixed effects α_{iw} control for each member's average willingness to pay for leisure and for business travel. Since traveling for business is more common among the business travelers, segment heterogeneity must be teased apart from the changing nature of the decision problem for the same traveler when she is shouldering the costs of the ticket versus when she is not. Segment level heterogeneity in response is specified as $\beta_{sk} = \beta_k + \sigma_k 1(\text{Segment}_i = 1)$ and a member belongs to Segment 1 if her proportion of business travel is larger than or equal to 1/2. We estimate this specification using a roundtrip subsample of members who booked at least 2 leisure and 2 business trips.

After controlling for individual differences in the differential price paid for leisure versus business travel and segment differences in response to progress, we interpret parameters μ_k as reflecting the impact of asymmetric incentives when ticket costs are not shouldered fully. Admittedly, the Saturday-night proxy for leisure trip is noisy, which may bias the μ_k towards zero. Thus, a significant μ_k would be a conservative

indicator that members respond to status incentives differently based on whether they are paying for the trip themselves or not.

5.2 Results

5.2.1 Average Response to Progress

The β_k estimates from specifications (10) and (11) are reported in Tables 5 and 6, respectively. The coefficients (and associated confidence intervals) from the relative market share (Table 5, Column 1) and price differential per mile (Table 6, Column 2) analyses are plotted in the first row of Figure 4 for ease of interpretation.

A decline in the average *RMS* on routes consumers book with TA suggests that consumers become more likely to fly with TA on routes where TA has lower overall appeal - a pattern we expect as switching costs increase. Consistent with our hypothesis, we see a decline as progress rate increases to the pace required to attain status. Compared to the reference progress level when the consumer is just starting out (baseline), when the member is only 55-70% behind target, TA's relative market share is 6.5% lower (β_4) on routes booked with TA; and, when the member's progress is around the required progress rate, TA's relative market share is 9.6% lower (β_6) on routes booked with TA; and, at peak response, TA's relative market share on booked routes is 10.7% lower (β_7). Each of these coefficients are significantly lower than the preceding one at p-values < 0.001 . The progressive decline as progress increases shows that consumers become more willing to book trips with TA on routes even when TA is not as attractive for other consumers in the market. As Proposition 2 shows, this result suggests that the consumers are becoming more likely to sacrifice current utility in exchange for the possibility of attaining status. When the member's progress speed is significantly above target, the consumers become slightly less willing to sacrifice current utility, captured by the uptick in the β_8 coefficient compared to β_7 (p-value of contrast < 0.001).

Consumers with increasing switching costs may pay for that loyalty by booking TA even when it offers less desirable flight options than its competition, when it is more expensive than its competition for what it offers, or both. The result that consumers are increasingly more likely to book on routes where TA has low appeal overall is consistent with all three of these possibilities. To show that consumers trade off not just prices but also convenience, we also provide supplemental evidence that consumers become more likely to book TA on routes where it offers less convenient flight options. We offer two pieces of evidence on this point. In Column 2 of Table 5, we present estimates from specification (10) where the dependent variable is

Table 5: Impact of Progress Towards Status on Booking Behavior: Attractiveness of TA.

	(1)	(2)	(3)
Units	<i>RMS</i> %	Rel.Direct Svc.Ratio %	Indirect Strong %
β_1	1.37*** (0.147)	-0.764*** (0.265)	-0.987*** (0.0850)
β_2	-2.79*** (0.126)	-5.23*** (0.243)	-0.338*** (0.0622)
β_3	-5.77*** (0.143)	-6.84*** (0.273)	0.178*** (0.0667)
β_4	-6.51*** (0.145)	-6.59*** (0.277)	0.565*** (0.0671)
β_5	-8.10*** (0.148)	-7.81*** (0.282)	0.829*** (0.0692)
β_6	-9.64*** (0.160)	-8.61*** (0.305)	1.12*** (0.0761)
β_7	-10.7*** (0.198)	-10.2*** (0.382)	1.51*** (0.0957)
β_8	-9.40*** (0.209)	-10.1*** (0.418)	1.22*** (0.102)
Observations	20,319,403	16,801,042	7,430,299
R^2	0.563	0.507	0.553

This table reports estimates from specification (10) for domestic bookings where TA has competition on the route. Dependent variables indicated as column titles. Relative Direct Service Ratio is only well-defined for domestic routes on which at least one carrier provides direct service. Indirect Strong is only well-defined for indirect bookings on domestic routes. In parentheses, we report robust standard errors clustered at the member level. *** p<.01, ** p<.05, * p<0.1.

TA’s direct to total service passenger ratio relative to the average direct to total passenger ratio of all carriers on the route, if direct service on the route is provided by any carrier. On routes where TA’s competitors’ direct service is more appealing than that of TA’s, this relative directness ratio is negative. The estimates show that consumers become more likely to book with TA on routes where its direct service is less desirable than its competitors. We also provide evidence that consumers become more likely to book indirect trips with TA on routes even when its competitors offer more popular direct options than TA. To this end, we define *Indirect Strong* as an indicator for an indirect flight being booked with TA on a route where the majority of travelers fly direct. This means that some carriers have desirable direct flight options but TA is not one of them.²⁰ Column 3 of Table 5 presents results from specification (10), where the dependent variable is *Indirect Strong* and the estimation sample is restricted to indirect bookings. The estimates show an increase in the propensity that an indirect flight is booked with TA on a route even when there are more popular competitive direct flights. This finding suggests that consumers become increasingly likely to prefer an indirect flight with TA over its competition even when the competition offers more popular direct flights.

