

Understanding Consumer Dynamic Decision Making Under Competing Loyalty Programs

Journal of Marketing Research
2020, Vol. 57(3) 422-444
© American Marketing Association 2020
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/0022243720911894
journals.sagepub.com/home/mrj



Jia Liu  and Asim Ansari

Abstract

The authors develop an incentive-aligned experimental paradigm to study how consumer purchase dynamics are affected by the interplay between competing firms' loyalty programs and their pricing and promotional strategies. In this experiment, participants made sequential choices between two competing airlines in a stylized frequent traveler task for which an optimal dynamic decision policy can be numerically computed. The authors find that, on average, participants are able to partially realize the long-term benefits from loyalty programs, though most are sensitive to price. They also find that participants' preferences and levels of bounded rationality depend on the nature of the competitive environment, the particular state of each decision scenario, and the type of optimal action. Accordingly, the authors use an approximate dynamic programming model to incorporate boundedly rational decision making. The model classifies participants into five segments that exhibit variation in their performance and decision strategies. Importantly, they find that participants are able to adapt their decision strategies to the environment they face, and thus the overall market outcome and the performance of each firm are influenced by both the competitive environment and the assumption on the extent of consumer optimality.

Keywords

approximate dynamic programming, bounded rationality, competition, experiments, loyalty programs, promotions

Online supplement: <https://doi.org/10.1177/0022243720911894>

Firms routinely use promotions in the battle for market share. In recent years, loyalty programs, which play a prominent role in the promotional mix of firms, have become ubiquitous. Hotels, airlines, department stores, credit card companies, and even grocery stores and cafeterias use loyalty programs to promote and maintain consumer loyalty (Dowling and Uncles 1997). Recent statistics (Autry 2017) indicate that U.S. consumers held approximately 3.8 billion memberships in customer loyalty programs in 2017, a 15% increase from 2015. Furthermore, it is estimated that a consumer belongs to more than 14 loyalty programs, on average. The widespread availability of loyalty programs means that customer purchasing decisions are influenced by the interplay among the prices, promotions, and loyalty programs of competing firms. In turn, firms face the challenge of disentangling this interactive impact to make informed pricing and promotional decisions, especially in competitive contexts.

However, much of the empirical work on loyalty programs (which is reviewed in the next section) is in the context of a single loyalty program because data are often available from a single firm or from a single loyalty program that is shared by

multiple firms. Thus, little is known about customers' choices when they hold simultaneous membership in multiple loyalty programs of competing firms, and how these choices inform loyalty program design. While some researchers have tried to circumvent this challenge by supplementing individual-level data with aggregate data on competition (also reviewed in the next section), this limits the insights that one can derive about customer purchasing dynamics. We address this challenge by using an experimental approach.

Figure 1 lays out a road map to the entire article and our research focus. We develop an incentive-aligned, controlled, experimental paradigm to study how consumer purchase dynamics vary across market designs (i.e., competitive landscapes) involving different types of loyalty programs, pricing, and promotional strategies of two competing firms. We use the realistic setting of frequent travelers making sequential choices

Jia Liu is Assistant Professor of Marketing, Hong Kong University of Science and Technology, Hong Kong (email: jialiu@ust.hk). Asim Ansari is William T. Dillard Professor of Marketing, Marketing Division, Columbia Business School, Columbia University, USA (email: maa48@columbia.edu).

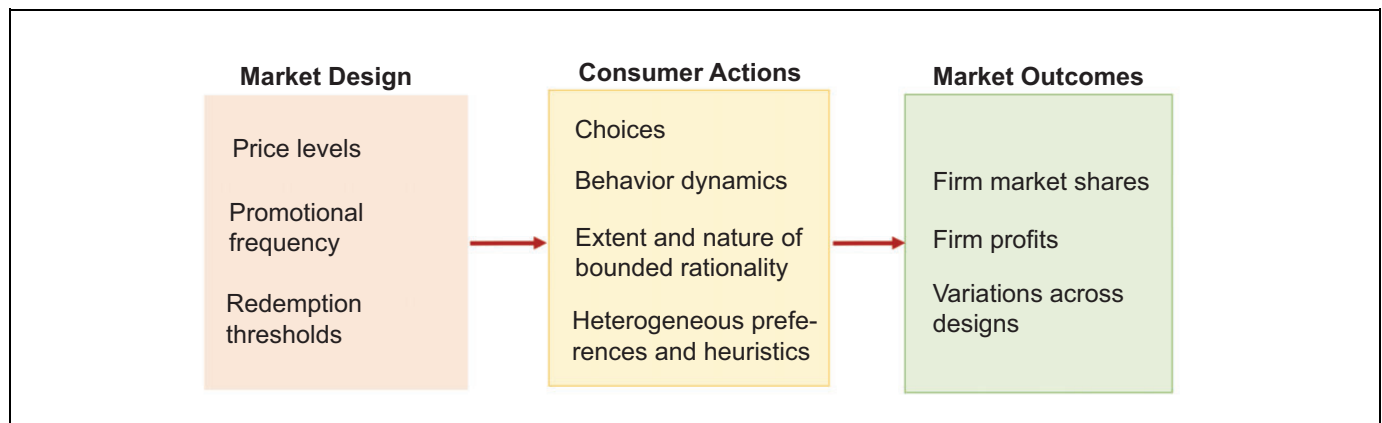


Figure 1. Road map.

between two competing airlines for a given number of trips. The finite period decision problem captures real-world dynamics that stem from consumer churn, or point expiry periods and membership renewal/expiration of reward programs. In our experiment, a participant is randomly assigned to one of eight markets that differ along three dimensions: (1) the loyalty program, in terms of the point threshold needed for a free ticket; (2) the pricing strategy, in terms of the sale price levels for the competing firms; and (3) the promotion strategy, in terms of the sales promotions frequency. As in natural settings, after accumulating enough points for a free ticket, participants can decide whether to redeem these for a trip. Our experiment is also incentive aligned, which is the state of the art in choice experiments and induces truth-telling in particular (Becker, DeGroot, and Marschak 1964; Ding 2007).

For any given set of design parameters for our duopoly market, the optimal strategy for consumers can be obtained by numerically solving their dynamic programming (DP) problem. This enables us to contrast the actual performance of participants with the performance under the optimal strategy and thus measure participants' degree of bounded rationality. In addition, our experimental paradigm enables us to investigate how consumers trade off immediate benefits stemming from sales promotions with the delayed, but potentially larger, benefits resulting from loyalty programs. In most loyalty programs, consumers have to accumulate enough program currency, such as stamps or points, to reach the redemption threshold. Such a commitment encourages consumers to consider the future consequences of their current purchase decisions (Ching and Ishihara 2017; Hartmann and Viard 2008; Kopalle et al. 2012; Lewis 2004) and may create a potential conflict with leveraging immediate benefits from sales promotions for consumers. In summary, we focus on the following questions: (1) How does the design of the loyalty programs and the nature of price competition affect consumer decision making?; (2) Do consumers realize the long-term benefits of program rewards as members of multiple competing programs?; (3) To what extent do consumers behave optimally in navigating competing programs, and what factors influence deviations

from optimality?; (4) What types of decision rules (e.g., heuristics) do consumers use when making sequential decisions?; and (5) Do consumers adjust their decision rules to the market environment? To the best of our knowledge, we are the first to study these research questions.

Answering these questions can further our understanding of how consumers navigate different offerings from competing firms and can provide insights that firms could use to design effective loyalty programs, based on the actual (i.e., boundedly rational, rather than optimal) behavior of customers within a competitive market. A comparison of consumer decisions across the eight markets enables us to understand how consumer purchase decisions as well as firm market shares and profits are affected by different combinations of redemption thresholds, prices, and the promotional strategies of the two firms. Moreover, the use of experimental control enables us to uncover the marginal impact of a specific factor while holding other factors constant, something that is clearly not possible in observational data, even if such data were available.

We analyze the sequential decisions of participants by building a dynamic choice model. However, instead of using a structural approach based on consumer rationality, we adopt a boundedly rational perspective using the methodology in Houser, Keane, and McCabe (2004) and Geweke and Keane (2000). This enables us to infer the degree of bounded rationality and the heterogeneity in the decision rules that consumers actually use in navigating competitive loyalty programs. This focus is consistent with the behavioral economics and marketing literature on reward programs that emphasizes descriptive outcomes (Ericson and Laibson 2019). It is also in line with the perspective in Meyer and Hutchinson (2016), which encourages researchers to design quasi-optimal models of consumer decision making. We augment this by studying how firm outcomes are affected by the actual behavior of participants as opposed to rational behavior.

Our key findings are as follows. First, we find that participants are not fully forward looking. Although they realize the delayed, but larger, benefits from the loyalty programs to a certain extent, their decisions systematically deviate from the

Table 1. Classification of the Loyalty Program Literature.

Approach	Single Program	Within-Program Competition	Cross-Program Competition
Descriptive or bounded rational	Kivetz, Urminsky, and Zheng (2006), Nunes and Drèze (2006), Shu and Gneezy (2010), Zhang and Breugelmans (2012), Stourm, Bradlow, and Fader (2015), Wang et al. (2016)	Evanschitzky et al. (2012), Pancras, Venkatesan, and Li (2015)	Liu and Yang (2009), this article
Perfectly rational	Lewis (2004), Hartmann and Viard (2008), Kopalle et al. (2012), Rossi (2018), Ching and Ishihara (2017)	N.A.	This article
Theoretical	Singh, Jain, and Krishnan (2008), Sun and Zhang (2019), Chun and Ovchinnikov (2019)	N.A.	Caminal and Matutes (1990), Kim, Shi, and Srinivasan (2001)

optimal. This finding adds to the body of psychological work on single loyalty programs (Evanschitzky et al. 2012; Kivetz, Urminsky, and Zheng 2006; Shu and Gneezy 2010). Specifically, we find that changes in redemption threshold have a more consistent impact on participants' decisions or preferences across contexts, when compared with the impact of promotion frequency and prices. We also find that trading off the immediate benefits from sales promotions with the delayed benefits resulting from loyalty programs is especially difficult for participants. This can leave participants worse off relative to under a myopic decision strategy that solely relies on observed prices. However, we also find that participants do better than under the simplest heuristic strategies.

Second, we find that, on average, participants underappreciate the firm with a higher regular price but a lower redemption threshold. However, this tendency varies across individuals and market designs. More generally, both the preferences of participants and their degree of optimality vary significantly across market designs. Even within a specific market design, the tendency to deviate from the optimal action is influenced by the particular state that is experienced on a decision point and by the nature of its corresponding optimal action. For example, we find that consumers are more likely to make mistakes when faced with free tickets from the nonoptimal airline, and they redeem the free tickets at the wrong occasions.

Third, we find considerable heterogeneity in the decision rules used by participants. We identify five segments that vary in their performance and decision strategies. Even though we find that all segments are present in each of the eight designs, the segment sizes vary across designs and are consistent with the optimal strategy under the design. For example, if the optimal strategy favors airline A under a certain design, then the size of the segment that mostly favors airline A also tends to be larger for the same design. This indicates that participants, to a certain extent, are able to adapt their decision strategies to the competitive environment they face. Finally, we show through a simulation that both the competitive environment and the degree of consumer optimality influence the revenues and costs of the firms and of the overall market. Therefore, assumptions about consumer decision strategies are critical in determining which competitive environment is most, or least, appealing for a firm.

The article is organized as follows. We first review the relevant literature. We then describe the sequential decision problem and the algorithm for computing the optimal strategy. This is followed by sections on the experimental study and the evidence for bounded rationality. We then describe the boundedly rational dynamic model, analyze the heterogeneity in decision rules, and present the results from our simulation study. Finally, we conclude with a discussion of the implications for researchers and practitioners.

Literature Review

Our focus is on studying how consumers balance prices, promotions, and reward points over time under competing loyalty programs, how consumers deviate from optimal behavior, and how this has implications for the design of competing loyalty programs. These research questions are relevant to multiple research streams that are summarized in Table 1 based on the research approach and the research context.

The first stream contains descriptive studies that focus on understanding the impact of loyalty program design on consumer purchasing behavior. Depending on the research context, a study can be classified into a single program, within-program competition, or cross-program competition. Most existing findings are within the single-program context and document how psychological mechanisms and boundedly rational behavior affect consumer actions. For instance, Nunes and Drèze (2006) and Kivetz, Urminsky, and Zheng (2006) show how "points pressure" propels consumers to increase their purchase amount or frequency as their accumulated points approach the reward threshold. Shu and Gneezy (2010) study the impact of deadlines on redemptions using a time discounting model of procrastination. Stourm, Bradlow, and Fader (2015) show that redemption behavior within a linear reward program is mostly driven by cognitive and psychological mechanisms, including mental accounting. Wang et al. (2016) use a field experiment to study the impact of success in obtaining a reward or failing to get one on subsequent purchase behavior of customers and showed that this effect depends on the loyalty status of members. More recent descriptive research examines within-program competition where multiple companies

participate in a common loyalty program through a single card (Danaher, Sajtos, and Danaher 2017; Evanschitzky et al. 2012; Gardete and Lattin 2018; Pancras, Venkatesan, and Li 2015; Stourm, Bradlow, and Fader 2017).

Researchers have sparingly studied the more common situation of cross-program competition, where consumers are members of different reward programs within an industry. Existing studies have only used aggregate or partial information about competitors. For example, Liu and Yang (2009) analyze firm-level sales data from the airline industry. Zhang and Breugelmans (2012) study the effect of competitive price discounts and reward point promotions across five major supermarkets, but use customer transaction and reward program data from only one supermarket. While these studies have yielded valuable insights, it is difficult to fully understand how competing loyalty programs affect purchasing when competitive variables are missing or are only available in aggregate form.

The second research stream focuses on dynamic decision making involving perfectly rational consumers and in the context of a single loyalty program. Several researchers have developed structural models involving consumers who have rational expectations and are able to optimally “solve” dynamic decision problems. In this stream, Hartmann and Viard (2008) study whether frequency-based reward programs create switching costs and find evidence of their heterogeneous impact. Kopalle et al. (2012) use purchase data from a single hotel chain and show that the reward program influences the purchase frequency of a segment of customers. Recently, Rossi (2018) developed a structural model for gasoline purchases and found that a substantial segment of consumers extracts more value from a dollar’s worth of rewards than from one dollar spent on gasoline. However, this research did not study competing loyalty programs, as it focused on a chain of stations belonging to a single oil firm. Finally, Ching and Ishihara (2017) explore the identification conditions needed for the estimation of discount factors in dynamic models and show that points-pressure effects can be explained by a rational dynamic decision-making model with state-dependent discount factors. Whereas their focus was theoretical in scope and concentrated on infinite-horizon stationary contexts, we study a finite-horizon nonstationary environment that does not rely on a discount factor.