Now we turn to discussing the changes in consumer’s willingness to accept higher prices than usual. The β_k estimates from specification (11) are reported in Column 1-2 of Table 6. The results also suggest that

²⁰In particular, we define an indirect booking to be *Indirect Strong* if at least 50% of the passengers prefer direct flights, but where TA carries less than 10% share of those direct passengers. Our results are not sensitive to these cutoff choices.

the consumers are becoming more likely to sacrifice current utility as their progress increases towards levels required to make steady gains towards status. Coefficients reported in Column 1 show that compared to the price differential the member is willing to pay at the very beginning of the point accumulation cycle, the price differential she is willing to pay compared to others on the same flight gradually increases, ranging from 4.7% to 7.6% ($\beta_2 - \beta_5$). It reaches 8% (β_6) when the member's progress is closest to the required pace to attain status. Each of these coefficients are significantly lower than the preceding one at p-values < 0.001 . When progress is close to being on track, consumers on average pay 1.95 cents more per mile (β_6) compared to the prices they pay when starting out. This corresponds to 8% of the average 24 cents paid per mile, or 12% of the median 16 cents paid per mile across all flights in the data. As Proposition 3 shows, these results suggest that the consumers are becoming more likely to sacrifice current utility in exchange for the possibility of attaining status.

We also observe a downturn in consumer willingness to pay when consumers are significantly ahead of the required pace. The relative price differential (Column 1) declines to 7% (β_7) and then to 5% (β_8) after the progress measure becomes positive (contrast of β_8 and β_6 p-value < 0.001). Similarly, price differential per mile (Column 2) also takes a dip from an increase of 1.95 (β_6) to 1.89 (β_7) and then to 1.40 (β_8) cents per mile (contrast of β_8 and β_6 p-value < 0.001).

Table 6: Impact of Progress Towards Status on Booking Behavior: Price and Search.

Units	(1) <i>RPD</i> %	(2) <i>PDPM</i> cents	(3) <i>RPD</i> %	(4) <i>PDPM</i> cents	(5) TA.com %
β_1	3.64*** (0.0899)	0.398*** (0.0201)	1.11*** (0.0896)	-0.112*** (0.0200)	-0.545*** (0.0481)
β_2	4.67*** (0.0676)	0.876*** (0.0144)	2.31*** (0.0679)	0.400*** (0.0144)	0.408*** (0.0392)
β_3	6.13*** (0.0753)	1.254*** (0.0160)	3.68*** (0.0754)	0.760*** (0.0160)	1.30*** (0.0439)
β_4	6.80*** (0.0756)	1.495*** (0.0161)	4.24*** (0.0755)	0.980*** (0.0160)	2.22*** (0.0435)
β_5	7.61*** (0.0753)	1.748*** (0.0161)	4.89*** (0.0757)	1.199*** (0.0161)	3.26*** (0.0435)
β_6	8.00*** (0.0832)	1.951*** (0.0180)	5.14*** (0.0831)	1.375*** (0.0179)	4.20*** (0.0472)
β_7	6.96*** (0.112)	1.892*** (0.0246)	4.44*** (0.110)	1.385*** (0.0243)	4.63*** (0.0614)
β_8	5.19*** (0.116)	1.401*** (0.0250)	2.69*** (0.115)	0.899*** (0.0247)	4.04*** (0.0662)
<i>adobook</i>			-11.3*** (0.0828)	-2.276*** (0.0155)	
Observations	23,942,392	23,942,392	23,941,883	23,941,883	23,942,392
R^2	0.411	0.424	0.429	0.440	0.658

In parentheses, we report robust standard errors clustered at the member level. *** p<.01, ** p<.05, * p<.01. This table reports estimates from specification (11). Sample includes only coach-class bookings. The number of observations in Columns 3 and 4 are slightly lower due to *adobook* not being well-defined for some observations.

Prices within a flight may vary for several reasons. One well known reason is that prices generally increase as time to departure decreases. In order to examine how much of the increase in the relative price differential can be explained by changes in booking timing, we re-estimate specification (11) while controlling for changes in a traveler’s propensity to book earlier or later. Column 3-4 of Table 6 report the results for *RPD* and *PDPM* as dependent variables, respectively. As expected, the impact of progress on the relative price differential travelers pay is more muted but still quite large. Even after controlling for changes in advance booking as a driver of price changes, at peak values, we still see a 5% increase in the relative differential price and 1.38 cents increase in the differential price per mile travelers are willing to pay when they stretch to attain status (β_6 , Columns 3 and 4, respectively).

There is also significant variation in prices across ticket class types within the economy cabin, even conditional on booking time. Discounted fares have more restrictions. Higher fares are associated with more flexibility, higher chance of business class upgrades, and a higher point multiplier. Finally, due to airline revenue management practices, prices increase as a function of remaining inventory of seats. The arrival of bookings for a flight may have a large random component. A flight may have more (fewer) bookings than its competitors for a given departure date, resulting its price to be higher (lower) than comparable products on the market. Finally, prices may also vary across different channels. Therefore, *ceteris paribus*, the more a consumer is willing to search, the lower a price she is able to find in a competitive market.