The last stream of research uses theoretical models to analyze competitive market conditions that enable different types of symmetric and asymmetric equilibrium outcomes. Singh, Jain, and Krishnan (2008) study a duopoly market when only one firm offers loyalty program. More recently, Sun and Zhang (2019) investigated the impact of expiration terms on firm profits for a monopolist and Chun and Ovchinnikov (2019) considered program design for a monopolist, when customers are strategic. Theoretical work on cross-program competition includes Caminal and Matutes (1990) and Kim, Shi, and Srinivasan (2001), who study duopolies with both firms offering loyalty programs.

Sequential Decision Problem and Optimal Strategy

Competitive Landscape

We are interested in understanding the sequential purchasing decisions of consumers who hold membership in competing loyalty programs. We consider a basic competitive environment that captures the essence of more complex real-world contexts but is easier to explain to participants. We use a competitive scenario involving two airlines that differ in their price offerings and loyalty program design, and hold other aspects, such as quality, the same across firms. We use an airline industry context because almost all airlines compete by offering loyalty programs, and most consumers have experience using these loyalty programs. In our experiment, consumers make sequential choice decisions between two competing airlines, A and B, over a finite horizon involving a fixed number of trips, T . The finite number of trips relates to scenarios where the number of trips are determined based on some underlying need for repeated travel on a route and also captures real-life scenarios stemming from points or membership expiration. Having a fixed number of trips also allows us to compare purchasing behavior across participants and across market designs, and to assess participants’ deviations from the rationality.

The two airlines differ along two aspects: pricing strategy and loyalty program design. While in reality each firm can have multiple price levels, for simplicity, we assume a high–low price combination involving a regular price level and a sale price level, for both airlines. Specifically, we let (p_{Ar}, p_{As}) denote the regular and sale prices, respectively, for airline A, and (p_{Br}, p_{Bs}) represent the regular and sale prices for airline B. The uncertainty in the consumer decision problem arises from the two airlines choosing their price levels probabilistically over trips. We characterize this uncertainty with three parameters: γ_A is the probability that airline A chooses its promotional price, γ_B is the probability that airline B chooses its promotional price, and ρ is the probability that both airlines simultaneously use the same price level (i.e., both choose regular or promotional). These three parameters $\{\gamma_A, \gamma_B, \rho\}$ can uniquely determine the probability f_{ij} that a consumer will face the set of prices (p_{Ai}, p_{Bj}) from airlines A and B, on a given purchasing decision, where $i, j \in \{r, s\}$.¹ For each airline, consumers earn one reward point per dollar spent for buying tickets. Consumers can gain free tickets by redeeming these reward points. The two airlines differ in the redemption thresholds needed for a free ticket. Airline A requires rt_A points, while airline B requires rt_B points. Any unused points at the end of a time period roll over to the next period, except for points remaining unused at the end of the task, which are forfeited.

¹ Given that $\sum_{i,j} f_{ij} = 1$, and because $\gamma_A = f_{sr} + f_{ss}$, $\gamma_B = f_{rs} + f_{ss}$, and $\rho = f_{rr} + f_{ss}$, we have $f_{rr} = (1 + \rho - \gamma_A - \gamma_B)/2$, $f_{ss} = (\gamma_A + \gamma_B + \rho - 1)/2$, $f_{sr} = (1 - \rho + \gamma_A - \gamma_B)/2$, and $f_{rs} = (1 - \rho - \gamma_A + \gamma_B)/2$. These probabilities $\{f_{ij}\}$ are used to obtain the optimal decision strategy.

This can be understood as the consequence of the membership expiring after the last trip, or of consumer churn. The design of the competitive context can, therefore, be summarized using ten parameters in the set $\Theta = \{T, p_{Ar}, p_{As}, p_{Br}, p_{Bs}, \gamma_A, \gamma_B, \rho, r_{tA}, r_{tB}\}$.

Optimal Decision Strategy for Consumers

Consistent with real world scenarios, consumers in our experiment were given the objective of minimizing the total cost over all the T trips. The optimal strategy for this forward-looking task can be computed by solving a finite time-horizon DP problem. Under our market setup, the state variables that summarize the decision context faced by a consumer i at trip t include (1) a_{itA} and a_{itB} , which denote the number of accumulated reward points for the two airlines; (2) x_{itA} and x_{itB} , which represent the observed prices of the airlines; and (3) n_{it} , the number of trips remaining before the end of the task. The redemption thresholds and the distribution of prices are held fixed throughout the task and are therefore not part of the state variables. In addition, the availability of a free ticket from any airline on a purchase occasion can be directly inferred from the points inventory, so it is not included in the state variables. Therefore, the set of state variables useful for forecasting future utility is $\mathcal{I}_{it} = \{a_{itA}, a_{itB}, x_{itA}, x_{itB}, n_{it}\}$. The choice set of alternatives, $\mathcal{A}(\mathcal{I}_{it})$, varies across trips because a free ticket is available only when the point inventory reaches the redemption threshold.

The choice d_{it} made by consumer i at trip t is given by

$$d_{it} = \begin{cases} 1, & \text{purchase ticket from airline A,} \\ 2, & \text{purchase ticket from airline B,} \\ 3, & \text{use the free ticket by redeeming points from airline A,} \\ 4, & \text{use the free ticket by redeeming points from airline B.} \end{cases} \quad (1)$$

Given the state variables in \mathcal{I}_{it} , consumer i 's value function for alternative $j \in \mathcal{A}(\mathcal{I}_{it})$ at time period $t \in \{1, 2, \dots, T\}$ can be written as $V_j(\mathcal{I}_{it}) = c_{itj} + EV(\mathcal{I}_{i,t+1} | \mathcal{I}_{it}, j)$. Here, c_{itj} is the current period cost of alternative j , and $EV(\mathcal{I}_{i,t+1} | \mathcal{I}_{it}, j)$ is the expected future cost function that captures the consumer's beliefs about her expected optimal future cost, given her current state and choice j . The uncertainty stems from the fact that the prices for future trips are unknown to the consumer at any point in time.

In a rational model of dynamic decision making, researchers usually assume that the decision maker knows the functional form of $EV(\cdot)$ and has correct beliefs regarding the transition dynamics of the state variables in \mathcal{I}_{it} . The optimal decision rule is to choose the alternative j with the lowest value, $V_j(\mathcal{I}_{it})$, among all alternatives. The Bellman equation of this finite time horizon DP problem can therefore be written as

$$V_t(\mathcal{I}_{it}) = \min_{j \in \mathcal{A}(\mathcal{I}_{it})} \{c_{itj} + EV_{t+1}(\mathcal{I}_{i,t+1} | \mathcal{I}_{it}, j)\}. \quad (2)$$

Note that we assume here that the discount factor is one. That is because the payoffs are only realized at the end of the game; hence discounting is irrelevant in our game setup. Given the finite horizon, we used backward induction to solve for the optimal strategy, as described in Web Appendix A. Unfortunately, the optimal strategy is not available in an analytical form, but it can be computed numerically for any given set of parameters.²

Empirical Study

We now describe our experimental design, the choice of our design parameters, Θ , and our data collection procedure from Amazon's Mechanical Turk (MTurk). Even though our controlled experiment lacks the full realism of secondary data or a natural field study, we believe that it offers many advantages. First, we are able to obtain data on consumer purchasing activity under competing loyalty programs. In contrast, real world purchasing data from competing loyalty programs are not readily available. Second, we are able to control for many other factors that influence purchasing in the real world, such as advertising and branding, and can therefore isolate the impact of our treatment variables. Third, we are able to manipulate the competitive landscape to study consumer dynamics under different market structures, something that is not possible in natural environments. Finally, we are able to assess the optimality of consumer behavior, as the optimal strategy is computable under different competitive scenarios.

Game Specification

We set airline A's regular price to be higher than that of B and A's reward redemption threshold to be lower than that of B (i.e., $p_{Ar} > p_{Br}$ and $r_{tA} < r_{tB}$). This comes with no loss of generality, as preferences are always relative and we want to avoid situations in which one firm is clearly dominant. In this setup, consumers face a trade-off between a delayed benefit arising from the more attractive loyalty program of airline A and an immediate benefit stemming from the lower regular price of airline B. Specifically, we fix $p_{Ar} = 480$ and $p_{Br} = 440$, which means that airline B's regular price is approximately 9% cheaper than A's; we also fix $r_{tB} = 2,400$ points, which means that at least six purchases are required for getting a free ticket reward from airline B. To make the setting realistic for our participants, we chose these values based on the average ticket price and required spending for free tickets in the U.S. airline industry (Airlines.org 2017; Seaney 2013). We also fix the number of trips $T = 24$, which is sufficient to allow

² When solving for the optimal policy using many different combinations of design parameters, we find that when a free ticket is not available, the optimal action is more likely to be either the choice of the airline with the lower observed price or the higher point inventory. In contrast, when a free ticket is available, the optimal action is more likely to involve using the free ticket, if only a few trips are remaining, or if both airlines are charging their regular prices.

Table 2. Parameterization of Six Designs.

Design	RP Threshold		Pricing Strategy (\$)				Promotion		Expected Price (\$)			
	rt_A	rt_B	p_{Ar}	p_{As}	p_{Br}	p_{Bs}	γ	ρ	$E(p_A)$	$E(p_B)$	$\frac{rt_A}{E(p_A)}$	$\frac{rt_B}{E(p_B)}$
1	1,840	2,400	480	360	440	400	.3	.7	444	428	4.1	5.6
2	1,840	2,400	480	360	440	400	.4	.7	432	424	4.3	5.7
3	1,840	2,400	480	400	440	360	.3	.7	456	416	4.0	5.8
4	1,840	2,400	480	400	440	360	.4	.7	448	408	4.1	5.9
5	2,000	2,400	480	360	440	400	.3	.7	444	428	4.5	5.6
6	2,000	2,400	480	360	440	400	.4	.7	432	424	4.6	5.7
7	2,000	2,400	480	400	440	360	.3	.7	456	416	4.4	5.8
8	2,000	2,400	480	400	440	360	.4	.7	448	408	4.5	5.9

Notes: RP = reward program. The first eight parameters describe the competitive environment of a duopoly market with firms A and B. Specifically, rt_i is the number of points that firm i requires for a reward, p_{ir} represents firm i 's regular price, and p_{is} represents firm i 's sale price, where $i \in \{A, B\}$. The probability $\gamma = \gamma_A = \gamma_B$ denotes the probability that each airline uses a discounted price and ρ denotes the probability that the two airlines offer the same price level. The last four columns are the expected prices and the expected numbers of purchases required for a reward (i.e., a free airline ticket) for the two firms.

participants three or four opportunities to redeem their points but is not so large as to induce fatigue or boredom. In addition, for the parameters that capture the price uncertainty across trips, we fix $\rho = .7$ (i.e., the two airlines offer the same price level 70% of the time). This captures the effect of common factors such as seasonality and industry norms. For simplicity we also set $\gamma_A = \gamma_B = \gamma$, which implies that the two airlines have the same probability of offering their promotional price on any given trip.

We then vary the focal parameters $\{rt_A, \gamma, p_{As}, p_{Bs}\}$ across the market designs to reflect different competitive environments that imply different optimal strategies. In particular, we create eight scenarios using a $2 \times 2 \times 2$ full factorial design. The first factor is airline A's point redemption threshold $rt_A \in \{1,840, 2,000\}$, which is either 400 or 560 points lower than that of airline B. The second factor is the probability of running a sales promotion $\gamma \in \{.3, .4\}$. When $\gamma = .3$, the realized promotional distribution is $(f_{rr}, f_{rs}, f_{sr}, f_{ss}) = (.55, .15, .15, .15)$; when $\gamma = .4$, the distribution becomes $(.45, .15, .15, .25)$. The third factor represents the promotional prices of the two airlines, such that $(p_{As}, p_{Bs}) \in \{(400, 360), (360, 400)\}$. We combine p_{As} and p_{Bs} into a single factor to avoid situations wherein customers may be indifferent between the two airlines in the optimal decision strategy. Given that $p_{Ar} = 480$ and $p_{Br} = 440$, the tuple $(p_{As}, p_{Bs}) = (400, 360)$ means that both the regular and the discounted prices of airline A are higher than the corresponding prices for B. The level $(p_{As}, p_{Bs}) = (360, 400)$ implies that airline A has a higher regular price level but a lower promotional price compared with airline B.

Table 2 presents the exact parameter combinations of the eight designs, along with the resulting expected price of each airline and the expected number of purchases required for getting a reward from each airline. For ease of illustration, we use D_k to denote design k throughout the article. The differences across these designs lead to different optimal strategies. For instance, the optimal choice share of airline A decreases from 70% in D_1 to 20% in D_8 (based on 1,000 simulated price

sequences for each design). Therefore, the extent to which participants need to be forward looking to realize the superiority of A and to improve their payoff in the task varies across designs. Our experimental design enables us to measure the impact of pricing, promotion, and reward redemption thresholds on consumer choices and preferences by comparing the results across multiple designs. We now describe our experimental setup in greater detail.