As additional evidence of increased willingness to book with the airline, we investigate whether travelers become more likely to make a booking on the airline’s website (TA.com). If the consumer makes the booking on TA.com, instead of calling an agent or originating the booking on price aggregator websites like Travelocity, she is less likely to be informed of competitors’ prices. We expect consumers to be more likely to book on TA.com as their switching costs increase. In Column 5 of Table 6, we report results from specification (11), where the dependent variable is an indicator of having originated the booking on TA.com. The results are again suggestive of consumers becoming more willing to book with TA as they progress towards attaining status. The estimates indicate that the consumers are 4.2% to 4.6% (β_6, β_7) more likely to initiate booking directly on TA.com when their progress is on or slightly ahead of track compared to when they are just starting out earning points. Similar to earlier results, when the consumer is significantly ahead of the required pace to attain status, their increased propensity to book on TA.com slightly decreases ($\beta_8 = 4.04\%$, contrast with β_7 p-value < 0.001).

5.2.2 Response heterogeneity

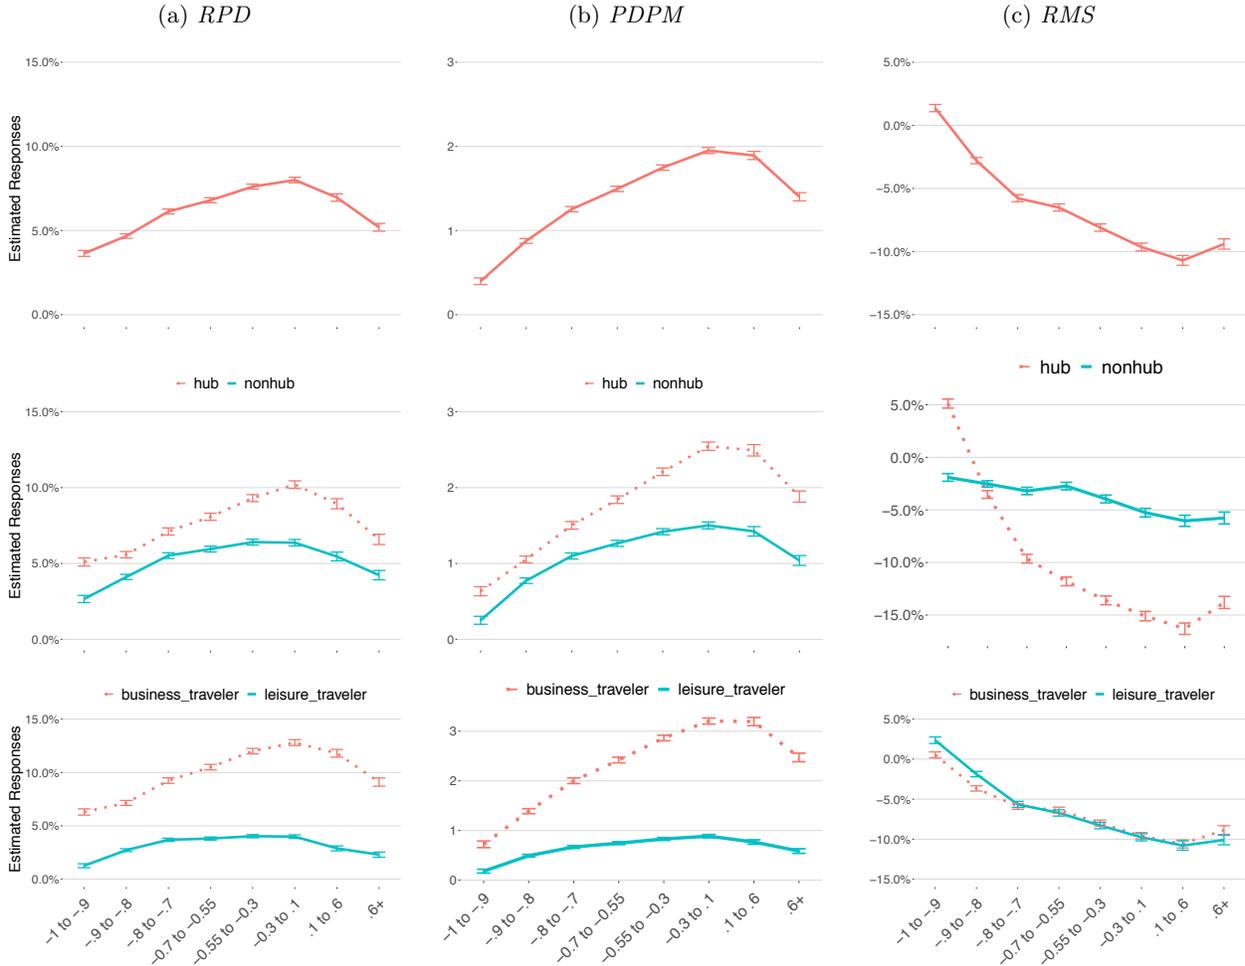
These main effects reflect average effects in response to status incentives, aggregating over different response curves across heterogeneous consumers. We explore two sources of heterogeneity. The first source of consumer heterogeneity pertains to whether the home airport of the consumer is a TA hub.²¹ The airline industry literature has recognized the consumer preference heterogeneity across distance to airline’s hubs, market power airlines have at their hubs, and the role frequent-flyer programs play in that market power (Borenstein, 1990, 1991, Lederman, 2007, 2008). In our data, 38% of the consumers’ home airport is hub of TA. If consumers who live close to a TA hub have a stronger response to status incentives than consumers who are far from a TA hub, it would amplify the role of status incentives in creating market power for the airline at its hubs. There are two reasons why we may expect a stronger response. First, consumers at a TA hub are more able to substitute into TA flights as their switching costs increase because TA offers service across a wider selection of routes at a hub. Second, and relatedly, these consumers may have more opportunities to enjoy the perks of attaining status. In the second row of Figure 4 we plot the heterogeneous responses of consumers who live close to a hub versus not (coefficients are also reported in Table A.1). We find that for the same level of increase in progress, consumers living close to a TA hub more strongly increase the relative price and more strongly decrease the relative market appeal they are willing to accept. This heterogeneity result suggests that a higher responsiveness to status incentives may be a contributing factor to airlines having more market power in their hubs.

As a second source of response heterogeneity, we examine whether leisure and business travelers differ in the extent to which they modify their purchase behaviors as they progress towards status. Among the consumers in the data, 36% are denoted business travelers, who account for 58% of the bookings. The fact that business and leisure travelers have different purchasing patterns has been widely recognized in the airline industry. For example, business travelers fly more on average and are less price sensitive (Brons et al., 2002). Therefore, we may expect differential responses to status incentives across the two segments because of differences in the benefits of status and differences in tradeoffs of these benefits with ticket prices.

In the third row of Figure 4 we plot the heterogeneous responses of business versus leisure travelers (coefficients also reported in Table A.2). We find that business travelers increase the relative price they pay more aggressively as they make progress towards status. However, we do not find meaningful differences across business and leisure travelers in the extent to which they become more willing to choose TA on

²¹We thank an anonymous reviewer for this suggestion.

Figure 4: Response Heterogeneity for RPD, PDPM, and RMS.



Notes: This figure plots the estimated coefficients from Specifications (10) and (11) and associated 95% confidence intervals. The top row depicts the average consumer response to progress, plotting coefficients reported in Column 1 of Table 5 and Column 2 & 3 of Table 6. The second row depicts response heterogeneity across hub (red, dotted line) vs. nonhub consumers (blue, solid line), plotting coefficients reported in Table A.1. The third row depicts response heterogeneity across business travelers (red, dotted line) and leisure travelers (blue, solid line), plotting coefficients reported in Table A.2.

routes where it is less appealing. We caution the reader that heterogeneous price responses across business and leisure travelers cannot immediately be attributed to differences in consumer price sensitivity. By construction, business travelers are more likely to take business trip, which are commonly not paid by the traveler. In the next section, we separate the impact of not shouldering the costs of travel and preference heterogeneity.

5.2.3 Differential Response when Traveling for Business vs. Leisure

In order to separately study the impact of not shouldering the costs of travel, we estimate specifications (12) and (13). Our estimates of β , σ , and μ , reported in Table A.5, respectively capture the baseline response to progress, how this response differs for business travelers, and how it differs within individuals based on whether they are flying for business or for leisure. Based on these estimates, Figure 5 plots the inferred responses to progress by business travelers and leisure travelers, when each segment is traveling for business versus for leisure.

It is immediately clear that the asymmetric incentives across leisure and business travel occasions account for a much larger proportion of the price response heterogeneity compared to differences across segments conditional on the type of trip.²² The highest price responses are from the business traveler segment taking business trips closely followed by the leisure traveler segment taking business trips. The responses are much more muted, for both segments, when consumer are taking leisure trips.

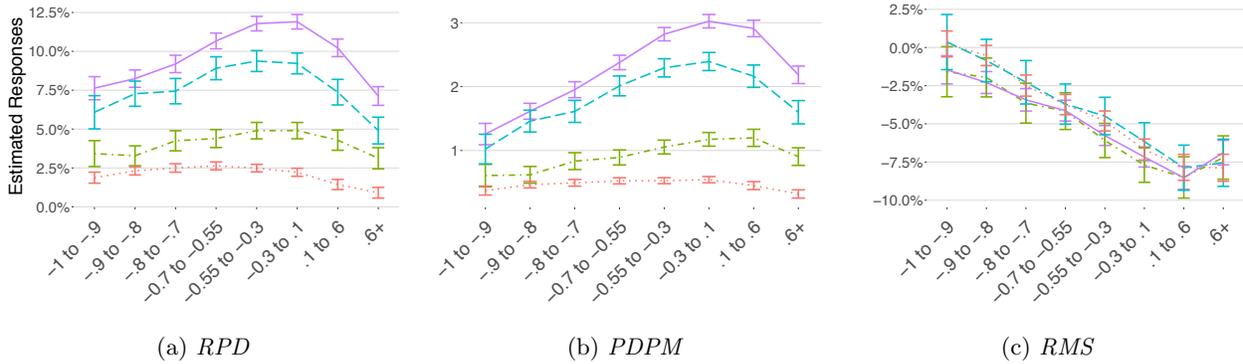
The results indicate that asymmetric incentives in business travel contribute to a substantial increase in members' willingness to pay more as they make progress towards status. For example, when the member's progress is just within reach of the pace required to attain status (progress between $-.3$ and $.1$), she pays 7% more or alternatively 1.85 cents more per mile when she is booking a business trip rather than a leisure trip. Note that this additional willingness to pay is above and beyond the usual price difference she pays for business versus leisure trips (captured by α_{iw}). Overall, the results clearly show that on average members' willingness to pay response to increased progress is higher when someone else is likely paying for these members' trips.