Experiment

In our experiment, each participant first played the game and then answered a short survey about their experience. We offered two types of performance-based bonus payments, in addition to a \$1 participation fee to motivate participants in exerting sufficient effort. Every participant started with \$9,100 of fictional money in an account that they could use to buy tickets over the game. After finishing the 24 trips, every \$5 left in the account was converted to a \$.01 cash bonus. In addition, participants were told that the top three performers would each receive a \$10 prize. Participants played a practice game that was the same as the actual game before playing the actual for the bonus.³

Each participant was randomly assigned to one design. Participants first read the instructions and then proceeded to a page that summarized all the market information. The Appendix, for example, shows the instructions and the information presented to D_7 participants. The tables in the Appendix that show the prices, reward programs, and the current point inventories for

³ Our base rate, bonus amount, and prizes were determined on the basis of (1) prior research on the use of MTurk for experimental studies (Hara et al. 2018) and (2) the responses of 50 participants in a pilot study that was run before the actual study. In the pilot study, participants played the same game but without bonus payments. At the end, they were asked how much bonus amount would be reasonable as an incentive for them to exert sufficient effort in the game. Our pilot study participants also indicated that one practice game (with 24 trips) was enough to understand the game. Note that participants involved in the pilot study were not allowed to participate in the study with bonus payment.

the two airlines were available to participants throughout the game to assist their decision making. Thus, all state variables were known to our participants at every trip. After making a choice, and before proceeding to the next trip, participants viewed the cost for the current trip, the accumulated total spending, and the current account balance. At the end, participants filled a short survey that assessed (1) whether they understood the main differences between the two airlines correctly, (2) whether our game instructions were clear, (3) whether they thought that the game was realistic, and (4) their overall preference between the two airlines. Web Appendix B shows the actual questions used in the survey.

We recruited 1,012 participants using MTurk. Our analysis is based on 836 participants who answered at least two of the three quiz questions (see Table S1 of Web Appendix B) correctly after the game and were therefore able to recall the main differences between the two airlines. This results in about 100 participants for each design. On average, each participant spent about nine minutes playing the game (with a standard deviation of four minutes) and about three minutes to finish the survey questions (with a standard deviation of seven minutes). The average bonus amount that participants won was \$.77 (with a standard deviation of \$.50). The survey responses indicated that our participants understood the game instructions well and perceived the game setup to be realistic. We also found a large variation in participants' self-reported decision strategies, both within and across the market designs. For supporting results, see Web Appendix B.

Exploratory Evidence for Bounded Rationality

Here, we showcase some preliminary model-free evidence regarding the extent of bounded rationality. This analysis forms the basis for a more formal model-based approach described later in the article. We focus on the choices made by our participants. We first compare the actual performance of participants to that under four benchmark strategies to understand the extent to which these choices are consistent with rational play. Second, we contrast the preferences and the performance of participants across the eight market designs to understand how participants respond to marketing strategies at the aggregate level. Third, we look at the pattern of choices over time to investigate the nature of choice dynamics. Finally, we study the factors that could underlie any deviations from optimality.

Benchmark Strategies

The first benchmark is the "optimal" strategy, which is defined in two ways, depending on the research objective. The first optimal strategy is called global optimality (GO). It refers to the unique optimal solution across all the trips throughout the entire game. The second optimal strategy, which we call local optimality (LO), refers to choosing the forward-looking optimal choice on a given trip, based on its current state, taking into account what is optimal to do from the current trip onward.

That is, for a sequence of 24 realized price pairs for each participant, GO yields a unique choice sequence, whereas LO is specific to each trip given the participant's current point inventory.⁴ If a participant is locally optimal at every trip, then (s)he is also globally optimal; LO is therefore a less stringent definition of optimality. We examine GO to evaluate the overall outcome of a participant's decisions in the game, such as total cost and choice share of airline A. In contrast, we examine LO to evaluate a participant's level of rationality at any point in the game.

The second benchmark is the myopic strategy, which implies always choosing the option that minimizes the cost for the current trip. Thus, consumers who use a myopic strategy focus on the immediate cost and are insensitive to both loyalty programs. In terms of the Bellman Equation 2, a myopic choice results from ignoring the future utility component. The last two benchmarks are naive strategies that involve always purchasing from the same airline across all the trips (i.e., always choosing A or B). In these two naive strategies, we assume that a free ticket is redeemed on the chosen airline only when the regular price is offered on it, unless it is the last trip of the game, in which case the free ticket is always chosen. For simplicity, we label these strategies as naive A and naive B. A comparison of the actual choices with those implied by the myopic or naive strategies can inform us about the conditions in which participants tend to mimic these simple decision strategies.

Aggregate-Level Analysis

We begin with an aggregate-level analysis to understand at a high level how our participants chose compared with the aforementioned benchmark strategies. Table 3 reports the average behavior of participants within each design on three metrics. The first set of results pertains to the percentage difference between the cost of a benchmark strategy and the participant's actual overall cost. This is computed using the same set of exogenously realized price sequences across all strategies. As expected, we find that the globally optimal strategy always yields a significantly lower total cost than that under the actual strategy ($p < .001$), for all designs. Interestingly, the same is also true for the myopic strategy, except for Design D₁, where there is no significant difference ($p = .31$). The naive A strategy is significantly more costly ($p < .001$) than the actual strategy of participants under four designs {D₃, D₄, D₇, D₈}. Finally, the cost for the naive B strategy was either significantly higher or lower than the actual cost ($p < .005$), depending on the design. These observations suggest that participants are neither fully optimal nor completely naive, and a myopic strategy would have been more beneficial for participants compared with their actual boundedly rational choices.

⁴ For example, on the first trip, the locally optimal choice is the same as the globally optimal one, as the starting point inventory is zero in both cases; but they may be different later on, as the participant's actual point inventory may deviate from the point inventory under GO behavior (e.g., the participant's first choice is not the optimal one).

Table 3. Average Performance Across Designs.

Design	Participants	% Difference in Cost from Actual				A's Choice Share (%)			Participants' Deviation (%)			
		GO	Myopic	Naive A	Naive B	Participant	GO	Myopic	LO	Myopic	Naive A	Naive B
1	101	-4.06 (3.80)	-.30 (4.09)	.02 (3.96)	2.03 (4.52)	49.50 (27.27)	69.30 (9.07)	27.97 (11.49)	40.68 (12.12)	37.00 (22.69)	57.30 (24.74)	60.27 (22.12)
2	110	-4.74 (3.58)	-1.11 (3.82)	-1.18 (4.19)	2.96 (4.21)	53.07 (26.59)	72.15 (5.10)	37.69 (10.85)	40.93 (14.82)	38.45 (22.75)	56.55 (21.38)	62.50 (21.37)
3	106	-3.65 (3.80)	-1.60 (3.93)	2.50 (4.33)	-1.58 (3.89)	32.27 (31.05)	58.31 (14.59)	11.56 (7.61)	43.95 (14.35)	40.57 (24.76)	73.98 (25.55)	42.22 (27.25)
4	103	-3.63 (3.85)	-1.29 (4.61)	2.28 (4.37)	-1.88 (3.99)	32.77 (31.09)	63.83 (6.88)	13.59 (7.96)	36.27 (13.91)	42.44 (23.05)	73.06 (26.22)	42.68 (28.66)
5	103	-3.22 (3.04)	-1.67 (3.19)	.11 (3.66)	1.51 (3.62)	50.85 (31.53)	70.05 (8.80)	28.68 (10.38)	30.72 (15.26)	41.26 (25.32)	57.36 (26.06)	59.95 (26.27)
6	102	-3.91 (4.00)	-1.96 (4.83)	.28 (4.60)	1.72 (4.86)	50.20 (28.41)	41.91 (12.28)	36.48 (10.93)	27.82 (20.44)	41.18 (23.65)	56.66 (24.96)	60.17 (23.77)
7	111	-2.70 (3.49)	-1.02 (4.15)	3.22 (4.16)	-1.39 (3.69)	23.24 (25.50)	15.84 (10.06)	13.03 (7.43)	24.66 (20.73)	33.15 (20.73)	81.76 (18.90)	34.72 (23.66)
8	100	-2.87 (3.01)	-1.00 (3.60)	4.15 (4.00)	-1.12 (3.43)	24.67 (25.53)	17.12 (9.94)	14.42 (8.59)	25.25 (18.51)	34.00 (20.00)	78.54 (22.92)	36.38 (22.28)
All	836	-3.60 (3.63)	-1.24 (4.07)	1.42 (4.51)	.28 (4.43)	39.51 (30.70)	49.63 (23.67)	22.90 (14.03)	33.95 (17.25)	38.48 (23.07)	66.98 (25.94)	49.81 (26.89)

Notes: The values in parentheses are the standard deviation across participants from a given design. Myopic strategy involves always choosing the option that minimizes the cost for the current trip; Naive A (B) means always choosing from airline A (B).

The second set of columns in Table 3 reports the average and the standard deviation (in parentheses) of the choice share of airline A under the different strategies. These statistics are computed based on all the observed sequences of prices for a given design. We omit naive A and naive B here, as the choice share is always one or zero for these strategies. We can see that A's average choice share under the optimal strategy (denoted as M_k , for the k th design) ranges from 15% to 72%. It varies significantly across the designs such that $M_1 \approx M_2 \approx M_5 > M_3 \approx M_4 > M_6 > M_7 \approx M_8$. In contrast, under the myopic strategy, the overall choice share for A is much smaller and is below 40% in all designs. We also see that the average choice share of A under the participants' actual strategies almost always falls between that for the myopic and optimal strategies, except in D_6 , D_7 , and D_8 , where participants overchoose A. Importantly, the actual choice share is never close to 100% or 0%. These findings also suggest that participants are neither completely myopic/naive nor fully forward looking. They therefore exhibit boundedly rational behavior that is more sophisticated than that implied by simple heuristics.

The last set of four columns in the table report the extent to which the actual choices of the participants within a design deviate from those under the different benchmark strategies across the 24 trips. Specifically, we compute the percentage of the 24 trips in which the actual choice deviates from the choice under a benchmark strategy. Thus, a higher percentage indicates smaller alignment between a participant's actual decision and that of a benchmark strategy. Aggregating across all designs, we can order the strategies as follows: naive A > naive B > myopic > optimal in terms of their deviation from the actual choices. The averages, however, vary considerably

across designs and benchmark strategies, ranging from 25% to 82%. These findings suggest that even though, on average, our participants' actual decision strategies are closer to the optimal behavior than to the other three benchmarks, this tendency varies significantly across designs.

The Impact of Different Design Factors

We chose our design parameters to enable us to isolate the impact of four experimental factors (i.e., promotion probability, promotional prices, redemption threshold, and long-term vs. short-term reward) on the participants' choices, and their deviations from optimal behavior, by comparing different design pairs. We now use Table 3 to investigate the nature and magnitude of this impact.

Promotion probability. The four design pairs, (D_1, D_2), (D_3, D_4), (D_5, D_6), and (D_7, D_8), can be used to study the impact of promotion probability. Table 2 shows that the two designs within each of these pairs differ in their promotion probabilities (.3 for the first design and .4 for the second) but share the same values on all other design factors. Focusing on the last set of columns in Table 3, we see that, of these four pairs, only the (D_3, D_4) pair shows a significant difference ($p < .001$) in the percentage of trips on which participants deviate from the locally optimal choice across its two designs. This suggests that the impact of the promotion probability is context dependent.

Promotional prices. We also have four pairs whose members differ only in their promotional prices: (D_1, D_3), (D_2, D_4), (D_5, D_7), and (D_6, D_8). Comparing the outcomes associated

with these pairs enables us to study the impact of changing the promotional prices (p_{As}, p_{Bs}) from (360, 400) to (400, 360). We see that the participants' deviation from the optimal behavior differs significantly ($p < .01$) only within one pair (D_5, D_7). The other three pairs show no significant differences ($p > .1$). These results also suggest that, depending on the context, promotional prices may have some impact on the boundedly rational behavior of participants.

Redemption threshold. We have four design pairs whose members differ only in airline A's point redemption threshold: (D_1, D_5), (D_2, D_6), (D_3, D_7), and (D_4, D_8). Within each pair, A's redemption threshold is 1,840 for the first design and 2,000 points for the second. Not surprisingly, for all these pairs, the optimal solution selects A significantly more frequently in the first design (i.e., when the point threshold is lower) than in the second. The differences in the actual choice share of A between the two designs within a pair are, however, not that high. This implies that participants often fail to realize the long-term benefit that A offers through its lowered redemption threshold. This is also consistent with our second observation that is based on the LO column in Table 3. We can see that the average percentage of trips on which participants deviate from the LO choices is significantly higher ($p < .01$) for the first design within each pair (i.e., when A has the lower redemption threshold). Thus, we can conclude that the redemption threshold has a consistent impact on the decision strategies of participants, across contexts, and that slightly lowering the redemption threshold may not be effective in increasing A's market share.

Long-term versus short-term reward. We have one pair (D_4, D_6) that allows us to study how participants trade off prices and loyalty program benefits. Compared with D_6 , in D_4 , airline A offers a lower redemption threshold but charges a higher sale price, and its optimal strategy picks airline A more frequently. Thus, it is more difficult for participants to realize airline A's superiority under D_4 , as it imposes a stronger tendency to trade off the long-term benefit from A's loyalty program with the short-term gain from B's lower price. Indeed, we find that participants deviate significantly more from the optimal under D_4 than under D_6 . That is because they fail to realize the "better" delayed rewards from A and place too much weight on the immediate saving from the lower promotional price of B. As a result, the actual choice share of A in D_4 is much lower than that under the optimal strategy (compared with being slightly above the optimal in D_6). Therefore, we can infer, in general, that offering a loyalty program with a significantly lower points threshold, but with very high prices may not be an effective strategy for increasing market share, when taking into account the actual behavior of customers—even though optimal behavior may suggest otherwise.

Behavior Dynamics over Trips

Next, we look at the behavior across the 24 trips to understand the preference dynamics and the temporal pattern of the

deviations from optimality. Figure 2 shows the proportion of the participants who deviate from each of the four benchmark strategies (i.e., the LO, myopic, naive A, and naive B) at each trip, under each design. We find variations across the designs that are consistent with the results in Table 3. Interestingly, we can see that the overall pattern of results is similar within the pairs (D_1, D_2), (D_3, D_4), (D_5, D_6), and (D_7, D_8) but differs across pairs. The two designs within a pair differ only in their promotion probability γ , while the remaining parameters are the same. This observation again suggests that our participants are not very sensitive to changes in the promotion frequency. In addition, one may notice that their deviation levels from naive B and the myopic strategy are very similar for designs D_3, D_4, D_7 and D_8 , which are markets in which airline A's regular/sale price is higher than that of B. So probabilistically, the myopic strategy should mostly choose B as well, and is thus very similar to naive B.

Furthermore, when comparing actual choices with optimal behavior (the solid line) in Figure 2, we see that the deviation from the optimal decreases significantly toward the end of the game for the first five designs (i.e., from D_1 to D_5 ; $p < .001$). However, there is no significant time trend for the remaining (i.e., D_6 to D_8 ; $p > .5$). We also see that the deviation level from the myopic strategy (the dashed line) increases significantly toward the end of the game, for all the designs ($p < .05$). These results can be explained by the fact that it is easier for participants to be forward looking or optimal towards the end of the game. In addition, one can see that the benchmark strategy from which the participants deviate the most, or least, depends both on the design and the trip.