In contrast, we do not see differences across business and leisure travelers, or across business and leisure bookings within an individual, in the extent to which they become more willing to choose TA on routes where it is less appealing. This result is entirely congruent with asymmetric incentives arising from the consumer not shouldering the monetary costs of the ticket. While the traveler bears little (or none) of the monetary costs of travel when traveling for business, she still bears the inconvenience just the same regardless of who pays for the ticket. Moreover, during business travel, consumer may even be more sensitive to inconvenience. Therefore, we would not expect consumers to show a higher willingness when they travel for business to book with TA even when it offers less appealing flight options in the market.

The differential willingness to pay response to progress when consumers are flying for business versus

²²In models where these asymmetric incentives are not accounted for, all the trip-level heterogeneity would be mis-attributed to segment differences, since by construction, business travelers are more likely to take business trips.

Figure 5: Estimated Responses to Progress by Business and Leisure Segments and Trip Type.



Notes: This figure plots the estimated response to progress (and associated 95% confidence intervals) for segments 1 and 2, when they are flying for business versus leisure, as inferred from the estimates in Table A.5. The purple solid line (at the very top in panels (a) and (b)) plots the response of Segment 1 and the blue long-dashed line (right below it) represents the response of Segment 2 when flying for business. The green dot-dashed line plots the response of Segment 1 (third from the top in panels (a) and (b)) and the red dotted line (at the bottom) plots the response of Segment 2 when flying for leisure.

leisure is consistent with the hypothesis put forth by earlier literature that loyalty programs may create moral hazard as the decision-maker (employee) accrues more benefits than the payer (employer) (e.g., see Basso et al., 2009; Klemperer and Png, 1986; Levine, 1986; Tretheway, 1989; Shugan, 2005). It also provides additional support in favor of dynamic incentives created by the potential status award in shaping consumer choice and consumers responding to asymmetries in incentives, rather than the observed increases in loyalty being driven completely by a state dependency in preferences.

5.3 Discussion

We find evidence consistent with switching costs increasing in progress for most of the progress range. We also document that switching costs decrease once progress becomes sufficiently large. Finding this pattern empirically suggests substantial sophistication on the part of the travelers: They are more likely to sacrifice current utility for the possibility of attaining status when returns to doing so are higher. When their progress is substantially slow, such that their chances of attaining status is low, they are less likely to sacrifice current utility than they otherwise would in order to fly with TA. When their progress is faster than the pace required to make steady progress, they also become less likely to sacrifice current utility, since their chances of attaining status are high and not substantively altered by additional purchases. In the negative region of progress where the probability of attaining status increases quickly with additional points, consumers are willing to pay more and are willing to travel with TA even if it does not offer an attractive option. Our results also show that this response pattern is particularly pronounced for consumers whose home airport

is a TA hub, and for business travelers. However, much of the heterogeneity across business and leisure travelers actually comes from asymmetric responses to progress when consumers fly for business vs. leisure, rather than consumer heterogeneity.

The estimates suggest economically substantial reactions to status incentives. We document that on average frequent-flyer members increase their willingness to pay by up to 8% or 2 cents per mile as their switching costs peak. As we will show in the next section, much of this increase is associated with business travel. Conservative calculations suggest an \$820 million increase in travel expenditures each year due to people’s increased willingness to pay associated with status incentives.²³ Of course, the total impact is likely much larger since switching costs also impact the extensive margin of consumer choice. These switching costs are likely to be a significant source of market power for TA with regard to consumers who are motivated to attain status.

The previous literature discussing moral hazard in the context of loyalty programs has suggested that much of the success of these programs rests on the asymmetric incentives present in business travel. Given our estimates, we can provide an approximate answer to how much of the overall increase in the willingness to pay is associated with business versus leisure travel. For drawing appropriate comparisons, Table A.4 in the Appendix replicates our main findings and our heterogeneity results for the subsample roundtrip bookings of members with at least two leisure and two business bookings. To provide a back-of-the-envelope calculation, we compare estimates of β_6 (10%) from Column 2 of Table A.4, and μ_6 (7%) from Column 3 of Table A.5. In this subsample, 51% of the bookings are characterized as business trips. Therefore, the contribution of asymmetric incentives present in business travel to increased willingness to pay on bookings overall can be approximated as $3.57\% = 51\% * 7\%$ at this level of progress. This suggests that travelers not shouldering the ticket costs themselves accounts for about 36% ($\frac{3.57\%}{10\%}$) of the overall price inflation due to status incentives. This result indicates that status rewards indeed owe a large part of their success to asymmetric incentives present in business travel.

Of course, asymmetric incentives arising from the decision-making not shouldering the entire cost while benefiting from perks associated with purchase is not specific to the airline industry. Such asymmetries could be present in other purchase decisions that employees make and employers pay for, such as hotels, car rentals and office equipment, etc. Basso et al. (2009, page 117) state “we think it is no coincidence that these very

²³The three top legacy airlines in the U.S. generated revenues of 119 billion USD in 2017. Conservatively, we assume that only 1 in 5 passengers on these airlines were motivated to attain status. We used the estimated increase in the willing to pay of the average consumer in our data across the distribution of different progress levels (weighted average equal to 4.1%) to arrive at this number.

large loyalty programs exist in areas in which a large fraction of purchases are work related.” Our results provide a sense of the magnitude of the impact of asymmetric incentives inherent in business purchases.

Citing concerns of employees taking advantage of these asymmetric incentives, some governmental organizations have attempted to prevent their employees from accumulating miles on official travel.²⁴ That said, the fact that other companies could save money by restricting employee choice does not suggest that their current policies are suboptimal. We recognize that companies may view status perks (and other travel perks) as a side benefit to their employees. Our paper offers an empirical measurement of the amount the companies are currently forgoing to do so, which we approximate at 7% based on our estimates.²⁵

We should note that our estimates are potentially attenuated by measurement error in the progress metric. The reason for the potential measurement error in the progress metric is that a consumer who appears in the data as having a certain number of miles may behave as if they had far more because they know with near certainty that they will travel extensively with TA on a trip later in the year, even if they have not booked the trip yet. While this may average out with consumers who know that they cannot travel at the end of the year, a plausible concern remains that some consumers in reality have more miles than observed in the data. This would manifest itself in our progress metric as a non-negative measurement error that could influence our estimates in two distinct ways. First, in contrast to classical measurement error, a non-negative measurement error causes the error term in our regressions to have nonzero mean. The sign of that mean is ambiguous and determined by the interaction of the slope parameters and the expected measurement error. Regardless of the sign, this mean is soaked up by our fixed effects, which are uniformly pushed in the direction of the mean of the error term to recenter it zero. Second, because the measurement error now has mean zero, our estimates of the β_k are attenuated exactly as in a classical measurement error problem. While the effect on the estimation of the conditional mean function is therefore ambiguous, the estimates of the marginal effects are unambiguously attenuated by the measurement error. Hence, in the presence of this type of consumer heterogeneity, one can interpret our estimates of the β_k as lower bounds. Still, if this attenuation bias exists, we do not expect it to be very large because people do not wait very long to book the trips they are near certain they will be taking. Appendix [A.6](#) discusses this point in more detail and presents empirical evidence from our data and a supplemental survey.

In all our analyses, we use the progress metric to account for the role of timing of point accumulation. However, it is also instructive to evaluate the purchase behavior with regard to the total point accumulation

²⁴Frequent-flyer Program. In Wikipedia, retrieved Sept 1, 2018.

²⁵Since the employer pays all of the business trips, from their perspective, the asymmetric incentives account for 68% of increased expenditures (a 10% increase) due to their employees’ willingness to spend more to attain status.

as some of the prior work has done. We present this analysis in Appendix [A.7](#). This analysis allows us to observe the extent to which the individual's pattern of behaviors reverse from the end of one point earning cycle to the beginning of another. The results are reported in Appendix Table [A.3](#). There are two main takeaways. First, these results parallel those based on progress: Consumers become more willing to book trips with the airline at higher prices and on markets where the airline does not offer the most appealing flight options as they get closer to achieving their goal. Second, consumers display significant resets in their behavior, going from the end of one point earning cycle to the beginning of another. This result suggests that the increase in switching costs reverses once the point accumulation period refreshes, which is consistent with the actions of forward-looking consumers. Importantly, this result adds further support to the notion that our findings are not driven by changes in consumer behavior due to spurious serial correlation or state dependency arising from cumulative experience or habit formation with the airline.

The reader may wonder about the role of heterogeneity in consumers' yearly travel expectations at the beginning of a point-earning cycle. Some consumers generally travel more on average, while others know that they will have only a few possibilities to earn points towards status. In Appendix [A.8](#) we theoretically explore how differences in travel expectations may change optimal switching costs as a function of progress within the context of our model and plot simulation results in Figure [A.4](#). We find that when expectations of total travel are higher, consumer switching costs are substantially higher for low levels of progress. The intuition is that an individual who expects more point earning opportunities in the year perceives a higher probability of attaining status even at low point accumulation levels, therefore has stronger incentives to forgo current utility for the potential of attaining status. However, at higher levels of progress, the impact of expectations of future travel are mostly negligible. In Figure [A.5](#), we depict the results of our empirical examination of this heterogeneity using a consumer's 2010 travel as proxy for her 2011 expectations. We see the same pattern of response to progress across the different groups of consumers: switching costs increase in progress before slightly declining at very high levels of progress. However, switching costs are larger for low levels of progress for individuals who traveled more in the past and therefore are likely to travel more in the current year.

In Appendix [A.9](#), we present several robustness checks. The first set of analyses check for aggregation bias. In our main specifications, we bin progress into fine regions of progress to capture response nonlinearity. We recognize that there may be a concern that our results may suffer from an aggregation bias due to this fine binning. We check for robustness of our findings by 1) fitting a quadratic functional form (instead of binning)

to capture response to progress, 2) fitting a multi-level mixed effect model assuming quadratic responses on the progress metric to capture consumer heterogeneity more flexibly, and 3) reducing the number of progress intervals from 9 to 4 to reduce data fragmentation across bin. As results reported in Tables [A.6](#)-[A.8](#) show, our conclusions are robust to these different specifications²⁶

Finally, in Appendix Table [A.9](#) and Table [A.10](#), we report results from several replications using different member subsamples, using different ways of constructing statistics from DB1B, or after accounting for outliers and/or punch code errors in the data. Our results replicate across all of these robustness checks. In addition, we check for robustness of our results regarding asymmetric incentives in business travel by modifying specifications [\(12\)](#) and [\(13\)](#) such that segment level heterogeneity is specified continuously as $\beta_{sk} = \beta_k + \sigma_k BizPerc_i$ where $BizPerc_i$ is the proportion of trips a consumer takes that do not span a Saturday night. Table [A.11](#) reports the coefficient estimates. The conclusions remain unchanged. Overall, this body of evidence strongly suggests that consumers respond significantly to the dynamic incentives created by the status award.