Figure 3 shows, for each design, the choice share of airline A across the 24 trips under the actual, globally optimal, and myopic strategies. One can discern a number of interesting patterns from these plots. We see that airline A's choice share under the optimal strategy tends to decrease as the design number increases, which is consistent with the aggregate results in Table 3. The same patterns is also observed in the actual decisions of the participants. This observation suggests that across the designs the actual behavior of our participants is directionally aligned with the optimal behavior. In addition, for the same reason as for Figure 2, we find that the actual choice shares are similar within the pairs (D_1, D_2), (D_3, D_4), (D_5, D_6), and (D_7, D_8). Furthermore, for the first five designs, D_1 to D_5 , participants tend to underchoose A relative to the optimal strategy, especially at the beginning of the game. However, for the remaining designs (i.e., D_6 to D_8), participants tend to overchoose A slightly, compared with that under the optimal strategy. Participants overall have a slightly higher preference for A, in all cases, compared with the choices implied by the myopic strategy. This gap, however, appears to diminish with the design number. This observation implies that participants are not completely myopic in their decisions and are therefore not entirely insensitive to the loyalty programs. Finally, we can see that their preferences are quite stable over the trips, even though the choice share of A decreases over time under the optimal strategy for designs D_1 to D_5 .

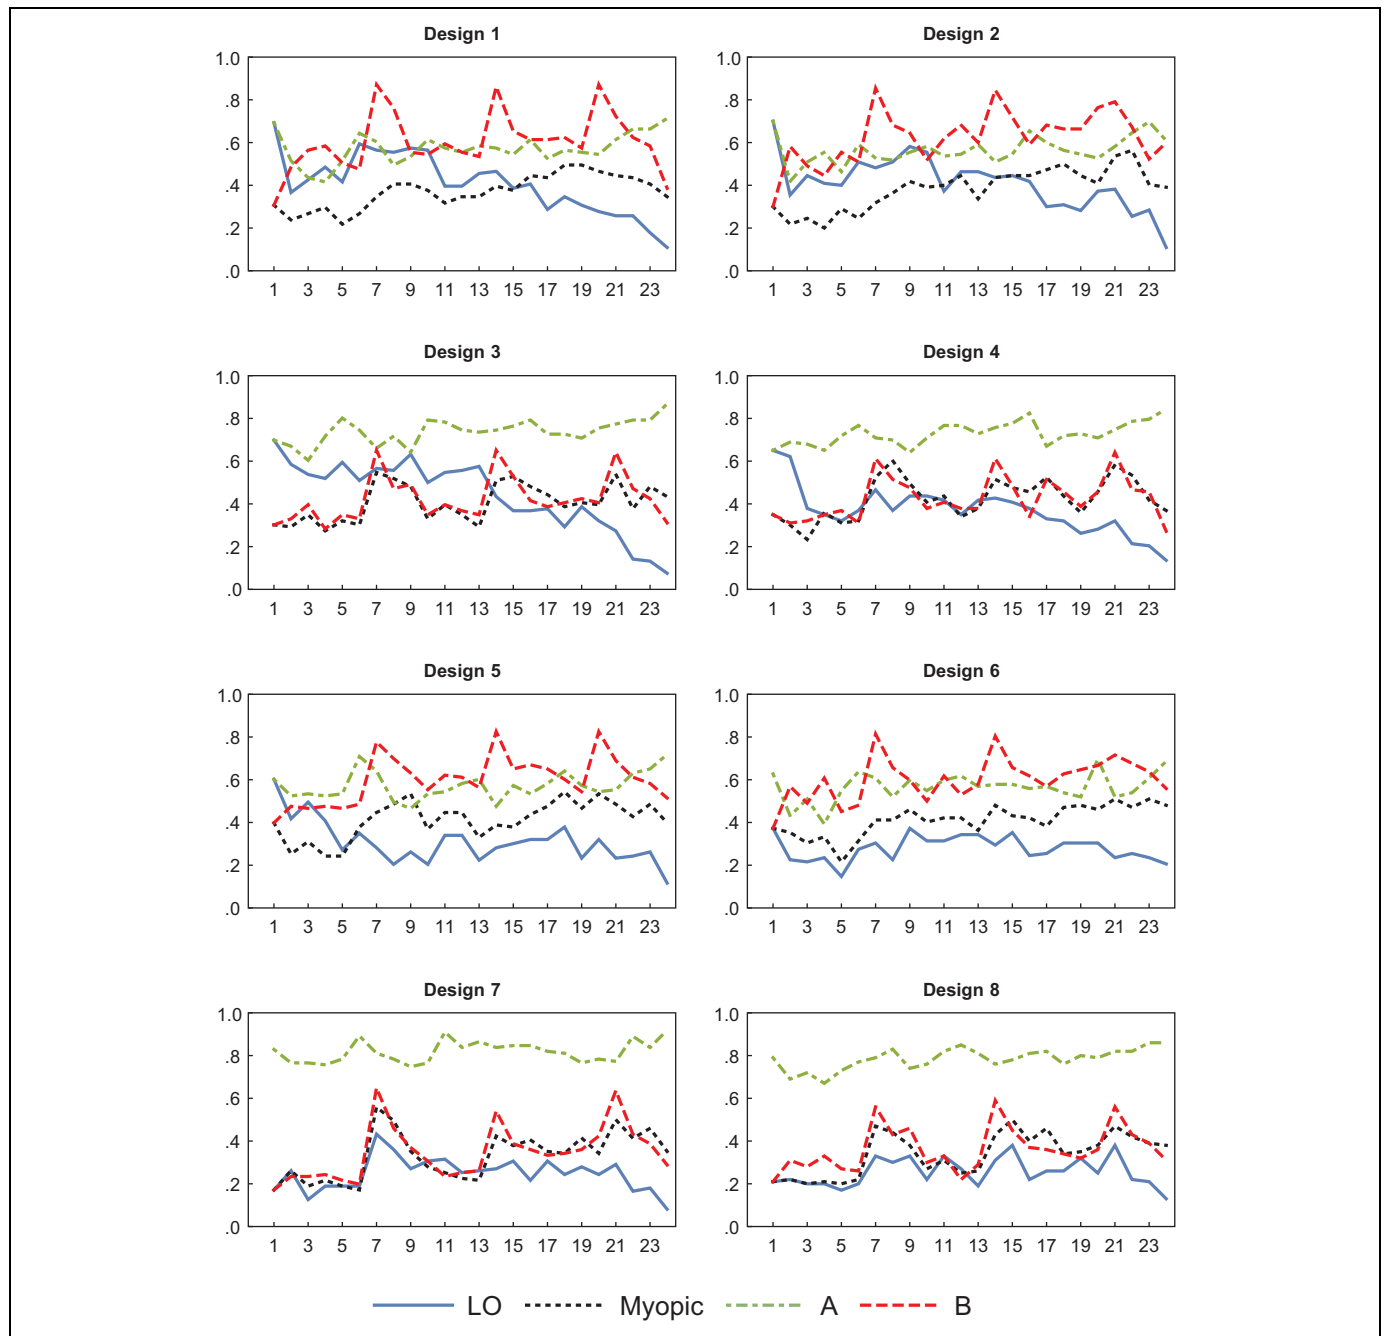


Figure 2. Deviation from benchmark strategies over trips and across designs.
 Notes: The curves show the proportion of participants in a design whose decisions deviate from each of the four benchmark strategies at each trip.

Determinants of Bounded Rationality

To further understand the boundedly rational behavior of participants, we use an individual-level model to investigate the conditions under which they are more likely to deviate from optimality. In our context, the conditions are characterized by the observed prices, the status of a consumer in the two competing loyalty programs, and the nature of the optimal choices, all of which also relate to the design that each participant is assigned to. We let d_{ijt} be a binary variable that is equal to 1

when participant i , who is assigned to D_j , deviates from the locally optimal decision on trip t and is equal to 0 otherwise. Then the probability of a deviation can be modeled as follows:

$$\Pr(d_{ijt} = 1) = \Phi(\gamma_t + \alpha_i + \eta_{1j} \Delta\text{Price}_{it} + \eta_{2j} \Delta\text{PointDist}_{it} + \eta_{3j} \Delta\text{Price}_{it}^2 + \eta_{4j} \Delta\text{PointDist}_{it}^2 + \beta'_j w_{it}), \tag{3}$$

where $\Phi()$ is the cumulative distribution function of the standard normal distribution. The γ_t term captures the time-specific

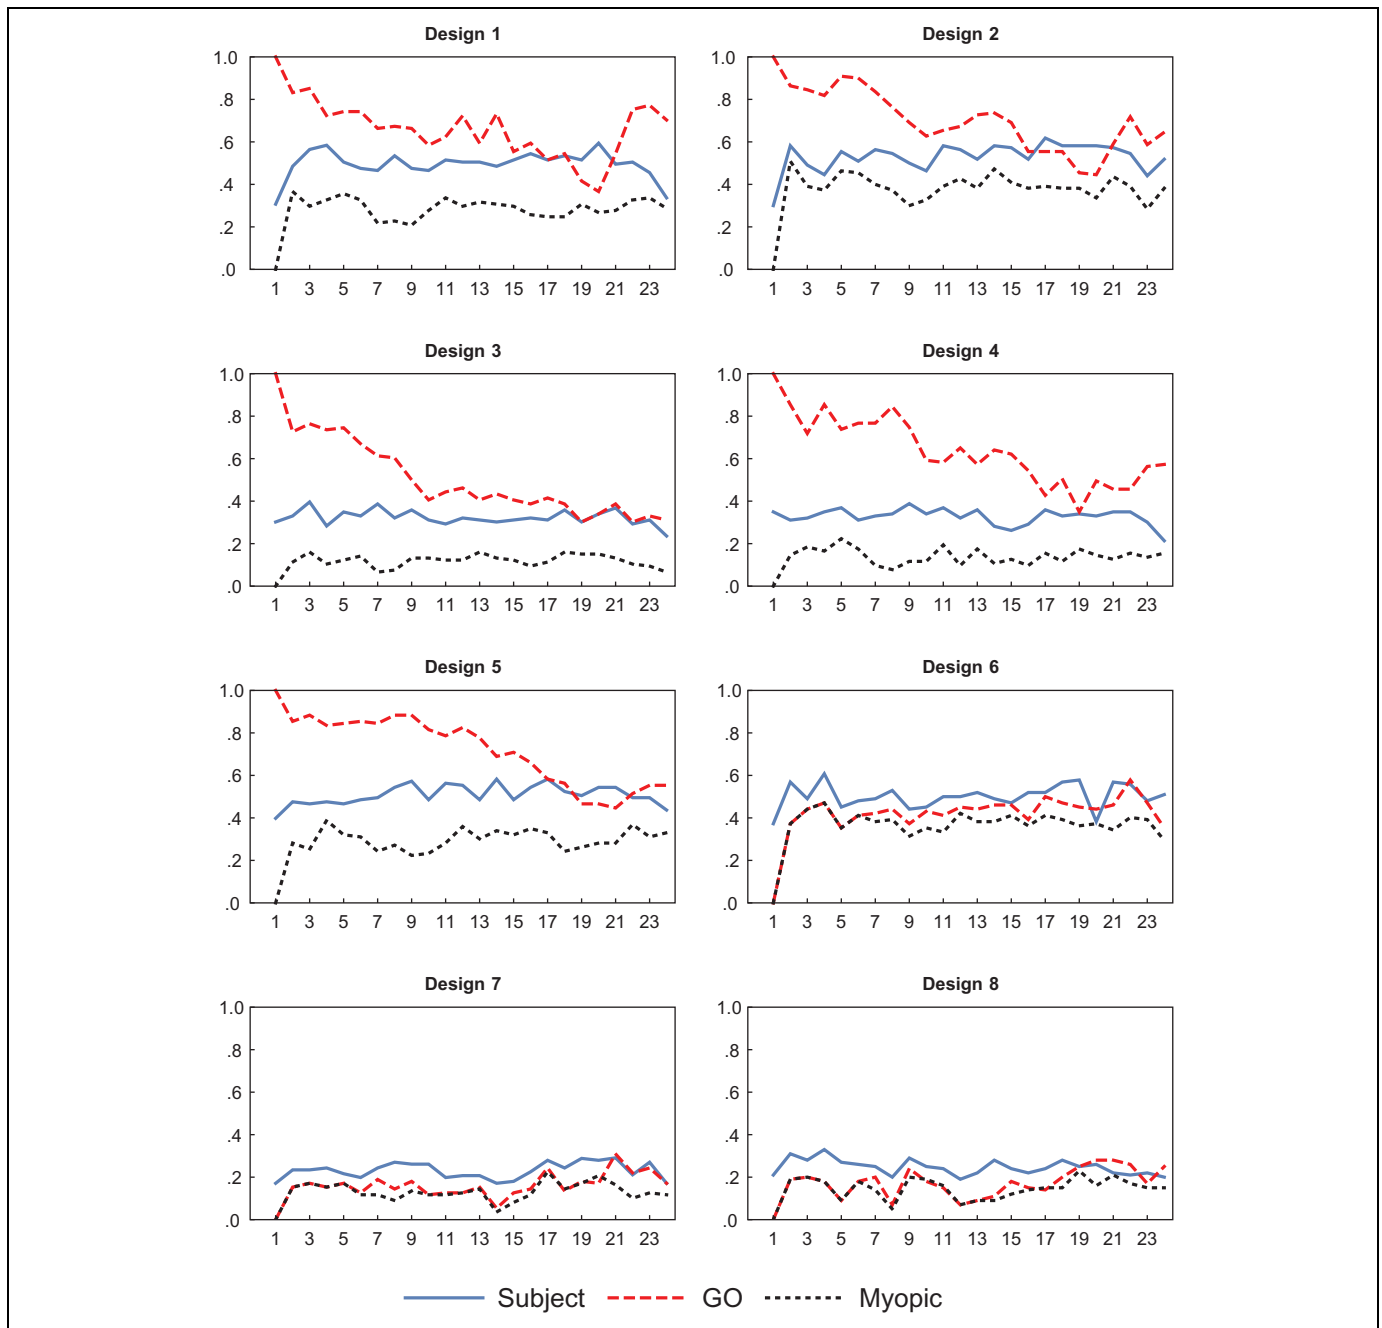


Figure 3. Airline A’s choice share over trips and across designs.