6 Conclusion

Airlines may have several purposes for their frequent-flyer programs, including price discrimination, creating behavioral loyalty by increasing switching costs, and taking advantage of the asymmetric incentives between employers and business travelers. In this paper, we show that progress towards status attainment leads to significant increases in consumer switching costs. Conservative calculations suggest a \$820 million increase in travel expenditures each year due to people’s increased willingness to fly with the sponsor airline. We argue that asymmetric incentives created by business travel have a large impact on the extent to which members increase their willingness to pay to attain status. Our estimates suggest that about a third of the increase in prices paid in response to dynamic incentives created by status rewards arise from the additional willingness to pay consumers exhibit when traveling for business. These results lend novel empirical support to a large body of theoretical work on frequent-flyer programs that have highlighted these aspects, and elucidate the magnitude of the impact these aspects have on the marketplace.

While the financial expenditures driven by switching costs induced by status incentives already seem substantial, for three reasons, we are likely to be underestimating the total impact. First, the estimated increases in the willingness to pay do not completely reflect the financial impact of the increased willingness

²⁶We thank an anonymous referee for suggesting these checks.

to travel with the airline over other transportation options. The fact that we only observe bookings with TA precludes us from quantifying how many bookings would not be made with TA if it were not for members' desire to attain status. Second, our estimates are likely to be attenuated by measurement error in the progress metric as well as in the business trip classification, as discussed in Section 5.3. Third, and more fundamentally, we can only speak to how consumers respond to changes in progress. Clearly, having a goal of attaining status may have an overall impact on consumers' purchase behaviors, even at the beginning of a point-earning cycle. The overall impact can be estimated if the behavior of consumers before and after the introduction of these status incentives could be observed. We hope that future research will be able to remedy these shortcomings and measure the total impact of frequent-flyer programs in general, and status incentives in particular, on consumers' purchase behaviors.

We recognize that frequent-flyer members may also be motivated by psychological reasons. For example, the mere perceived progress hypothesis (Nunes and Dreze, 2006) and the goal gradient hypothesis (Kivetz et al., 2006) predict that consumer valuations increase as consumers make progress towards their status goal. Our results are consistent with these general predictions. However, a rational model of forward-looking consumers better reflects the totality of our findings because members' responses vary with not only the amount of points they accumulated, but also the timing of these accumulations. In particular, we find that travelers scale back their efforts when their progress is faster than the pace required to attain their goal. Surely, evidence in favor of dynamic incentives does not rule out additional psychological factors that may be present. Our empirical context does not allow us to separately identify such factors. We hope that our work inspires future research in this area.

Future research should also examine the overall customer satisfaction implications of reward programs that require customers to make a price or product sacrifices along the way. If consumers are time-consistent in their preferences, the fact that they are trading off current consumption utility for these status benefits suggests that the utility from receiving the status benefits is larger than the disutility of the sacrifice. Therefore, at the end of the day, we may not expect a negative backlash. However, expectations of quality and experience quality may differ in their impact on customer satisfaction. Furthermore, consumers may not be time-consistent in their preferences. Therefore, more research is needed to fully evaluate the customer satisfaction and long-term loyalty effects of such programs.

Finally, we would like to mention an important change in the status qualification requirements in this industry since the period of our data. Delta added minimum spending thresholds for elite status tiers in

January 2014, United in March 2015, and American Airlines in August, 2016. These changes may have been responses to the observation that incremental sales arising from status incentives are more likely to be on shorter and cheaper routes that gave travelers a relatively high boost in miles per dollar spent. We predict that this new policy will help airlines better take advantage of asymmetries present in business travel. Under the new regime, road warriors flying on the company's dime not only will keep diverting business from competitors to their preferred airline, but they will also be incentivized to book full-priced tickets and/or to unnecessarily delay/advance their bookings to meet the new spending requirements. Although we lack the necessary competitive bookings data to identify whether travelers in our time period are gaming booking timing in this manner, we hope that our research spurs follow-on work that studies this aspect of booking behaviors and documents the impact of the change in program rules on consumers, airlines, and employers.

References

- Basso, L. J., Clements, M. T., and Ross, T. W. (2009). Moral hazard and customer loyalty programs. *American Economic Journal: Microeconomics*, 1(1):101–123.
- Bhaskara, V. (2015). The economics of revenue-based frequent flyer programs. *Airways Magazine*. <https://airwaysmag.com/industry/analysis-the-economics-of-revenue-based-frequent-flyer-programs/>
- Bolton, R. N., Kannan, P., and Bramlett, M. D. (2000). Implications of loyalty program membership and service experiences for customer retention and value. *Journal of the Academy of Marketing Science*, 28(1):95–108.
- Borenstein, S. (1990). Airline mergers, airport dominance, and market power. *The American Economic Review*, 80(2):400–404.
- Borenstein, S. (1991). The dominant-firm advantage in multiproduct industries: Evidence from the us airlines. *The Quarterly Journal of Economics*, 106(4):1237–1266.
- Breugelmans, E., Bijmolt, T., Zhang, J., Basso, L., Dorotic, M., Kopalle, P., Minnema, A., Mijnlief, W. J., and Wunderlich, N. V. (2015). Advancing research on loyalty programs: A future research agenda. *Marketing Letters*, 26(2):127–139.
- Brons, M., Pels, E., Nijkamp, P., and Rietveld, P. (2002). Price elasticities of demand for passenger air travel: a meta-analysis. *Journal of Air Transport Management*, 8(3):165–175. Developments in Air Transport Economics.
- Caminal, R. and Matutes, C. (1990). Endogenous switching costs in a duopoly model. *International Journal of Industrial Organization*, 8(3):353–373.
- Chung, D. J., Steenburgh, T., and Sudhir, K. (2013). Do bonuses enhance sales productivity? a dynamic structural analysis of bonus-based compensation plans. *Marketing Science*, 33(2):165–187.
- Copeland, A. and Monnet, C. (2009). The welfare effects of incentive schemes. *The Review of Economic Studies*, 76(1):93–113.