Notes: The curves show, for each design, the choice share of airline A at each trip under the three strategies (i.e., the actual decisions of participants, globally optimal strategy, and myopic strategy).

fixed effect (we set the baseline to be the last trip; i.e., $t = 24$); the intercept, $\alpha_i \sim N(0, \sigma_\alpha^2)$, is a random effect of individual i . Apart from these, we have two variables that describe the competitive environment faced by the participant at time t and three variables that characterize the optimal action that the participant should take at time t . We allow the effect of these variables to vary across designs using random coefficients, $\{\eta_j, \beta_j\}$, that come from independent normal population distributions. Note that the heterogeneity here is over the designs, whereas, α_i varies over participants.

For ease of illustration, let $O_{it} \in \{A, B\}$ denote the identity of the optimal airline for participant i at time t , and N_{it} denote the nonoptimal airline. Similarly, let $Price_{O_{it}}$ ($Price_{N_{it}}$) be the observed price for the optimal (nonoptimal) airline, and $Point_{O_{it}}$ ($Point_{N_{it}}$) be i 's accumulated points for the optimal (nonoptimal) airline. The variable $\Delta Price_{it}$ in Equation 3 denotes the difference between the prices of the optimal and the nonoptimal airline (i.e., $Price_{O_{it}} - Price_{N_{it}}$). Similarly, $\Delta PointDist_{it} = (rt_{O_j} - Point_{O_{it}}) - (rt_{N_j} - Point_{N_{it}})$ denotes the difference between the remaining distance to the redemption

Table 4. Determinants of Deviations from LO.

	Mean	SD	Het. Var.
ΔPrice_{it} : price difference	.341	.055	.033
$\Delta\text{PointDist}_{it}$: point distance difference	.687	.040	.016
$\Delta\text{Price}_{it}^2$: squared price difference	.129	.036	.009
$\Delta\text{PointDist}_{it}^2$: squared point distance difference	.046	.025	.004
w_{it1} : purchase, when only the optimal airline is free	1.481	.110	.077
w_{it2} : purchase, when only the nonoptimal airline is free	.256	.074	.023
w_{it3} : redeem, when only one airline has free ticket	-.834	.157	.175
w_{it4} : when both airlines have free tickets	.542	.145	.091
α_i : individual random effect	—	—	.174
Trip fixed effects (not shown)			

Notes: The baseline scenario for the dummy variables refers to the condition when it is optimal to purchase a ticket when no airline is free. Mean = the posterior mean; SD = the posterior standard deviation; Het. Var. = the variance of the heterogeneity distributions.

threshold for the optimal and the nonoptimal airline and can be used to quantify the effect of points pressure (Kivetz, Urminsky, and Zheng 2006). Therefore, larger values of $\Delta\text{PointDist}_{it}$ mean that participants are closer to getting a free ticket from the non-optimal airline than from the optimal one. These two variables are standardized before their inclusion in the model. We also include their squared terms to capture potential nonlinear effects. The vector w_{it} contains three indicator variables that describe all possible optimal scenarios. The variable w_{it1} denotes whether the optimal decision involves purchasing a ticket from O_{it} , when a free ticket is only available from O_{it} ; w_{it2} denotes whether the optimal choice requires purchasing a ticket from O_{it} , when a free ticket is only available from N_{it} ; w_{it3} denotes whether the optimal decision is to redeem a ticket when it is available only from the optimal airline. The variable w_{it4} denotes whether the optimal decision involves redeeming or purchasing a ticket when free tickets are available from both airlines. Here, we combine redeeming and purchasing for ease of empirical identification, as we have approximately 1% of observations in which there were free tickets from both airlines at the same time. The baseline scenario is that in which the optimal action is to purchase a ticket from O_{it} when no free ticket is available from either airline.

Table 4 reports the estimation results. The coefficient for ΔPrice_{it} implies, not surprisingly, that when the observed price from the optimal airline is higher than the price for the other alternative, participants are more likely to pick the nonoptimal airline. Its squared term also has significantly positive coefficients ($p < .01$), suggesting that this effect increases as the difference becomes larger, at least within the range of our data. Similarly, the coefficient for $\Delta\text{PointDist}_{it}$ implies that when participants are closer to getting a free ticket from the nonoptimal airline than from the optimal one, they are more likely to deviate from optimality and choose the alternative airline. Its squared term has marginally significant positive effect. A

comparison of the effect sizes of these two variables suggests that differences in point distances have a greater impact on the tendency to deviate. This is consistent with the fact that participants fail to realize the long-term benefit of loyalty programs, and thus, they tend to deviate more when they should take advantage from it.

Focusing on the four variables that characterize the nature of the locally optimal decision, we see that only w_{it3} is negatively associated with the chance of deviating from optimality, while the other three increase the probability of deviation. This suggests that the tendency to deviate from the optimal choice is affected by the presence of a free ticket and also by the airline on which it is available. The coefficient for w_{it1} indicates that the tendency to deviate is especially strong when the optimal action is to purchase a ticket when only the optimal airline has a free ticket. Similarly, the coefficient of w_{it2} suggests that people deviate from the optimal if it is optimal to purchase a ticket when the nonoptimal airline is free. Together, this implies that participants are tempted by the presence of a free ticket and tend to redeem their free tickets at the wrong occasions (i.e., they are not good at successfully navigating the two loyalty programs). The negative sign for the coefficient of w_{it3} suggests that the deviation from optimality is lower if the optimal action is to redeem when the optimal airline has a free ticket. The cross-design variances in the last column suggests that the aforementioned effect sizes vary significantly across the eight competitive environments. Finally, the large magnitude of the variance of the random effect α_i indicates that the participants vary considerably in their tendency to deviate from the optimal.

In conclusion, these results are generally consistent with our aggregate-level descriptive findings reported previously. More importantly, we find that the tendency to deviate from optimality depends on the overall competitive environment that participants experience. Furthermore, within a given environment, the level and the type of bounded rationality is influenced by the particular state of each decision scenario and the nature of the optimal action for that state.

Dynamic Choice Model with Bounded Rationality

Our descriptive analysis suggests that participants are clearly boundedly rational, making it necessary to account for this finding when modeling the decision-making process. This will enable us to further understand the rules that participants use when making dynamic choices and their responses to different marketing strategies. Next, we investigate the dynamic decision rules of individuals using the model proposed in Houser, Keane, and McCabe (2004).

Approximate Value Function

We extend the DP framework by allowing for the possibility that participants are boundedly rational. Specifically, we assume that participants rely on heuristics and can make mistakes when implementing these heuristics. We also assume that

there are K types/segments of heuristics that participants could use in their decision making. Let the latent variable $z_i \in \{1, 2, \dots, K\}$ denote individual i 's type. Then the original value function under the rationality assumption can be amended to include z_i ,

$$V_{ij}(\mathcal{I}_{it}, z_i) = c_{ij} + EV(\mathcal{I}_{i,t+1}|\mathcal{I}_{it}, z_i, j). \tag{4}$$

Rather than assuming that EV is the mathematical expectation operator, we approximate the future component of each choice alternative's value function with a flexible parametric function $F(\cdot)$. The future component for segment k can then be written as

$$EV(\mathcal{I}_{i,t+1}|\mathcal{I}_{it}, z_i, j) = F(\mathcal{I}_{it}, j|z_i) + \epsilon_{ij}, \tag{5}$$

where ϵ_{ij} is an optimization error that accounts for the chance that a participant may deviate from the exact implementation of particular decision rule. We assume that the distribution of ϵ_{ij} varies by segment to allow for differences in the implementation of decision rules.

Note that similar to Houser, Keane, and McCabe (2004), we assume that discount factor equals one, as the time frame involved in our experiment is very short and the entire payment is realized only at the end of the game, not during the game. Moreover, participants were instructed to minimize the total undiscounted cost for all the trips. Given this objective function, the timing of the per-period costs does not matter in our game. However, discounting could be relevant in contexts involving reward programs, especially when costs are spread out over time. Recent research by Ching and Ishihara (2017) and Ching and Osborne (2019) shows that the discount factor can be identified in a reward program setting similar to ours but involving an infinite horizon. Rossi (2018) has applied the identification argument in Ching and Ishihara (2017) to estimate the discount factor in the context of a single loyalty program. Given these developments, future researchers could estimate the discount factor in empirical contexts involving competing loyalty programs if the costs are spread out over time and discounting is salient. In situations where discounting is relevant, if researchers compute the benchmark optimal DP solution assuming no discounting, but customers in reality use a discount factor less than one and otherwise are fully forward looking, then conclusions made regarding how, and how much, behavior deviates from rationality could be affected by this assumption about the discount rate.

The most popular strategy for specifying the future component involves using a flexible polynomial specification that represents important state variables (Geweke and Keane 2000, 2001; Houser, Keane, and McCabe 2004; Imai and Keane 2004; Keane and Wolpin 1994; Krusell and Smith 1996; Powell 2009). Research has shown that value functions can be approximated accurately using lower-order polynomials, thus mitigating the curse of dimensionality that arises from a large state-space. However, it is important to note that if this approximation is not accurate, then conclusions regarding the nature and extent of bounded rationality can be

affected. For a choice j , we specify the future component using the terms in

$$\begin{aligned} \mathcal{M}(\mathcal{I}_{it}) = & (a_{itA} + x_{itA}\mathbb{I}(j = 1) - r_{tA}\mathbb{I}(j = 3), a_{itB} + x_{itB}\mathbb{I}(j = 2) \\ & - r_{tB}\mathbb{I}(j = 4), x_{itA}, x_{itB}, n_{it} - 1), \end{aligned}$$

which govern the Markovian law of motion for the state variables. The first two items denote the points inventory for the two airlines, and the last one is the number of remaining trips. We approximate the future utility component with a polynomial of order two over $\mathcal{M}(\mathcal{I}_{it})$. Because choices depend only on the relative values of all the alternatives, not all the parameter in the polynomial can be identified. After removing terms that are common to all alternatives, the polynomial function $F(\mathcal{M}(\mathcal{I}_{it})|j, \beta_{z_i})$ for segment k has 11 coefficients to be estimated. Web Appendix C shows the derivation and the simplified functional form. Then, the value function for individual i of type z_i for alternative j at time t can be written as

$$V_{ij}(\mathcal{I}_{it}, z_i) = c_{ij} + F[\mathcal{M}(\mathcal{I}_{it})|j, \beta_{z_i}] + \epsilon_{ij}. \tag{6}$$

We can finally write the decision rule for participant i of type k at time t as

$$\arg \min_{j \in \mathcal{A}(\mathcal{I}_{it})} \{c_{ij} + F[\mathcal{M}(\mathcal{I}_{it})|j, \beta_k] + \epsilon_{ij}\}. \tag{7}$$

We also assume that the error term ϵ_{ij} follows a double-exponential distribution with a scale parameter μ_k for participants with type k . Then the probability that a participant who belongs to the k th segment will choose alternative $j \in \mathcal{A}(\mathcal{I}_{it})$ can be computed as

$$\Pr(d_{it} = j | \mathcal{A}(\mathcal{I}_{it}), \beta_k) = \frac{\exp(\mu_k \{-c_{ij} - F[\mathcal{M}(\mathcal{I}_{it})|j, \beta_k]\})}{\sum_{j' \in \mathcal{A}(\mathcal{I}_{it})} \exp(\mu_k \{-c_{ij'} - F[\mathcal{M}(\mathcal{I}_{it})|j', \beta_k]\})}, \tag{8}$$

where $\mathcal{A}(\mathcal{I}_{it})$ denotes the current choice set. In the preceding equation, a smaller logit scale parameter means a larger variance of the error term, and as all the variables in the future utility component are positive, a larger β_k leads to a smaller utility for all the alternatives.

Identification, Inference, and Estimation

Our game design and model specification enable us to identify all the parameters in the approximated future utility component using the choice data of the participants. According to Geweke and Keane (2000), it is important to have at least one of two conditions for identification. The first is that the current payoff is observed (Geweke and Keane 2000), which is true in our case. The second is that the model satisfies exclusion restrictions (Ching, Erdem, and Keane 2014), such that some state variables for some specific values have no impact on current payoffs but can influence the expected future payoffs. In our case, the accumulated reward points have no impact on the current payoff unless they exceed the reward threshold. This

Table 5. Estimation Results for ADP Model.

Variable	Segment 1 "Kind of Rational Toward Price"	Segment 2 "Price Sensitive"	Segment 3 "Reward Sensitive"	Segment 4 "Conflicted"	Segment 5 "Kind of Rational Toward Reward"
A's point inventory (PI)	1,104.73 (485.33)	1,719.73 (692.10)	-136.27 (106.35)	-.20 (.04)	-.20 (.02)
B's PI	1,402.07 (481.73)	-64.77 (83.82)	733.34 (305.37)	.03 (.04)	-.07 (.01)
Square of A's PI	200.70 (160.04)	-788.36 (241.26)	-42.80 (14.13)	-.02 (.00)	-.03 (.00)
Square of B's PI	385.16 (192.19)	-26.76 (9.34)	-219.66 (82.84)	-.04 (.00)	-.00 (.00)
(A's PI) × (B's PI)	1,095.67 (466.86)	6.54 (33.38)	1,501.94 (481.93)	-.02 (.00)	.00 (.00)
(A's PI) × (A's price)	735.21 (861.31)	1,058.55 (964.67)	1,311.06 (429.27)	.11 (.10)	.09 (.05)
(A's PI) × (B's price)	578.38 (913.00)	-245.72 (948.40)	-11.40 (374.50)	.30 (.13)	.25 (.06)
(A's PI) × (trips left)	-19.15 (17.13)	-18.17 (12.63)	-6.58 (2.43)	-.01 (.00)	-.00 (.00)
(B's PI) × (A's price)	-657.35 (853.54)	-633.50 (249.95)	-1,689.75 (713.56)	-.05 (.09)	-.11 (.02)
(B's PI) × (B's price)	7.58 (864.86)	1,370.56 (432.67)	725.14 (806.92)	-.02 (.16)	-.09 (.03)
(B's PI) × (trips left)	-13.89 (17.41)	-7.00 (2.43)	14.03 (5.71)	-.01 (.00)	-.00 (.00)
Scale parameter	.04 (.01)	.01 (4.77)	.01 (.00)	23.36 (.00)	118.58 (1.87)
Segment size	.11 (.01)	.14 (.02)	.17 (.01)	.24 (.01)	.35 (.02)
Fitting accuracy	84%	82%	87%	61%	84%

Notes: The parentheses contain the posterior standard deviations. Values in boldface are statistically significant, based on the 95% posterior intervals.

is explained by Ching and Ishihara (2017) within the context of frequent buyer reward programs. Therefore, our model satisfies both identification conditions. Note that in the previous work by Geweke and Keane (2000) and Houser, Keane, and McCabe (2004), only the first condition is satisfied, and in other DP models on frequent buyer reward programs (Ching and Ishihara 2017; Hartmann and Viard 2008; Kopalle et al. 2012; Lewis 2004), only the second condition is satisfied.