- Dreze, X. and Hoch, S. J. (1998). Exploiting the installed base using cross-merchandising and category destination programs. *International Journal of Research in Marketing*, 15(5):459–471.
- Dreze, X. and Nunes, J. C. (2011). Recurring goals and learning: The impact of of successful reward attainment on purchase behavior. *Journal of Marketing Research*, 48(2):268–281.
- Gerardi, K. S. and Shapiro, A. H. (2009). Does competition reduce price dispersion? new evidence from the airline industry. *Journal of Political Economy*, 117(1):1–37.
- Gopalakrishnan, A., Jiang, Z., Nevskaya, Y., and Thomadsen, R. (2021). Can non-tiered customer loyalty programs be profitable? *Marketing Science*, forthcoming.
- Hartmann, W. R. and Viard, V. B. (2008). Do frequency reward programs create switching costs? a dynamic structural analysis of demand in a reward program. *Quantitative Marketing and Economics*, 6(2):109–137.
- Hess, S., Adler, T., and Polak, J. W. (2007). Modelling airport and airline choice behaviour with the use of stated preference survey data. *Transportation Research Part E: Logistics and Transportation Review*, 43(3):221–233.
- Joakim Karlsson, A. O. and Yamanaka, S. (2004). The impact of infrastructure-related taxes and fees on domestic airline fares in the u.s. *Journal of Air Transport Management*, 10(4):283–291.
- Kim, B.-D., Shi, M., and Srinivasan, K. (2001). Reward programs and tacit collusion. *Marketing Science*, 20(2):99–120.
- Kivetz, R., Urminsky, O., and Zheng, Y. (2006). The goal-gradient hypothesis resurrected: Purchase acceleration, illusionary goal progress, and customer retention. *Journal of Marketing Research*, 43(1):39–58.
- Klemperer, P. (1987). Markets with consumer switching costs. *The Quarterly Journal of Economics*, 102(2):375–394.
- Klemperer, P. (1995). Competition when consumers when have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade. *The Review of Economic Studies*, 62(4):515–539.
- Klemperer, P. and Png, I. (1986). Frequent flyer plans: Marketing device with insidious effects. *The Los Angeles Times*, 8(3).
- Kopalle, P. and Neslin, S. (2003). The economic viability of frequency reward programs in a strategic competitive environment. *Review of Marketing Science*, 1(1).
- Kopalle, P., Sun, Y., Neslin, S. A., and Sun, B. (2012). The joint sales impact of frequency reward and customer tier components of loyalty programs. *Marketing Science*, 31(2):216–235.
- Lal, R. and Bell, D. E. (2003). The impact of frequent shopper programs in grocery retailing. *Quantitative Marketing and Economics*, 1(2):179–202.
- Lederman, M. (2007). Do enhancements to loyalty programs affect demand? the impact of international frequent flyer partnerships on domestic airline demand. *The RAND Journal of Economics*, 38(4):1134–1158.
- Lederman, M. (2008). Are frequent-flyer programs a cause of the hub premium? *Journal of Economics and Management Strategy*, 17(1):35–66.
- Leenheer, J., van Heerde, H. J., Bijmolt, T. H., and Smidts, A. (2007). Do loyalty programs really enhance behavioral loyalty? an empirical analysis accounting for self-selecting. *International Journal of Research in Marketing*, 24(1):31–47.

- Levine, M. E. (1986). Airline competition in deregulated markets: Theory, firm strategy, and public policy. *Yale Journal*, 4(2):393–494.
- Lewis, M. (2004). The influence of loyalty programs and short-term promotions on customer retention. *Journal of Marketing Research*, 41(3):281–292.
- Mian Dai, Q. L. and Serfes, K. (2014). Is the effect of competition on price dispersion nonmonotonic? evidence from the u.s. airline industry. *Review of Economics and Statistics*, 96(1).
- Misra, S. and Nair, H. S. (2011). A structural model of sales-force compensation dynamics: Estimation and field implementation. *Quantitative Marketing and Economics*, 9(3):211–257.
- Nunes, J. C. and Dreze, X. (2006). The endowed progress effect: How artificial advancement increases effort. *Journal of Consumer Research*, 32(24):504–512.
- Proussaloglou, K. and Koppelman, F. (1995). Air carrier demand. *Transportation*, 22(4):371–388.
- Rossi, F. (2017). Lower price or higher reward? measuring the effect of consumers' preferences on reward programs. *Management Science*, 64(9):3971–4470.
- Shugan, S. M. (2005). Brand loyalty programs: Are they shams? *Marketing Science*, 24(2):185–304.
- Tretheway, M. W. (1989). Frequent flyer programs: Marketing bonanza or anti-competitive tool? *Journal of Transportation Research Forum*, 30(1):195–201.
- 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.
- Zhang, J. and Breugelmans, E. (2012). The impact of an item-based loyalty program on consumer purchase behavior. *Journal of Marketing Research*, 49(1):50–65.