Next, we describe our estimation method. Let π_k denote the probability that a participant is of type k . Then the likelihood function of the observed data can be written as

$$L(\Upsilon) = \prod_{i=1}^N \left(\sum_{k=1}^K \left\{ \pi_k \prod_{t=1}^T \Pr[d_{it} | \mathcal{A}(\mathcal{I}_{it}), \beta_k] \right\} \right), \quad (9)$$

where $\pi_k \in (0, 1)$ represents the segment size, such that $\sum_k \pi_k = 1$. We rely on Bayesian inference and use Markov chain Monte Carlo methods to estimate these parameters $\Upsilon = \{\beta_k, \pi_k, \mu_k\}_{k=1}^K$. We assume the following priors for these model parameters:

$$\beta_k \sim N(0, 100^2 I), \mu_k \sim \text{Cauchy}^+(0, 5),$$

$$(\pi_1, \pi_2, \dots, \pi_K) \sim \text{Dirichlet}_K(1),$$

where, Cauchy^+ is a half-Cauchy prior and $\text{Dirichlet}_K(1)$ refers to a K -variate symmetric Dirichlet distribution, with its parameter set to unity. In particular, we use Hamiltonian Monte Carlo sampling (Neal 2011) for inference, which uses the gradient of the log-joint distribution in simulating parameter draws from the posterior distribution. It therefore avoids the random walk behavior of ordinary Metropolis-Hasting (Metropolis et al. 1953) methods and explores the parameter space more efficiently.

We estimated the dynamic choice model with K ranging from two to six segments and used the deviance information

criteria (DIC) proposed by Spiegelhalter et al. (2002) to select the number of segments. The DIC values are 20,133 for two segments, 18,429 for three segments, 18,013 for four segments, 17,911 for five segments, and 17,981 for six segments. Thus, the five-segment solution has the highest support.⁵ Table 5 reports the posterior means and standard deviations of the logit-scale parameters, the segment sizes, and the parameters in the future utility components for the five segments. We see that the segments vary in size, with smallest being at 11%. In addition, the number of variables that have a significant effect varies across segments, ranging from 4 to 9 out of a total of 11 variables. The coefficients in these decision rules are difficult to interpret as the underlying variables appear in a nonlinear and interactive fashion as part of the polynomial specification. We therefore analyze the differences across segments by first assigning participants to segments and then comparing behaviors across segments.

Understanding Consumers' Decision Rules

Segment membership and fitting. To better characterize the behavior of the segments and to connect these with the eight market designs, we used the vector of posterior segment probabilities to simulate each participant's segment membership. The highest posterior segment probability is at least .90 for 85% of our participants. This indicates that the choices are very informative about a participant's type and one can clearly classify a majority of the participants. We used multiple draws of each participant's segment membership to capture the uncertainty in

⁵ In our empirical analysis, we also tried an alternative polynomial expansion that included a participant's current choice as an additional variable in the transition function $\mathcal{M}(\cdot)$. This yielded a specification with 31 parameters per segment. However, this resulted in much higher values of the DIC, compared with the model we report herein.

the segment allocations. We then used these simulated segment memberships to evaluate how well the model parameters for the segment that a participant is assigned to predict the choices made by the participant. Specifically, given the realizations of the price sequence that each participant actually observed, we simulated the participant's choice decisions for all the trips either separately or in sequence. At each trip, we simulated the choice using the parameters of the participant's assigned segment and the accumulated point inventories on the two airlines using the past observed (or simulated) choices. The resulting average prediction accuracy for segments 1 to 5 is 93%, 87%, 89%, 63%, and 91%, respectively, when all the choices are predicted separately and is 77%, 84%, 71%, 52%, and 85% when the choices are predicted in sequence.⁶ Not surprisingly, we see that the accuracy decreases when the point inventories are computed using the past simulated choices instead of the actual choices. At the end of this section, we explain why the prediction accuracy is relatively lower for segment 4.

Aggregated preferences and performance. Table 6 compares the average behavior of the participants that are assigned to each segment along a few dimensions. First, on average, at least 93 participants are assigned to each segment, and the posterior segment size for each segment k is almost identical to the prior size π_k . Second, there is a large variation across segments in the deviation level from a given benchmark strategy. For example, Segments 1 and 5 deviate from LO much less frequently than the other three segments, whereas Segments 2 and 3 deviate the most. Actually, all segments except Segment 4 exhibit relatively clear deviation tendencies to certain benchmark strategies. For instance, Segment 2 deviates the most from naive A, and Segment 3 deviates the most from naive B. Third, we see that the actual probabilities of choosing A vary a great deal across the segments, even though the corresponding choice shares of the globally optimal strategy are much more similar (due to random assignment). In particular, Segment 2 tends to underchoose A by $41 - 4 = 37\%$ relative to the optimal strategy, whereas Segment 3 tends to overshoot by $84 - 55 = 29\%$. Interestingly, though the choice share of A in Segment 4 deviates the least from GO compared with other segments, the actual sequential choices deviate a lot more from LO. This suggests that Segment 4 participants are conflicted between the two alternatives and do not have explicit preferences in their decision making. Finally, it seems that Segments 1 and 5 are similar to each other across many dimensions, but Segment 5 prefers airline A (i.e., reward) slightly more than Segment 1, based on the choice shares.

⁶ To further validate our model, we also conducted posterior predictive checks (Gelman, Meng, and Stern 1996). Specifically, we dynamically simulated multiple choice sequence of each participant from the posterior predictive distribution using the post convergence Hamiltonian Monte Carlo draws for all parameters to generate multiple replicated data sets. We then compared these replicated data sets with the actual data set using four test statistics to assess how well our model recovers important aspects of the data-generating process. Details are provided in Web Appendix E.

Dynamics in performance and preference. Figure 4 shows the average deviations from each of the four benchmark strategies over the 24 trips for all segments separately. The overall level of the deviation across the segments is consistent with the aggregate results in Table 6, so here we only highlight new findings. We see a significantly decreasing trend (over the trips) in the percentage of participants that deviated from the optimal choices ($p < .05$) for all segments except Segment 3. The decreasing trend reflects the fact that it is easier to be forward looking when the horizon is short. For the same reason, deviation from myopic strategy exhibits a significantly positive trend ($p < .05$) for all segments except Segment 3.

Figure 5 shows airline A's average choice share across trips among participants assigned to each segment, compared with the share under the optimal and myopic strategies, using the same observed price sequences. First, we see that the overall pattern is also consistent with the aggregate results reported in Table 6, and the choice share pattern is mostly stable across trips. Second, Segment 3 consistently chooses A with a percentage that is even higher than that implied by GO. This further confirms that Segment 3 is very much loyal to airline A. Yet we see that this choice share drops slightly towards the end of the game. These observations suggest that participants in Segment 3 are more loyal to A than they should be, but they are able to realize the benefit of switching to A's competitor when the need for such forward looking is less important. In contrast, Segment 2 is even less likely to choose A than implied under a myopic strategy. This suggests that consumers in Segment 2 show loyalty to B, even though they make decisions heavily based on prices. Finally, the choice share patterns of Segment 1 and 5 are very similar—both consistently mimicking the myopic strategy—although they both are relatively more “rational” than the other segments.

Design and segments. We now examine the relationship between the decision rules used by our participants and the competitive environment to which a participant belongs. Given each participant's segment membership and assigned game design, we can obtain a 8×5 contingency table that records the number of participants for all design-segment combinations. For ease of illustration, we present in Figure 6 the proportion of participants that belong to each of the five segments within a given design. All segments appear in every design, suggesting that the five decision segments are generalizable across market designs. However, the chi-square independence test rejects the null hypothesis that the market design is independent of participants' decision rules ($p < .001$). Therefore, the market design that is assigned to each participant may affect its probability of adopting a certain decision rule.

In particular, based on the average proportions, we summarize which designs are most/least likely to trigger which behavioral segments in Table 7. Note that the choice share of A under the optimal strategy tends to be lower for higher design numbers (see Table 3). Thus, we find that under D_7 and D_8 (D_1 and D_2), when optimal behavior should choose B (A) for more than 70% of the trips, participants are more likely to come from Segments 1 and 2 (3 and 5), which exhibit relatively stronger

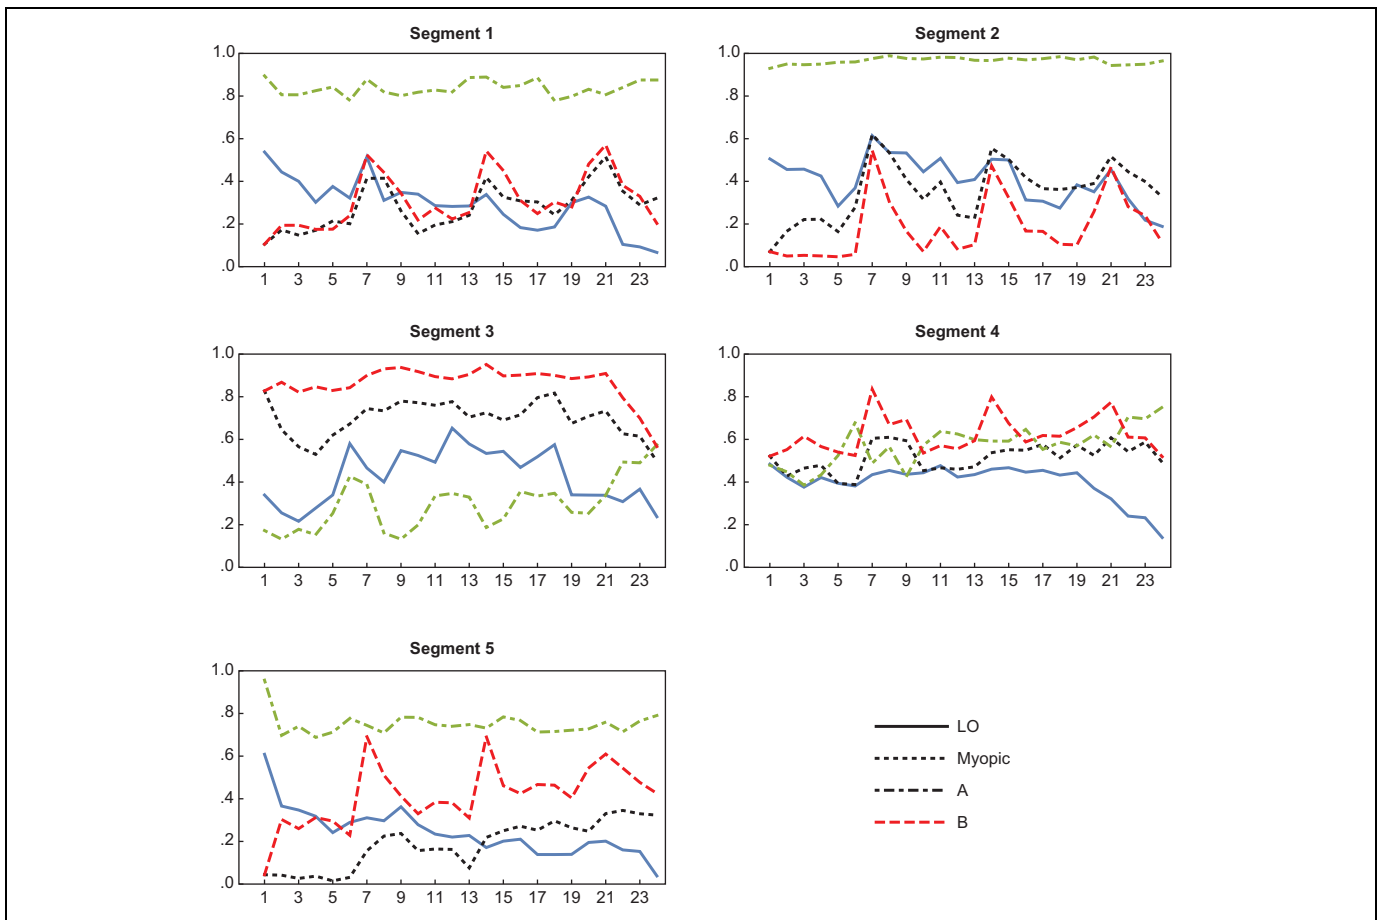


Figure 4. Deviation from benchmark strategies over trips and across segments.
 Notes: The curves show the average percentage of participants whose decisions deviate from each of the four benchmark strategies at each trip, for the five segments.

Table 6. Descriptive Statistics (in %) Across Segments.

	Segment 1 "Kind of Rational Toward Price"	Segment 2 "Price Sensitive"	Segment 3 "Reward Sensitive"	Segment 4 "Conflicted"	Segment 5 "Kind of Rational Toward Reward"
Percentage of participants	10	14	17	24	35
Deviation from LO	30	41	43	40	24
Deviation from myopic	28	36	70	52	19
Deviation from naive A	83	97	29	57	75
Deviation from naive B	31	19	86	62	42
Pr(A) for GO	46	41	55	52	50
Pr(A) for participants	20	4	84	55	28
Pr(A) for myopic	18	19	27	25	23

Notes: The first row is the average percentage of participants assigned to each segment. Starting from the second row, each column reports the average values across participants assigned to the same segment. The second row is the average percent of participants' choices that deviate from LO. Pr(A) represents the choice share of A between A and B. All numbers represent percentages.

sensitivity to prices (rewards). In comparison, Segment 4 participants who exhibit conflicted preferences are more likely under D_6 , where the optimal strategy shows no clear preference toward one particular airline. In conclusion, we find that the market designs may lead to different optimal strategies and thereby affect the likelihood of adopting particular decision rules.

Given these findings, we present these five segments' relative preferences to the two airlines in Figure 7. We label them as "kind of rational toward price," "price sensitive," "reward sensitive," "conflicted," and "kind of rational toward reward." These names reflect each segment's overall level of rationality and its sensitivity to rewards or prices. Segment 1 and 5 are

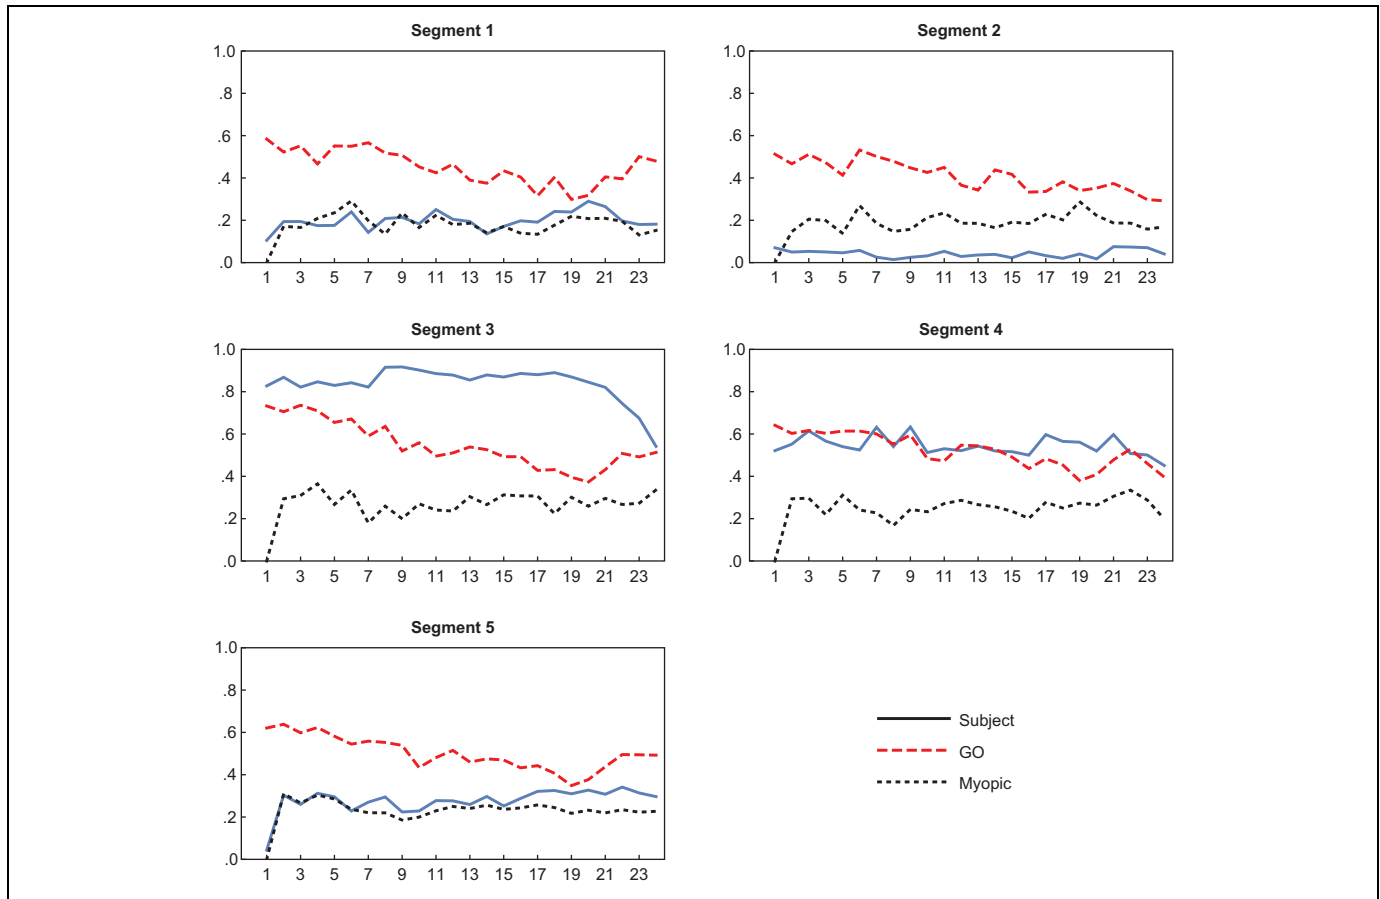


Figure 5. Choice share over trips and across segments.

Notes: Each panel displays A’s average choice share at each trip among participants who are assigned to the same segment. As benchmarks, we also derive A’s choice share under the globally optimal strategy and the myopic strategy using the same set of prices that were observed by actual participants at each trip in the study.




Figure 6. Distribution of segments within each design.

both kind of rational, mimicking the pattern of the optimal behavior, but with a consistently lower share for A. Their main difference is whether they care for price more than reward. Segment 2 is “price sensitive” in its extreme loyalty to B. Segment 3, which we label “reward sensitive,” mostly favors airline A, but its “reward sensitivity” in favor of A is not as extreme as the “price sensitivity” of Segment 2 to B, as participants in Segment 3 do realize more situations that they should

Table 7. How Market Design Affects Participant Decision Strategy.

Segment	Most Likely	Least Likely
1. “Kind of rational toward price”	D ₇	D ₁ , D ₆
2. “Price sensitive”	D ₈	D ₁ , D ₂
3. “Reward sensitive”	D ₁ , D ₂	D ₇ , D ₈
4. “Conflicted”	D ₆	D ₇ , D ₈
5. “Kind of rational toward reward”	D ₁	D ₄

instead purchase from B. Finally, Segment 4 is most “conflicted” in its preferences and thus is relatively harder to predict (the prediction accuracy is lowest among all segments). We also confirm our labeling of these segments using participants’ self-reported decision rules. For details, see Web Appendix D.

Consumer Optimality and Market Outcomes

We conducted a simulation study to understand whether and how different consumer decision strategies affect market outcomes and to determine which competitive environment is best for a firm. Specifically, we compared the airline profits across our eight designs, assuming that consumers are optimal,




Figure 7. The relative preferences of segments.

myopic, or boundedly rational. Under bounded rationality, we used two scenarios. The first, “boundedly rational aggregate,” assumes that a hypothetical consumer follows one of the previously identified five segments with a probability equal to the overall segment size (as in Table 5). The second boundedly rational scenario, “boundedly rational adaptive,” assumes that the consumer follows one of the five segments based on the segment sizes within a particular market design. This is consistent with our finding that participants’ choice of heuristics depends on the market conditions they face (see Figure 5 and Table 7). Given eight designs and the four assumptions regarding consumer optimality, our simulation considered $8 \times 4 = 32$ scenarios.

We computed the profits under each design based on 10,000 simulated sequences of realized prices for the 24 trips. These prices were simulated using the price distribution f for the design. We then used these design-specific price sequences to simulate 10,000 choice histories for each of the four consumer rationality assumptions: optimal, myopic, boundedly rational, and adaptive. Therefore, the same price sequences were used to compute the market outcomes under each design across the four optimality assumptions. We computed the revenue of each firm as the dollar value of the tickets purchased by consumers across all the trips and obtained the profit by subtracting the cost of offering the loyalty program from the revenue. For simplicity, we report the results when the cost is defined as the total dollar amount of the free tickets that the firm gives to its customers. However, we found that the overall qualitative conclusions remain the same with alternative assumptions that use only a proportion of this cost (e.g., 80%).

Figure 8 shows the normalized profits of airlines A and B under the 32 scenarios. As we only care about the relative differences across conditions, we normalized the profits by dividing each profit with the total profit under D_3 for the bounded rationality aggregate scenario. The figure uses a stacked bar chart where the length of each bar indicates the total market profit obtained by summing the profits for the two firms. The dark proportion of each bar indicates the profit of airline A, and the lighter part represents the profit for B. We first compare the results across all the 32 scenarios. Regardless of the market design, we can see from Figure 8 that the total market profit is the highest for the boundedly rational scenarios, followed by the adaptive, myopic, and optimal scenarios. These results are consistent with the aggregate analysis in Table 3 which shows that the myopic strategy can help participants save more money than they do with their actual decision rules. Again, this implies that consumers, in making their trade-offs between the rewards from the loyalty program and the promotional opportunities, end up worse off when

compared with choosing based solely on observed prices (i.e., promotions). Our findings also imply that adaptive strategies could lead to better outcomes for customers.

In addition, we compare the best and worst designs for firm A, firm B, and the whole market, under each of the four assumptions on consumer decision strategy. Table 8 summarizes the findings. It is clear from Table 8 that the overall market outcomes and the performance of each airline depend on both the competitive environment and the extent of consumer optimality. Thus, assumptions about consumer decision making strategies are critical in determining what is the best course of action and which type of market environment is most appealing for the entire market, and for each firm in the market.

Conclusions

We studied, experimentally, how consumer choices are dynamically shaped by the interplay of competing loyalty programs and the pricing and promotional strategies of firms. Our results highlight decision rules and heuristics that consumers use to simultaneously navigate multiple loyalty programs, and our findings are highly relevant for the development of pricing and promotional strategies by firms. We showed that consumers are unable to be fully forward looking in their purchasing activities and are especially vulnerable at particular decision states where free tickets are available. Our result that consumers are boundedly rational is consistent with the behavioral literature on dynamic decision making, which suggests that consumer optimality is especially difficult when feedback is limited or delayed (Meyer and Hutchinson 2016). Our results also show that both market-level outcomes and the performance of specific firms depend on the competitive environment as well as the level of consumer optimality. These factors jointly influence a firm’s revenue as well as its cost of running the loyalty program, and therefore, the assumption that one makes about consumers’ decision strategies is critical in determining which competitive environment is most, or least, appealing for the entire market or for a particular firm.

We showed that consumer decisions consistently deviate from the optimal ones even though some individuals are somewhat able to realize the delayed but larger benefits from the loyalty programs. We found that these deviations stem from the fact that individuals are not forward looking enough and cannot resist the temptation of immediate benefits accruing from lower prices. This makes consumers worse off than under a myopic decision strategy that only relies on observed prices. At the aggregate level, consumers tend to underappreciate a firm that offers a higher regular price, but a lower redemption threshold in its loyalty program. We found that, compared with




Figure 8. Normalized simulated profit of A and B.

Notes: Each stacked bar shows the normalized profits for A and B for each of the 32 scenarios: 8 (market design: D1, D2,..., D8) × 4 (consumer decision rule: optimal, myopic, boundedly rational aggregate, boundedly rational adaptive). The normalization is done by dividing each profit with the total profit under D₃ for the bounded rationality aggregate scenario.

Table 8. (Best, Worst) Design for Different Players.

Assumption	A	B	Total Market
Optimal	(D ₅ , D ₇)	(D ₇ , D ₅)	(D ₇ , D ₄)
Myopic	(D ₂ , D ₃)	(D ₃ , D ₂)	(D ₇ , D ₂)
Boundedly rational aggregate	(D ₆ , D ₈)	(D ₃ , D ₂)	(D ₃ , D ₂)
Boundedly rational adaptive	(D ₆ , D ₈)	(D ₇ , D ₂)	(D ₄ , D ₁)

promotion probability and sale prices, changes in redemption threshold have more consistent impact on participants' decisions/preferences across contexts. In addition, trading off the immediate benefits from sales promotions with delayed ones from loyalty programs is especially difficult for participants.

Importantly, we found that the preferences and the optimality of consumers varied significantly across our market designs. Even within the same competitive environment, the tendency to deviate is influenced by the particular state of each decision scenario (e.g., the observed prices, the loyalty program states, the number of remaining trips) and the nature of the optimal action required for that state. Specifically, consumers are more likely to deviate when faced with a large decision horizon (i.e., when a large number of decisions remain). Ironically, this is exactly when forward-looking behavior is most relevant but also difficult. In addition, consumers are more likely to be confused in the presence of free tickets, as they tend to redeem these at the wrong occasions.

We also found considerable heterogeneity in the decision rules used by individuals and in their ability to behave optimally. Drawing on our model estimates, we classified participants into five segments (“conflicted,” “reward sensitive,” “price sensitive,” “kind of rational toward price,” and “kind of rational toward reward”) that exhibit variation in their decision strategies. The “conflicted” segment has no clear preference over either airline, and thus, it is relatively harder to predict their choices. The “reward-sensitive” segment consists of individuals who exaggerate the benefits from the loyalty program to the extent that they forgo the immediate benefit from lower prices. Therefore, a firm that is focused on lower prices will not be effective in attracting this segment by merely offering slightly lower prices than its competitors. In contrast, the “price-sensitive” segment is unable to appreciate the delayed benefits from the loyalty program and focuses mostly on immediate prices. Thus, it is difficult for a firm that offers a better loyalty program but with much higher prices to attract individuals in this segment. Finally, we uncovered two “kind-of-rational” segments (approximately 46% of participants) who exhibit the lowest deviation from optimality and a stable pattern of preferences over time but differ in whether price matters more than reward. Even though we find that all these five segments are generalizable across market designs, the likelihood that a given participant will adopt a certain strategy varies. In particular, such likelihood is in the same direction as what the optimal strategy should do under the same design,

implying that customers could adapt the decision strategies to the market environment.

Our findings yield managerial insights about how firms could design their reward programs, pricing, and promotion strategies to leverage the boundedly rational behavior of customers. First, we observe that, on average, both firms can acquire significant market share under the eight market designs. This is consistent with theoretical work that shows the existence of equilibrium in symmetric markets where both firms offer loyalty programs or sale prices (Caminal and Matutes, 1990; Kim, Shi, and Srinivasan 2001; Klemperer 1995). Our work also clarifies how a firm can balance its promotional strategy with its loyalty program under different market conditions and in the presence of consumer heterogeneity and firm competition. However, it is important to note that because we do not use a structural model, our conclusions about firm outcomes are subject to the Lucas critique.

More broadly, our research contributes to the literature of bounded rationality on several fronts. It showcases individual-level heterogeneity in the extent of bounded rationality and the types of rules that consumers are likely to use in the substantive context of competing loyalty programs. Our results show that the extent and type of bounded rationality varies by context (competitive environments and decision states, in our case). Most importantly, our modeling approach can identify and quantify the factors that drive bounded rationality in consumer decision making. Policy makers can utilize such knowledge to assist consumers in making better decisions and positively affect consumer welfare.

Even though, for simplicity, we limited our study to a market with two competing firms, our findings enable us to conjecture about behavior in markets with multiple firms. First, adding more firms will make the dynamic decision problem much harder for people, and therefore, we expect greater

deviations from optimality and a wider reliance on simple heuristics. Second, we expect that participants will continue to make “mistakes” regarding when to use a free ticket and that their behavior will continue to reflect psychological effects such as the goal gradient effect. Regarding the other substantive results, we expect that our research findings about market outcomes should hold within-market designs that are similar to the current ones (i.e., they include firms that do not dominate either on the prices they offer or through their reward programs). Nevertheless, future researchers could extend our study to multiple firms to fully investigate these anticipated impacts.

This article has several limitations and offers important directions for further investigation. First, we admit that the inferred decision rules depend on our assumed parametric form of the future value function, which may not be expressive enough to cover all possible decision rules. Future research may explore this issue further assuming noncompensatory alternatives (Gilbride and Allenby 2004). Second, future research might explicitly compare self-stated rules with the decision rules inferred from a proposed method. While we included an open-ended question about participants’ actual decision strategies in the game, most participants did not provide an answer. Third, we fixed the number of trips T to obtain the optimal strategy and to ensure comparable market designs and an incentive-aligned experiment. Future research could investigate how the number of trips could depend on program design. Fourth, our experimental paradigm provides opportunities for studying other aspects of competing loyalty programs including vertical differentiation between competing firms, flight availability, reward availability, and programs with very different design features, such as currency, fee policy, tiered programs, among others. Finally, if there is strong evidence that participants learn over time, this could be explicitly included in the modeling framework.

Appendix. Actual instructions for participants

Instructions: How to Play

Please read carefully the following information. This will help you better achieve your goal in the game.

In this game, the two available airlines, A and B, differ in their ticket **prices** and **reward programs**. Each airline sells round-trip tickets at either a **regular** price or a **discounted (lower)** price on any given trip. **A's** regular price is **\$480**, and its discounted price is **\$400**. **B's** regular price is **\$440**, and its discounted price is **\$360**. Across all the trips, you should expect to receive the promotional price level from each airline **30%** of the time, and receive the same price level (i.e., both promotional or both regular) from the two airlines **70%** of the time. That also means, you should expect that:

- **55%** of the time, both will sell at their regular prices, i.e., **\$480** from **A**, **\$440** from **B**;
- **15%** of the time, both airlines will sell at their discounted prices, i.e., **\$400** from **A**, **\$360** from **B**;
- **15%** of the time, **A** will sell at its discounted price **\$400**, while **B** will sell at its regular price **\$440**;
- **15%** of the time, **A** will sell at its regular price **\$480**, while **B** will sell at its discounted price **\$360**.

You are a member of both airlines' reward programs. For each airline, you earn **1 reward point for each dollar** that you spend in buying its tickets. You will win **free tickets** for future trips by **redeeming these reward points** that you accumulate. **A** requires **2000 points**, while **B** requires **2400 points** for a free ticket. Any leftover points at the end of a trip will rollover to the next trip, except for points leftover at the end of the game, which will be forfeited.

Instructions: Summary

For your convenience, we summarize all the information in the tables below. We will present these tables throughout the game to help you make your decisions. When comparing A and B, you will find that

- **A** has a **lower point threshold** for free ticket than **B**
- **B** has a **lower regular** price level than **A**
- **B** has a **lower promotional** price level than **A**

Please try to memorize the above information for your decision making. In the past, we find that this could help increase participants' bonus amount.

Table 1: Price Distribution of the Two Airlines		Prices of B	
		Promotion: \$360	Regular: \$440
Prices of A	Promotion: \$400	15% †	15%
	Regular: \$480	15%	55%

† the chance of seeing a price of \$400 for Airline A and a price \$360 for Airline B at any given trip is 15%.

Table 2: Reward Programs	A	B
Number of Points for One Free Ticket	2000 points	2400 points

Associate Editor

Eric Bradlow


Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iD

Jia Liu  <https://orcid.org/0000-0002-0279-724X>

References

- Airlines.org (2017), "Domestic Round-Trip Fares and Fees," (accessed May 2, 2018), <http://airlines.org/dataset/annual-round-trip-fares-and-fees-domestic/>.
- Autry, Ashley (2017), "Customer Loyalty Statistics: 2017 Edition," Access (December 31), <https://blog.accessdevelopment.com/customer-loyalty-statistics-2017-edition#loyalty>.
- Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak (1964), "Measuring Utility by a Single-Response Sequential Method," *Behavioral Science*, 9 (3), 226–32.
- Caminal, Ramon and Carmen Matutes (1990), "Endogenous Switching Costs in a Duopoly Model," *International Journal of Industrial Organization*, 8 (3), 353–73.
- Ching, Andrew T., Tülin Erdem, and Michael P. Keane (2014), "A Simple Method to Estimate the Roles of Learning, Inventories and

- Category Consideration in Consumer Choice,” *Journal of Choice Modelling*, 13, 60–72.
- Ching, Andrew T. and Masakazu Ishihara (2017), “Identification of Dynamic Models of Rewards Program,” *Japanese Economic Review*, 69 (3), 306–23.
- Ching, Andrew T. and Matthew Osborne (2019), “Identification and Estimation of Forward-Looking Behavior: The Case of Consumer Stockpiling,” *Marketing Science* (forthcoming), <https://ssrn.com/abstract=2594032>.
- Chun, So Yeon and Anton Ovchinnikov (2019), “Strategic Consumers, Revenue Management, and the Design of Loyalty Programs,” *Management Science*, 65 (9), 3949–4450.
- Danaher, Peter J., Laszlo Sajtos, and Tracey S. Danaher (2017), “Dynamic Management of Rewards in Coalition Loyalty Programs,” *technical report, working paper* (May 17).
- Ding, Min (2007), “An Incentive-Aligned Mechanism for Conjoint Analysis,” *Journal of Marketing Research*, 44 (2), 214–23.
- Dowling, Grahame R. and Mark Uncles (1997), “Do Customer Loyalty Programs Really Work?” *Sloan Management Review*, 38 (4), 71–82.
- Ericson, Keith Marzilli and David Laibson (2019), “Intertemporal Choice,” in *Handbook of Behavioral Economics: Foundations and Applications 2* B. Douglas Bernheim, Stefano Della Vigna, and David Laibson, eds. Amsterdam: North-Holland.
- Evanschitzky, Heiner, Balasubramanian Ramaseshan, David M. Woisetschlager, Verena Richelsen, Markus Blut, and Christof Backhaus (2012), “Consequences of Customer Loyalty to the Loyalty Program and to the Company,” *Journal of the Academy of Marketing Science*, 40 (5), 625–38.
- Gardete, Pedro and James Lattin (2018), “Coalition Loyalty Program Not Working? Maybe You’re Doing It Wrong,” Working Paper 3645, Stanford University, <https://www.gsb.stanford.edu/faculty-research/working-papers/coalition-loyalty-program-not-working-maybe-youre-doing-it-wrong>.
- Gelman, Andrew, Xiao-Li Meng, and Hal Stern (1996), “Posterior Predictive Assessment of Model Fitness via Realized Discrepancies,” *Statistica Sinica*, 6 (4), 733–60.
- Geweke, John and Michael P. Keane (2000), “Bayesian Inference for Dynamic Discrete Choice Models Without the Need for Dynamic Programming,” in *Simulation-Based Inference in Econometrics: Methods and Applications*, Roberto S. Mariano, Til Schuermann and Melvyn J. Weeks, eds. Cambridge, UK: Cambridge University Press, 100–131.
- Geweke, John and Michael P. Keane (2001), “Computationally Intensive Methods for Integration in Econometrics,” in *Handbook of Econometrics*, Vol. 5, James J. Heckman and Edward E. Leamer, eds. Amsterdam: Elsevier, 3463–3568.
- Gilbride, Timothy J. and Greg M. Allenby (2004), “A Choice Model with Conjunctive, Disjunctive, and Compensatory Screening Rules,” *Marketing Science*, 23 (3), 391–406.
- Hara, Kotaro, Abigail Adams, Kristy Milland, Saiph Savage, Chris Callison-Burch, and Jeffrey P. Bigham (2018), “A Data-Driven Analysis of Workers’ Earnings on Amazon Mechanical Turk,” in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, Association for Computing Machinery, <https://dl.acm.org/doi/10.1145/3173574.3174023>.
- Hartmann, Wesley R. and V. Brian Viard (2008), “Do Frequency Reward Programs Create Switching Costs? A Dynamic Structural Analysis of Demand in a Reward Program,” *Quantitative Marketing and Economics*, 6 (2), 109–37.
- Houser, Daniel, Michael Keane, and Kevin McCabe (2004), “Behavior in a Dynamic Decision Problem: An Analysis of Experimental Evidence Using a Bayesian Type Classification Algorithm,” *Econometrica*, 72 (3), 781–822.
- Imai, Susumu and Michael P. Keane (2004), “Intertemporal Labor Supply and Human Capital Accumulation,” *International Economic Review*, 45 (2), 601–41.
- Keane, Michael P. and Kenneth I. Wolpin (1994), “The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence,” *Review of Economics and Statistics*, 76 (4), 648–72.
- Kim, Byung-Do, Mengze Shi, and Kannan Srinivasan (2001), “Reward Programs and Tacit Collusion,” *Marketing Science*, 20 (2), 99–120.
- Kivetz, Ran, Oleg Urminsky, and Yuhuang Zheng (2006), “The Goal-Gradient Hypothesis Resurrected: Purchase Acceleration, Illusory Goal Progress, and Customer Retention,” *Journal of Marketing Research*, 43 (1), 39–58.
- Klemperer, Paul (1995), “Competition When Consumers Have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade,” *Review of Economic Studies*, 62 (4), 515–39.
- Kopalle, Praveen K., Yacheng Sun, Scott A. Neslin, Baohong Sun, and Vanitha Swaminathan (2012), “The Joint Sales Impact of Frequency Reward and Customer Tier Components of Loyalty Programs,” *Marketing Science*, 31 (2), 216–35.
- Krusell, Per and Anthony A. Smith Jr. (1996), “Rules of Thumb in Macroeconomic Equilibrium: A Quantitative Analysis,” *Journal of Economic Dynamics and Control*, 20 (4), 527–58.
- Lewis, Michael (2004), “The Influence of Loyalty Programs and Short-Term Promotions on Customer Retention,” *Journal of Marketing Research*, 41 (3), 281–92.
- Liu, Yuping and Rong Yang (2009), “Competing Loyalty Programs: Impact of Market Saturation, Market Share, and Category Expandability,” *Journal of Marketing*, 73 (1), 93–108.
- Metropolis, Nicholas, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller (1953), “Equation of State Calculations by Fast Computing Machines,” *Journal of Chemical Physics*, 21 (6), 1087–92.
- Meyer, Robert J. and J. Wesley Hutchinson (2016), “(When) Are We Dynamically Optimal? A Psychological Field Guide for Marketing Modelers,” *Journal of Marketing*, 80 (5), 20–33.
- Neal, Radford M. (2011), “MCMC Using Hamiltonian Dynamics,” in *Handbook of Markov Chain Monte Carlo*, Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng, eds. New York: Chapman and Hall/CRC, 113–62.
- Nunes, Joseph C. and Xavier Drèze (2006), “The Endowed Progress Effect: How Artificial Advancement Increases Effort,” *Journal of Consumer Research*, 32 (4), 504–12.

- Pancras, Joseph, Rajkumar Venkatesan, and Bin Li (2015), "Investigating the Value of Competitive Mobile Loyalty Program Platforms for Intermediaries and Retailers," *Marketing Science Institute Working Paper Series*, 15–107, <https://www.msi.org/reports/investigating-the-value-of-competitive-mobile-loyalty-program-platforms-for/>.
- Powell, Warren B. (2009), "What You Should Know About Approximate Dynamic Programming," *Naval Research Logistics (NRL)*, 56 (3), 239–49.
- Rossi, Federico (2018), "Lower Price or Higher Reward? Measuring the Effect of Consumers' Preferences on Reward Programs," *Management Science*, 64 (9), 4451–70.
- Seaney, Rick (2013), "Airfare Expert: How Much Does a 'Free' Ticket Really Cost?" *USA Today* (February 5), <https://www.usatoday.com/story/travel/columnist/seaney/2013/02/22/frequent-flier-tickets-fees-surcharges/1891517/>.
- Shu, Suzanne B. and Ayelet Gneezy (2010), "Procrastination of Enjoyable Experiences," *Journal of Marketing Research*, 47 (5), 933–44.
- Singh, Siddharth S., Dipak C. Jain, and Trichy V. Krishnan (2008), "Research Note: Customer Loyalty Programs: Are They Profitable?" *Management Science*, 54 (6), 1205–11.
- Spiegelhalter, David J., Nicola G. Best, Bradley P. Carlin, and Angelika Van Der Linde (2002), "Bayesian Measures of Model Complexity and Fit," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64 (4), 583–639.
- Stourm, Valeria, Eric Bradlow, and Peter Fader (2015), "Stockpiling Points in Linear Loyalty Programs," *Journal of Marketing Research*, 52 (2), 253–67.
- Stourm, Valeria, Eric Bradlow, and Peter Fader (2017), "Market Positioning Using Cross-Reward Effects in a Coalition Loyalty Program," HEC Paris Research Paper No. MKG-2017-1242.
- Sun, Yacheng and Dan Zhang (2019), "A Model of Customer Reward Programs with Finite Expiration Terms," *Management Science*, 65 (8), 3889–3903.
- Wang, Yanwen, Michael Lewis, Cynthia Cryder, and Jim Sprigg (2016), "Enduring Effects of Goal Achievement and Failure Within Customer Loyalty Programs: A Large-Scale Field Experiment," *Marketing Science*, 35 (4), 565–75.
- Zhang, Jie and Els Breugelmans (2012), "The Impact of an Item-Based Loyalty Program on Consumer Purchase Behavior," *Journal of Marketing Research*, 49 (1), 50–65.